

Unbalancing Sets and an Almost Quadratic Lower Bound for Syntactically Multilinear Arithmetic Circuits

Noga Alon¹

Sackler School of Mathematics and Blavatnik School of Computer Science
Tel Aviv, 6997801, Israel, and
Center for Mathematical Sciences and Applications, Harvard University,
Cambridge, MA 02138, USA
nogaa@tau.ac.il

Mrinal Kumar²

Center for Mathematical Sciences and Applications, Harvard University
Cambridge, MA 02138, USA
mrinalkumar08@gmail.com

Ben Lee Volk³

Blavatnik School of Computer Science, Tel Aviv University
Tel Aviv, 6997801, Israel
benleevolk@gmail.com

Abstract

We prove a lower bound of $\Omega(n^2/\log^2 n)$ on the size of any syntactically multilinear arithmetic circuit computing some explicit multilinear polynomial $f(x_1, \dots, x_n)$. Our approach expands and improves upon a result of Raz, Shpilka and Yehudayoff ([31]), who proved a lower bound of $\Omega(n^{4/3}/\log^2 n)$ for the same polynomial. Our improvement follows from an asymptotically optimal lower bound for a generalized version of Galvin's problem in extremal set theory.

2012 ACM Subject Classification Theory of computation → Algebraic complexity theory

Keywords and phrases Algebraic Complexity, Multilinear Circuits, Circuit Lower Bounds

Digital Object Identifier 10.4230/LIPIcs.CCC.2018.11

Related Version <https://arxiv.org/abs/1708.02037>

Acknowledgements Part of this work was done while the second author was visiting Tel Aviv University. We thank Amir Shpilka for the visit, for many insightful discussions, and for comments on an earlier version of this text. We are also thankful to Andy Drucker for pointing out a correction in a previous version of this paper.

1 Introduction

An arithmetic circuit is one of the most natural and standard computational models for computing multivariate polynomials. Such circuits provide a succinct representation of

¹ Research supported in part by an ISF grant and by a GIF grant.

² Part of this work was done while visiting Tel Aviv University.

³ The research leading to these results has received funding from the Israel Science Foundation (grant number 552/16).



multivariate polynomials, and in some sense, they can be thought of as algebraic analogs of boolean circuits. Formally, an arithmetic circuit over a field \mathbb{F} and a set of variables $X = \{x_1, x_2, \dots, x_n\}$ is a directed acyclic graph in which every vertex has in-degree either zero or two. The vertices of in-degree zero (called *leaves*) are labeled by variables in X or elements of \mathbb{F} , and the vertices of in-degree two are labeled by either $+$ (called *sum gates*) or \times (called *product gates*). A circuit can have one or more vertices of out degree zero, known as the output gates. The polynomial computed by a vertex in any⁴ given circuit is naturally defined in an inductive way: a leaf computes the polynomial which is equal to its label. A sum gate computes the polynomial which is the sum of the polynomials computed at its children and a product gate computes the polynomial which is the product of the polynomials at its children. The polynomials computed by a circuit are the polynomials computed by its output gates. The size of an arithmetic circuit is the number of vertices in it.

It is not hard to show (see, e.g., [7]) that a random polynomial of degree $d = \text{poly}(n)$ in n variables cannot be computed by an arithmetic circuit of size $\text{poly}(n)$ with overwhelmingly high probability. A fundamental problem in this area of research is to prove a similar super-polynomial lower bound for an *explicit* polynomial family. Unfortunately, the problem continues to remain wide open and the current best lower bound known for general arithmetic circuits⁵ is an $\Omega(n \log n)$ lower bound due to Strassen [37] and Baur and Strassen [5] from more than three decades ago. The absence of substantial progress on this general question has led to focus on the question of proving better lower bounds for restricted and more structured subclasses of arithmetic circuits. Arithmetic formulas [19], non-commutative arithmetic circuits [26], algebraic branching programs [22], and low depth arithmetic circuits [27, 13, 14, 30, 15, 11, 20, 24, 23] are some such subclasses which have been studied from this perspective. For an overview of the definition of these models and the state of art for lower bounds for them, we refer the reader to the surveys of Shpilka and Yehudayoff [35] and Saptharishi [34].

Several of the most important polynomials in algebraic complexity and in mathematics in general are multilinear. Notable examples include the determinant, the permanent, and the elementary symmetric polynomials. Therefore, one subclass which has received a lot of attention in the last two decades and will be the focus of this paper is the class of *multilinear* arithmetic circuits.

1.1 Multilinear arithmetic circuits

For an arithmetic circuit Ψ and a vertex v in Ψ , we denote by X_v the set of variables x_i such that there is a directed path from a leaf labeled by x_i to v ; in this case, we also say that v *depends* on x_i ⁶. A polynomial P is said to be multilinear if the individual degree of every variable in P is at most one.

An arithmetic circuit Ψ is said to be *syntactically* multilinear if for every multiplication gate v in Ψ with children u and w , the sets of variables X_u and X_w are disjoint. We say that Ψ is *semantically* multilinear if the polynomial computed at every vertex is a multilinear polynomial. Observe that if Ψ is a syntactically multilinear circuit, then it is also semantically multilinear. However, it is not clear if every semantically multilinear circuit can be efficiently simulated by a syntactically multilinear circuit.

⁴ Throughout this paper, we will use the terms gates and vertices interchangeably.

⁵ In the rest of the paper, when we say a lower bound, we always mean it for an explicit polynomial family.

⁶ We remark that this is a syntactic notion of dependency, since it is possible that every monomial with x_i might get canceled in the intermediate computation and might not eventually appear in the polynomial computed at v .

A multilinear circuit is a natural model for computing multilinear polynomials, but it is not necessarily the most efficient one. Indeed, it is remarkable that all the constructions of polynomial size arithmetic circuits for the determinant [8, 6, 25], which are fundamentally different from one another, nevertheless share the property of being *non*-multilinear, namely, they involve non-multilinear intermediate computations which eventually cancel out. There are no subexponential-size multilinear circuits known for the determinant, and one may very well conjecture these do not exist at all.

Multilinear circuits were first studied by Nisan and Wigderson [27]. Subsequently, Raz [29] defined the notion of multilinear formulas⁷ and showed that any multilinear formula computing the determinant or the permanent of an $n \times n$ variable matrix must have super-polynomial size. In a follow up work [28], Raz further strengthened the results in [29] and showed that there is a family of multilinear polynomials in n variables which can be computed by a $\text{poly}(n)$ size syntactically multilinear arithmetic circuits but require multilinear formulas of size $n^{\Omega(\log n)}$.

Building on the ideas and techniques developed in [29], Raz and Yehudayoff [33] showed an exponential lower bound for syntactically multilinear circuits of constant depth. Interestingly, they also showed a super-polynomial separation between depth Δ and depth $\Delta+1$ syntactically multilinear circuits for constant Δ .

In spite of the aforementioned progress on the question of lower bounds for multilinear formulas and bounded depth syntactically multilinear circuits, there was no $\Omega(n^{1+\varepsilon})$ lower bounds known for general syntactically multilinear circuits for any constant $\varepsilon > 0$. In fact, the results in [28] show that the main technical idea underlying the results in [29, 28, 33] is unlikely to directly give a super-polynomial lower bound for general syntactically multilinear circuits. However, a weaker super-linear lower bound still seemed conceivable via similar techniques.

Raz, Shpilka and Yehudayoff [31] showed that this is indeed the case. By a sophisticated and careful application of the techniques in [29] along with several additional ideas, they established an $\Omega\left(\frac{n^{4/3}}{\log^2 n}\right)$ lower bound for an explicit n variate polynomial. Since then, this has remained the best lower bound known for syntactically multilinear circuits. In this paper, we improve this result by showing an almost quadratic lower bound for syntactically multilinear circuits for an explicit n variate polynomial. In fact, the family of hard polynomials in this paper is the same as the one used in [31]. We now formally state our result.

► **Theorem 1.** *There is an explicit family of polynomials $\{f_n\}$, where f_n is an n variate multilinear polynomial, such that any syntactically multilinear arithmetic circuit computing f_n must have size at least $\Omega(n^2/\log^2 n)$.*

For our proof, we follow the strategy in [31]. Our improvement comes from an improvement in a key lemma in [31] which addresses the following combinatorial problem.

► **Question 2.** *What is the minimal integer $m = m(n)$ for which there is a family of subsets $S_1, S_2, \dots, S_m \subseteq [n]$, each S_i satisfying $6 \log n \leq |S_i| \leq n - 6 \log n$ such that for every $T \subseteq [n]$, $|T| = \lfloor n/2 \rfloor$, there exists an $i \in [m]$ with $|T \cap S_i| \in \{\lfloor |S_i|/2 \rfloor - 3 \log n, \lfloor |S_i|/2 \rfloor - 3 \log n + 1, \dots, \lfloor |S_i|/2 \rfloor + 3 \log n\}$?*

Raz, Shpilka and Yehudayoff [31] showed that $m(n) \geq \Omega(n^{1/3}/\log n)$. For our proof, we show that $m(n) \geq \Omega(n/\log n)$.

⁷ For formulas, it is known that syntactic multilinearity and semantically multilinearity are equivalent (See, e.g., [29]).

In addition to its application to the proof of Theorem 1, Question 2 seems to be a natural problem in extremal combinatorics and might be of independent interest, and special cases thereof were studied in the combinatorics literature. In the next section, we briefly discuss the state of the art of this question and state our main technical result about it in Theorem 3.

1.2 Unbalancing Sets

The following question, which is of very similar nature to Question 2, is known as Galvin's problem (see [12, 9]): What is the minimal integer $m = m(n)$, for which there exists a family of subsets $S_1, \dots, S_m \subseteq [4n]$, each of size $2n$, such that for every subset $T \subseteq [4n]$ of size $2n$ there exists some $i \in [m]$ such that $|T \cap S_i| = n$?

It is not hard to show that $m(n) \leq 2n$. Indeed, let $S_i = \{i, i+1, \dots, i+2n-1\}$, for $i \in \{1, 2, \dots, 2n+1\}$, and let $\alpha_i(T) = |T \cap S_i| - |([4n] \setminus T) \cap S_i|$. Then $\alpha_i(T)$ is always an even integer, $\alpha_1(T) = -\alpha_{2n+1}(T)$, and $\alpha_i - \alpha_{i+1}(T) \in \{0, \pm 2\}$ if $i \leq 2n$. By a discrete version of the intermediate value theorem, it follows there exists $j \in [2n]$ such that $\alpha_j(T) = 0$, which implies that exactly n elements of S_j belong to T . Thus, the family $\{S_1, \dots, S_{2n}\}$ satisfies this property.

As for lower bounds, a counting argument shows that $m(n) = \Omega(\sqrt{n})$, since for each fixed S of size $[2n]$ and random T of size $2n$,

$$\Pr[|T \cap S| = n] = \frac{\binom{2n}{n} \cdot \binom{2n}{n}}{\binom{4n}{2n}} = \Theta\left(\frac{1}{\sqrt{n}}\right).$$

Frankl and Rödl [12] were able to show that $m(n) \geq \varepsilon n$ for some $\varepsilon > 0$ if n is odd, and Enomoto, Frankl, Ito and Nomura [9] proved that $m(n) \geq 2n$ if n is even, which implies that even the constant in the construction given above is optimal. Until this work, the question was still open for even values of n : in fact, Markert and West (unpublished, see [9]) showed that for $n \in \{2, 4\}$, $m(n) < 2n$.

For our purposes, we need to generalize Galvin's problem in two ways. The first is to lift the restriction on the set sizes. The second is to ask how small can the size of the family $\mathcal{F} = \{S_1, \dots, S_m\} \subseteq 2^{[n]}$ be if we merely assume each balanced partition T is " τ -balanced" on some $S \in \mathcal{F}$, namely, if $||T \cap S| - |S|/2|| \leq \tau$ for some S (the main case of interest for us is $\tau = O(\log n)$). Of course, since T itself is balanced, very small or very large sets are always τ -balanced, and thus we impose the (tight) non-triviality condition $2\tau \leq |S| \leq n - 2\tau$ for every $S \in \mathcal{F}$.

Once again, by defining $S_i = \{i, i+1, \dots, i+n/2-1\}$ (n is always assumed to be even), the family $\mathcal{F} = \{S_1, S_{1+\tau}, S_{1+2\tau}, \dots, S_{1+\lfloor n/(2\tau) \rfloor \cdot \tau}\}$ gives a construction of size $O(n/\tau)$ such that every balanced partition T is τ -balanced on some $S \in \mathcal{F}$.

It is natural to conjecture that, perhaps up to a constant, this construction is optimal. Indeed, this is what we prove here.

► **Theorem 3.** *Let n be any large enough even number, and let $\tau \geq 1$ be an integer. Let $S_1, \dots, S_m \subseteq [n]$ be sets such that for all $i \in [m]$, $2\tau \leq |S_i| \leq n - 2\tau$. Further, assume that for every $Y \subseteq [n]$ of size $n/2$ there exists $i \in [m]$ such that $||Y \cap S_i| - |S_i|/2| < \tau$. Then, $m \geq \Omega(n/\tau)$.*

In particular, Theorem 3 proves a linear lower bound $m = \Omega(n)$ for the original problem of Galvin, even when the universe size is of the form $4k$ for even k .

We remark that the relevance of problems of this form to lower bounds in algebraic complexity was also observed by Jansen [18] who considered the problem of obtaining a

lower bound on homogenous syntactically multilinear algebraic branching program (which is a weaker model than syntactically multilinear circuits), and essentially proposed Theorem 3 as a conjecture. In fact, a special case of this theorem (see Theorem 9), which has a simpler proof, is already enough to derive the improved lower bounds for syntactically multilinear circuits.

Alon, Bergmann, Coppersmith and Odlyzko [1] considered a very similar problem of balancing ± 1 -vectors: they studied families of vectors $\mathcal{F} = \{v_1, \dots, v_m\}$ such that $v_i \in \{\pm 1\}^n$ for $i \in [m]$, which satisfy the properties that for every $w \in \{\pm 1\}^n$ (not necessarily balanced), there exists $i \in [m]$ such that $|\langle v_i, w \rangle| \leq d$. They generalized a construction of Knuth [21] and proved a matching lower bound which together showed that $m = \lceil n/(d+1) \rceil$ is both necessary and sufficient for such a set to exist. Galvin's problem seems like "the $\{0, 1\}$ version" of the same problem, but, to quote from [1], there does not seem to be any simple dependence between the problems.

1.3 Proof overview

In this section, we discuss the main ideas and give a brief sketch of the proofs of Theorem 1 and Theorem 3. Since our proof heavily depends on the proof in [31] and follows the same strategy, we start by revisiting the main steps in their proof and noting the key differences between the proof in [31] and our proof. We also outline the reduction to the combinatorial problem of unbalancing set families in Question 2.

Proof sketch of [31]

The proof in [31] starts by proving a syntactically multilinear analog of a classical result of Baur and Strassen [5], where it was shown that if an n variate polynomial f is computable by an arithmetic circuit Ψ of size $s(n)$, then there is an arithmetic circuit Ψ' of size at most $5s(n)$ with n outputs such that the i -th output gate of Ψ' computes $f_i = \frac{\partial f}{\partial x_i}$. Raz, Shpilka and Yehudayoff show that if Ψ is syntactically multilinear, then the circuit Ψ' continues to be syntactically multilinear. Additionally, there is no directed path from a leaf labeled by x_i to the output gate computing f_i .⁸

Once we have this structural result, it would suffice to prove a lower bound on the size of Ψ' . For brevity, we denote the subcircuit of Ψ' rooted at the output gate computing f_i by Ψ'_i . As a key step of the proof in [31], the authors identify certain sets of vertices $\mathcal{U}_1, \mathcal{U}_2, \dots, \mathcal{U}_n$ in Ψ' with the following properties.

- For every $i \in [n]$, \mathcal{U}_i is a subset of vertices in Ψ'_i .
- For every $i \in [n]$ and $v \in \mathcal{U}_i$, the number of $j \neq i$ such that $v \in \mathcal{U}_j$ is not too large (at most $O(\log n)$).

Observe that at this point, showing a lower bound of $s'(n)$ on the size of each \mathcal{U}_i implies a lower bound of $\Omega(ns'(n)/\log n)$ on the size of Ψ' and hence Ψ . In [31], the authors show that there is an explicit f such that each \mathcal{U}_i must have size at least $\Omega(n^{1/3}/\log n)$, thereby getting a lower bound of $\Omega(n^{4/3}/\log^2 n)$ on the size of Ψ .

For our proof, we follow precisely this high level strategy. Our improvement in the lower bound comes from showing that each \mathcal{U}_i must be of size at least $\Omega(n/\log n)$ and not just $\Omega(n^{1/3}/\log n)$ as shown in [31]. We now elaborate further on the main ideas in this step in [31] and the differences with the proofs in this paper.

⁸ See Theorem 15 for a formal statement.

We start with some intuition into the definition of the sets \mathcal{U}_i in [31]. Consider a vertex v in Ψ' which depends on at least k variables. Without loss of generality, let these variables be $\{x_1, x_2, \dots, x_k\}$. From item 4 in Theorem 15, we know that the variable x_i does not appear in the subcircuit Ψ'_i . Therefore, the vertex v cannot appear in the subcircuits $\Psi'_1, \Psi'_2, \dots, \Psi'_k$. So, if we define the set \mathcal{U}_i as the set of vertices in Ψ'_i which depend on at least k variables, then \mathcal{U}_i must be disjoint from vertices in at least k of the subcircuits $\Psi'_1, \Psi'_2, \dots, \Psi'_n$. Picking $k \geq n - O(\log n)$ would give us the desired property. So, if we can prove a lower bound on the size of the set \mathcal{U}_i , we would be done. However, the definition of the set \mathcal{U}_i so far turns out to be too general: indeed, it is not even a priori clear that the \mathcal{U}_i has any other gates apart from the output gate of Ψ'_i .

As is often the case, the solution to this obstacle is to prove a stronger claim by imposing additional structure on the set \mathcal{U}_i . In [31], the set \mathcal{U}_i (called the *upper leveled* gates in Ψ'_i) is defined as the set of all vertices in Ψ'_i which depend on at least $n - 6 \log n$ variables and have a child which depends on more than $6 \log n$ variables and less than $n - 6 \log n$ variables. This additional structure is helpful in proving a lower bound on the size of \mathcal{U}_i . We now discuss this in some more detail.

For every $i \in [n]$, let \mathcal{L}_i be the set of vertices u in Ψ'_i , such that $6 \log n < |X_u| < n - 6 \log n$, and u has a parent in \mathcal{U}_i . These gates are referred to as *lower leveled* gates. Observe that $|\mathcal{L}_i| \geq \frac{|\mathcal{U}_i|}{2}$, since the in-degree of every vertex in Ψ'_i is at most 2. The key structural property of the set \mathcal{L}_i is the following (see Proposition 5.5 in [31]).

► **Lemma 4** ([31]). *Let $i \in [n]$, and let h_1, h_2, \dots, h_ℓ be the polynomials computed by the gates in \mathcal{L}_i . Then, there exist multilinear polynomials $g_1, g_2, \dots, g_\ell, g$ such that*

$$f_i = \sum_{j \in [\ell]} g_j \cdot h_j + g \tag{1}$$

where

- For every $j \in [\ell]$, h_j and g_j are variable disjoint.
- The degree of g is at most $O(\log n)$.

Observe that (1) is basically a decomposition of a potentially-hard polynomial f_i in terms of the sum of products of multilinear polynomials in an intermediate number of variables. The goal is to show that for an appropriate explicit f_i , the number of summands on the right hand side of (1) cannot be too small. A similar scenario also appears in the multilinear formula lower bounds and bounded depth multilinear formula lower bounds of [29, 28, 33] (albeit with some key differences). Hence, a natural approach at this point would be to use the tools in [29, 28, 33], namely the rank of the *partial derivative matrix*, to attempt to prove this lower bound. We refer the reader to Section 2.2 for the definitions and properties of the partial derivative matrix and proceed with the overview. For each $j \in [\ell]$, let the polynomial h_j in Lemma 4 depend on the variables $S_j \subseteq X$. The key technical step in the rest of the proof is to show that there is a partition of the set of variables $X = \{x_1, x_2, \dots, x_n\}$ into Y and Z such that $|Y| = |Z|$ and for every $j \in [\ell]$, $||S_j \cap Y| - |S_j \cap Z|| \geq \Omega(\log n)$. In [31], the authors show that there is an absolute constant $\varepsilon > 0$ such that if $\ell \leq \varepsilon n^{1/3} / \log n$, then there is an equipartition of X which *unbalances* all the sets $\{S_j : j \in [\ell]\}$ by at least $\Omega(\log n)$. Our key technical contribution (Theorem 3) in this paper is to show that as long as $\ell \leq \varepsilon n / \log n$, there is an equipartition which unbalances all the S_j 's by at least $\Omega(\log n)$. This implies an $\Omega(n / \log n)$ on the size of each set \mathcal{U}_i , and thus an $\Omega(n^2 / \log^2 n)$ lower bound on the circuit size.

Before we dive into a more detailed discussion on the overview and main ideas in the proof of Theorem 3 in the next section, we would like to remark that the lower bound question in (1) seems to be a trickier question than what is encountered while proving

multilinear formula lower bounds [29, 28] or bounded depth syntactically multilinear circuit lower bounds [33]. The main differences are that in the proofs in [29, 28, 33], the sets S_j have a stronger guarantee on their size (at least $n^{\Omega(1)}$ and at most $n - n^{\Omega(1)}$), and each of the summands on the right has *many* variable disjoint factors and not just two factors as in (1). For instance, in the formula lower bound proofs the number of variable disjoint factors in each summand on the right is $\Omega(\log n)$, and for constant depth circuit lower bounds it is $n^{\Omega(1)}$. Together, these properties make it possible to show much stronger lower bounds on ℓ . In particular, it is known that a *random* equipartition works for these two applications, in the sense that it unbalances sufficiently many factors in each summand, thereby implying that the rank of the partial derivative matrix of the polynomial is small. Hence, for an appropriate⁹ f_i , the number of summands must be large. However, since a set of size $O(\log n)$ is balanced under a random equipartition with probability $\Omega(1/\sqrt{\log n})$ and the identity in (1) involves just two variable disjoint factors, taking a random equipartition would not enable us to prove any meaningful bounds.

Proof sketch of Theorem 3

Recall that our task is, given a small collection of subsets of $[n]$, to find a balanced partition which is unbalanced on each of the sets. Equivalently, we would like to prove that if \mathcal{F} is a family of subsets such that every balanced partition balances at least one set in \mathcal{F} , then $|\mathcal{F}|$ must be large (of course, \mathcal{F} must satisfy the conditions in Theorem 3).

We first sketch the proof of a special case (which suffices for the main application here), when $n = 4p$ and p is a prime. For the sake of simplicity, suppose also that all subsets $S \in \mathcal{F}$ are of even size, and assume further that for every subset $T \subseteq [n]$ of size $n/2$ there exists $S \in \mathcal{F}$ such that T completely balances S , namely, $|T \cap S| = |S|/2$. One possible approach to obtain lower bounds on $|\mathcal{F}|$ is via an application of the polynomial method as done, for example, in [1]. Define the following polynomial over, say, the rationals:

$$f(x_1, \dots, x_n) = \prod_{S \in \mathcal{F}} (\langle x, \mathbf{1}_S \rangle - |S|/2).$$

By the assumption on \mathcal{F} , the polynomial f evaluates to 0 over all points in $\{0, 1\}^n$ with Hamming weight exactly $n/2$. We can also argue, using the assumption on the set sizes in \mathcal{F} , that f is not identically zero, and clearly $\deg(f) \leq |\mathcal{F}|$. Thus, a lower bound on $\deg(f)$ translates to a lower bound on $|\mathcal{F}|$.

This idea, however, seems like a complete nonstarter, since there exists a degree 1 non-zero polynomial which evaluates to 0 over the middle layer of $\{0, 1\}^n$, namely, $\sum_i x_i - n/2$.

A very clever solution to this potential obstacle was found by Hegedűs [16]. Suppose $n = 4p$ for some prime p . The main insight in [16] is to consider the polynomial f over \mathbb{F}_p , and to add the requirement that there exists some $z \in \{0, 1\}^{4p}$, of Hamming weight *exactly* $3p$, such that $f(z) \neq 0$. This requirement rules out the trivial example $\sum_i x_i - n/2$, and Hegedűs was able to show that the degree of any polynomial with these properties must be at least $p = n/4$ (see Lemma 5 for the complete statement).

We are thus left with the task of proving that our polynomial evaluates to a non-zero value over some point $z \in \{0, 1\}^{4p}$ of Hamming weight $3p$. This turns out to be not very hard to show, assuming each set is of size at least, say, $100 \log n$ and at most $n - 100 \log n$,

⁹ f_i is chosen so that the the partial derivative matrix for f_i is of full rank for *every* equipartition.

by choosing a random such vector z . Indeed, it is not surprising that it is much easier to directly show that a highly unbalanced partition of $[n]$ (into $3n/4$ vs $n/4$) unbalances all the sets \mathcal{F} .¹⁰

As mentioned earlier, the case $n = 4p$ and $\tau \geq 100 \log n$ in Theorem 3 is considerably easier to prove and suffices for the application to circuit lower bounds. Proving this theorem for every even n and every $\tau \geq 1$ requires further technical ideas which appear in the full version of this paper [2].

Even though Lemma 5 seems to be a fundamental statement about polynomials over finite fields and could conceivably have an elementary proof, the proof in [16] uses more advanced techniques. It relies on the description of Gröbner basis for ideals of polynomials in $\mathbb{F}[x_1, x_2, \dots, x_n]$ which vanish on all points in $\{0, 1\}^n$ of weight equal to $n/2$. A complete description of the reduced Gröbner basis for such ideals was given by Hegedűs and Rónyai [17] and their proof builds up on a number of earlier partial results [4, 10] on this problem.

To the best of our knowledge, the proof in [16] is the only known proof of Lemma 5, and giving a self contained elementary proof of it seems to be an interesting question.

Organization of the paper

In the rest of the paper, we set up some notation and discuss some preliminary notions in Section 2, prove Theorem 3 in Section 3 and complete the proof of Theorem 1 in Section 4. Throughout the paper we assume, whenever this is needed, that n is sufficiently large, and make no attempts to optimize the absolute constants.

2 Preliminaries

For $n \in \mathbb{N}$, we denote $[n] = \{1, 2, \dots, n\}$. For a prime p , we denote by \mathbb{F}_p the finite field with p elements. For two integers i, j with $i \leq j$, we denote $[i, j] = \{a \in \mathbb{Z} : i \leq a \leq j\}$. The characteristic vector of a set $S \subseteq [n]$ is denoted by $\mathbf{1}_S \in \{0, 1\}^n$.

As is standard, $\binom{[n]}{k}$ denotes the family $\{S \subseteq [n] : |S| = k\}$.

For an even $n \in \mathbb{N}$ and $Y \subseteq [n]$ such that $|Y| = n/2$, we call Y a *balanced partition* of $[n]$, with the implied meaning that Y partitions $[n]$ evenly into Y and $[n] \setminus Y$. The *imbalance* of a set $S \subseteq [n]$ under Y is $d_Y(S) := ||Y \cap S| - |S|/2|$. Observe the useful symmetry $d_Y(S) = d_Y([n] \setminus [S])$, which follows from the fact that $|Y| = n/2$. We say S is τ -unbalanced under Y if $d_Y(S) \geq \tau$.

We use the following lemma from [16].

► **Lemma 5** ([16]). *Let p be a prime, and let $f \in \mathbb{F}_p[x_1, \dots, x_{4p}]$ be a polynomial. Suppose that for all $Y \in \binom{[4p]}{2p}$, it holds that $f(\mathbf{1}_Y) = 0$, and that there exists $T \subseteq [4p]$ such that $|T| = 3p$ and $f(\mathbf{1}_T) \neq 0$. Then $\deg(f) \geq p$.*

2.1 Hypergeometric distribution

For parameters N, M, k , where $N \geq M$, by $\mathcal{H}(M, N, k)$, we denote the distribution of $|S \cap T|$, where S is any fixed subset of $[N]$ of size M , and T is a uniformly random subset of $[N]$ of

¹⁰In our case, we need to argue that the imbalance is non-zero modulo p , which adds an extra layer of complication, although again, one which is not hard to solve.

size equal to k . Clearly,

$$\Pr[|S \cap T| = i] = \frac{\binom{M}{i} \binom{N-M}{k-i}}{\binom{N}{k}}.$$

The expected value of $|S \cap T|$ under this distribution is equal to kM/N . We need the following tail bound of hypergeometric distribution for our proof.

► **Lemma 6** ([36]). *Let N, M, k , and $\mathcal{H}(M, N, k)$ be as defined above. Then, for every t*

$$\Pr[||S \cap T| - kM/N| \geq tk] \leq e^{-2t^2k}.$$

► **Lemma 7** (Hoeffding's inequality, [3]). *Let X_1, X_2, \dots, X_n be independent random variables taking values in $\{0, 1\}$. Then,*

$$\Pr \left[\left| \sum_{i=1}^n X_i - \mathbb{E} \left[\sum_{i=1}^n X_i \right] \right| \geq t \right] \leq 2 \exp(-2t^2/n).$$

2.2 Partial derivative matrix

For a circuit Ψ , we denote by $|\Psi|$ the size of Ψ , namely, the number of gates in it. For a gate v , we denote by X_v the set of variables that occur in the subcircuit rooted at v .

Let $X = \{x_1, \dots, x_n\}$ be a set of variables, $Y \subseteq X$ (not necessarily of size $n/2$) and let $Z = X \setminus Y$. For a multilinear polynomial $f(X) \in \mathbb{F}[X]$, we define the *partial derivative matrix* of f with respect to Y, Z , denoted $M_{Y,Z}(f)$, as follows: the rows of M are indexed by multilinear monomials in Y . the columns of M are indexed by multilinear monomials in Z . The entry which corresponds to (m_1, m_2) is the coefficient of the monomial $m_1 \cdot m_2$ in f . We define $\text{rank}_{Y,Z}(f) = \text{rank}(M_{Y,Z}(f))$.

The following properties of the partial derivative matrix are easy to prove and well-documented (see, e.g., [31]).

► **Proposition 8.** *The following properties hold:*

1. *For every multilinear polynomial $f(X) \in \mathbb{F}[X]$, $Y \subseteq X$ and $Z = X \setminus Y$, $\text{rank}_{Y,Z}(f) \leq \min \{2^{|Y|}, 2^{|Z|}\}$.*
2. *For every two multilinear polynomials $f_1(X), f_2(X) \in \mathbb{F}[X]$ and for every partition $X = Y \sqcup Z$, $\text{rank}_{Y,Z}(f_1 + f_2) \leq \text{rank}_{Y,Z}(f_1) + \text{rank}_{Y,Z}(f_2)$.*
3. *Let $f_1 \in \mathbb{F}[X_1]$ and $f_2 \in \mathbb{F}[X_2]$ be multilinear polynomials such that $X_1 \cap X_2 = \emptyset$. Let $Y_i \subseteq X_i$ and $Z_i = X_i \setminus Y_i$ for $i \in \{1, 2\}$. Set $Y = Y_1 \cup Y_2, Z = Z_1 \cup Z_2$. Then $\text{rank}_{Y,Z}(f_1 \cdot f_2) = \text{rank}_{Y_1,Z_1}(f_1) \cdot \text{rank}_{Y_2,Z_2}(f_2)$.*
4. *Let $f(X) \in \mathbb{F}[X]$ be a multilinear polynomial such that $X = Y \sqcup Z$ and $|Y| = |Z| = n/2$. Suppose $\text{rank}_{Y,Z}(f) = 2^{n/2}$, and let $g = \partial f / \partial x$ for some $x \in X$. Then $\text{rank}_{Y,Z}(g) = 2^{n/2-1}$.*
5. *Let $f(X) \in \mathbb{F}[X]$ be a multilinear polynomial of total degree d . Then for every partition $X = Y \sqcup Z$ such that $|Y| = |Z| = n/2$, $\text{rank}_{Y,Z}(f) \leq 2^{(d+1) \log(n/2)}$.*

3 Unbalancing sets under a balanced partition

In this section, we prove Theorem 3. We start by proving a special case (see Theorem 9 below) when n equals $4p$ for some prime p , and $\tau \geq \Omega(\log n)$. This special case already suffices for the application to the proof of Theorem 1 (for infinitely many values of n), and has a somewhat simpler proof. We then move on to prove the case for general n and τ , which while being similar to the proof of Theorem 9, needs some additional ideas and care.

3.1 Special case: $n = 4p$ and $\tau \geq \Omega(\log n)$

► **Theorem 9.** *Let p be a large enough prime, and let $\log p \leq \tau \leq p/1000$. Let $S_1, \dots, S_m \subseteq [4p]$ be sets such that for all $i \in [m]$, $100\tau \leq |S_i| \leq 4p - 100\tau$. Further, assume that for every balanced partition Y of $[4p]$ there exists $i \in [m]$ such that $d_Y(S_i) < \tau$. Then, $m \geq \frac{1}{2} \cdot p/\tau$.*

We start with the following lemma, which shows that a small collection of sets can be unbalanced (modulo p) by a partition which is very unbalanced.

► **Lemma 10.** *Let p be a large enough prime, and let $\log p \leq \tau \leq p/1000$. Let $S_1, \dots, S_m \subseteq [4p]$ be sets such that for all $i \in [m]$, $100\tau \leq |S_i| \leq 2p$. Assume further $m \leq p$. Then, there exists $T \subseteq [4p]$, $|T| = 3p$ such that for all $i \in [m]$ and for all $-\tau + 1 \leq t \leq \tau$, $|S_i \cap T| \not\equiv \lfloor |S_i|/2 \rfloor + t \pmod{p}$.*

To prove Lemma 10, we use the following two technical claims. Let $\mu_{3/4}$ denote the probability distribution on subsets of $[4p]$ obtained by putting each $j \in [4p]$ in T with probability $3/4$, independently of all other elements.

► **Claim 11.** *For a random set $T \sim \mu_{3/4}$, $\Pr[|T| = 3p] = \Theta(1/\sqrt{p})$.*

Proof. The probability that $|T| = 3p$ is given by $\binom{4p}{3p} \cdot (3/4)^{3p} \cdot (1/4)^p$, which is $\Theta(1/\sqrt{p})$, by Stirling's approximation. ◀

► **Claim 12.** *Let $\log p \leq \tau \leq p/1000$ and let $S \subseteq [4p]$ such that $100\tau \leq |S| \leq 2p$. For a random set $T \sim \mu_{3/4}$, the probability that for some integer $-\tau + 1 \leq t \leq \tau$ it holds that $|T \cap S| = \lfloor |S|/2 \rfloor + t \pmod{p}$ is at most $1/p^5$.*

Proof. Denote $s = |S|$. Then $\mathbb{E}[|T \cap S|] = 3s/4$. We say T is bad for S if $|T \cap S| = \lfloor s/2 \rfloor + t + kp$ for some $-\tau \leq t \leq \tau + 1$ and $k \in \mathbb{Z}$. We claim this in particular implies that $||T \cap S| - 3s/4| \geq s/5$. Indeed, since $|T \cap S|$ is an integer in the interval $[0, 2p]$, and by the bounds on s , the only cases needed to be analyzed are $k = 0, \pm 1$.

If $|T \cap S| = \lfloor s/2 \rfloor + t - p$, then clearly $|T \cap S| \leq \lfloor s/2 \rfloor$ which implies the statement.

If $|T \cap S| = \lfloor s/2 \rfloor + t + p$, then, as $s \leq 2p$ and $\tau \leq s/100$,

$$|T \cap S| - 3s/4 \geq -s/4 - 1 + t + p \geq p/2 + t - 1 \geq s/4 + t - 1 \geq s/5$$

(The “ -1 ” accounts for the fact that $s/2$ might not be an integer).

Finally, if $|T \cap S| = \lfloor s/2 \rfloor + t$, it holds that

$$|T \cap S| \leq s/2 + \tau \leq s/2 + 2s/100,$$

which again implies the statement.

By Chernoff Bound (see, e.g., [3]), $\Pr[||T \cap S| - 3s/4| \geq s/5] \leq 2^{-|S|/20} \leq 1/p^5$, hence T is bad for S with at most that probability. ◀

The proof of Lemma 10 is now fairly immediate.

Proof of Lemma 10. Pick $T \sim \mu_{3/4}$. By Claim 11, $|T| = 3p$ with probability $\Theta(1/\sqrt{p})$. Recall that T is bad for S_i if $|T \cap S_i| = \lfloor |S_i|/2 \rfloor + t \pmod{p}$ for $t \in \{-\tau + 1, \dots, \tau\}$. By Claim 11, for each S_i , T is bad for S_i with probability at most $1/p^5$. Hence, the probability that there exists $i \in [m]$ such that T is bad for S_i is at most $m/p^5 \leq 1/p^4$.

It follows that with probability at most $1 - \Theta(1/\sqrt{p}) + 1/p^4 < 1$, either $|T| \neq 3p$ or T is bad for some S_i , and hence there exists a selection of T such that $|T| = 3p$ and T is good for all S_i 's. ◀

We are now ready to prove Theorem 9.

Proof of Theorem 9. Let S_1, \dots, S_m be a collection of sets as stated in the theorem. Since $d_Y(S_j) = d_Y([n] \setminus S_j)$, we can assume without loss of generality, by possibly replacing a set with its complement, that $|S_j| \leq 2p$ for all $j \in [m]$. We may further assume $m \leq p$ as otherwise the statement directly follows. For $j \in [m]$, define the following polynomials over \mathbb{F}_p :

$$B_j(x_1, \dots, x_{4p}) = \prod_{t=-\tau+1}^{\tau} (\langle x, \mathbf{1}_{S_j} \rangle - \lfloor |S_j|/2 \rfloor - t),$$

where $x = (x_1, \dots, x_{4p})$ and $\langle u, v \rangle = \sum u_i v_i$ is the usual inner product. Further, define

$$f(x_1, \dots, x_{4p}) = \prod_{j=1}^m B_j(x_1, \dots, x_{4p}),$$

as a polynomial over \mathbb{F}_p .

By assumption, for every $Y \in \binom{[4p]}{2p}$, $f(\mathbf{1}_Y) = 0$. This follows because $\langle \mathbf{1}_Y, \mathbf{1}_{S_j} \rangle = |Y \cap S_j|$, and by assumption, for some j it holds that $d_Y(S_j) < \tau$, so it must be that $|Y \cap S_j| - \lfloor |S_j|/2 \rfloor \in \{-\tau + 1, \dots, 0, \dots, \tau\}$, so that $B_j(\mathbf{1}_Y) = 0$.

Furthermore, Lemma 10 guarantees the existence of a set $T \in \binom{[4p]}{3p}$ such that $f(\mathbf{1}_T) \neq 0$, as the set T from Lemma 10 satisfies the property that $(\langle \mathbf{1}_T, \mathbf{1}_{S_j} \rangle - \lfloor |S_j|/2 \rfloor - t) \neq 0 \pmod p$ for all $-\tau + 1 \leq t \leq \tau$ and for all $j \in [m]$.

By Lemma 5, $\deg(f) \geq p$, and by construction, $\deg(f) \leq 2\tau \cdot m$, which implies the desired lower bound on m . \blacktriangleleft

In the full version of the paper we extend Theorem 9 for a more general range of parameters, by proving the following.

► Theorem 13. *Let n be a large enough even natural number, and let $\tau \in \{1, 2, \dots, n/10^6\}$ be a parameter. Let $S_1, S_2, \dots, S_m \subseteq [n]$ be sets such that for each $i \in [m]$, $2\tau \leq |S_i| \leq n - 2\tau$. Furthermore, assume that for every balanced partition Y of $[n]$, there exists an i such that $d_Y(S_i) < \tau$. Then, $m \geq \frac{1}{10^5} \cdot n/\tau$.*

The proof of Theorem 13 appears in the full version of the paper [2]. We remark that Theorem 9 suffices for the application to circuit lower bounds.

4 Syntactically Multilinear Arithmetic Circuits

In this section, for the sake of completeness, we review the arguments of Raz, Shpilka and Yehudayoff [31], and show how Theorem 9 implies a lower bound of $\Omega(n^2/\log^2 n)$. We mostly refer for [31] for the proofs.

Specifically, we will show the following.

► Theorem 14. *Let n be an even integer, and $X = \{x_1, \dots, x_n\}$. Let $f(X) \in \mathbb{F}[X]$ be a multilinear polynomial such that for every balanced partition $X = Y \sqcup Z$, $\text{rank}_{Y,Z}(f) = 2^{n/2}$. Let Ψ be a syntactically multilinear circuit computing f . Then $|\Psi| = \Omega(n^2/\log^2 n)$.*

The first step in proof of Theorem 14 is to show that if f is computed by a syntactically multilinear circuit of size s , then there exists a syntactically multilinear circuit of size $O(s)$ that computes all the first-order partial derivatives of f , with the additional important property that for each i , the variable x_i does not appear in the subcircuit rooted at the output gate which computes $\partial f/\partial x_i$.

11:12 Unbalancing Sets and Lower Bounds for Multilinear Arithmetic Circuits

► **Theorem 15** ([31], Theorem 3.1). *Let Ψ be a syntactically multilinear circuit over a field \mathbb{F} and the set of variables $X = \{x_1, \dots, x_n\}$. Then, there exists a syntactically multilinear circuit Ψ' , over \mathbb{F} and X , such that:*

1. Ψ' computes all n first-order partial derivatives $\partial f / \partial x_i$, $i \in [n]$.
2. $|\Psi'| \leq 5|\Psi|$.
3. Ψ' is syntactically multilinear.
4. For every $i \in [n]$, $x_i \notin X_{v_i}$, where v_i is the gate in Ψ' computing $\partial f / \partial x_i$.

In particular, if v is a gate in Ψ' , then it is connected by a directed path to at most $n - |X_v|$ output gates.

The proof of Theorem 15 appears in [31], and mostly follows the classical proof of Baur and Strassen [5] of the analogous result for general circuits, with additional care in order to guarantee the last two properties.

Next we define two types of gates in a syntactically multilinear arithmetic circuits.

► **Definition 16.** Let Φ be a syntactically multilinear arithmetic circuit. Define $\mathcal{L}(\Phi, k)$, the set of lower-leveled gates in Φ , by

$$\mathcal{L}(\Phi, k) = \{u : u \text{ is a gate in } \Phi, k < |X_u| < n - k, \text{ and } u \text{ has a parent } v \text{ with } |X_v| \geq n - k\}.$$

Define $\mathcal{U}(\Phi, k)$, the set of upper-leveled gates in Φ , by

$$\mathcal{U}(\Phi, k) = \{v : v \text{ is a gate in } \Phi, |X_v| \geq n - k, \text{ and } v \text{ has a child } u \in \mathcal{L}(\Phi, k)\}.$$

The following lemma shows that if the set of lower-leveled gates is small, then there exists a partition $X = Y \sqcup Z$ under which the polynomial computed by the circuit is not of full rank.

► **Lemma 17.** *Let Φ be a syntactically multilinear arithmetic circuit over \mathbb{F} and $X = \{x_1, \dots, x_n\}$, for an even integer n , computing f . Let $\tau = 3 \log n$ and $\mathcal{L} = \mathcal{L}(\Phi, 100\tau)$. If $|\mathcal{L}| < n / (10^5 \tau)$, then there exists a partition $X = Y \sqcup Z$ such that $\text{rank}_{Y,Z}(f) < 2^{n/2-1}$.*

We first sketch how Theorem 14 follows from Lemma 17. The proof is identical to the proof given in [31] with slightly different parameters.

Proof of Theorem 14 assuming Lemma 17. Let Ψ' be the arithmetic circuit computing all n first-order partial derivatives of f , given by Theorem 15. Set $\tau = 3 \log n$ and let $\mathcal{L} = \mathcal{L}(\Psi', 100\tau)$ and $\mathcal{U} = \mathcal{U}(\Psi', 100\tau)$ as in Definition 16.

Denote $f_i = \partial f / \partial x_i$ and let v_i be the gate in Ψ' computing f_i , and Ψ'_i be the subcircuit of Ψ' rooted at v_i . Let $\mathcal{L}_i = \mathcal{L}(\Psi'_i, 100\tau)$. It is not hard to show (see [31]) that $\mathcal{L}_i \subseteq \mathcal{L}$, and by Lemma 17 and item 4 in Proposition 8, it follows that $|\mathcal{L}_i| \geq n / (10^5 \tau)$.

For every gate v in Ψ' define $C_v = \{i \in [n] : v \text{ is a gate in } \Psi'_i\}$ to be the set of indices i such that there exists a directed path from v to the output gate computing f_i . For $i \in [n]$, let $\mathcal{U}_i = \{u \in \mathcal{U} : u \text{ is a gate in } \Psi'_i\}$, so that $\sum_{u \in \mathcal{U}} C_u = \sum_{i \in [n]} |\mathcal{U}_i|$.

Since the fan-in of each gate is at most two, $|\mathcal{L}_i| \leq 2|\mathcal{U}_i|$, and since every $u \in \mathcal{U}$ satisfies $|X_u| \geq n - 100\tau$, it follows by Theorem 15 that $|C_u| \leq 100\tau$. Thus, we get

$$n \cdot \frac{n}{10^5 \tau} \leq \sum_{i \in [n]} |\mathcal{L}_i| \leq 2 \sum_{i \in [n]} |\mathcal{U}_i| = 2 \sum_{u \in \mathcal{U}} C_u \leq 2|\mathcal{U}| \cdot 100\tau.$$

By item 2 in Theorem 15, and $\tau = 3 \log n$,

$$|\Psi| = \Omega(|\Psi'|) = \Omega(|\mathcal{U}|) = \Omega\left(\frac{n^2}{\log^2 n}\right). \quad \blacktriangleleft$$

It remains to prove Lemma 17. As the proof mostly appears in [31], we only sketch the main steps.

Proof sketch of Lemma 17. Suppose $\mathcal{L} \leq n/(10^5\tau)$. By applying Theorem 13 to the family of sets $\{X_v : v \in \mathcal{L}\}$, it follows that there exists a balanced partition $Y \sqcup Z$ of X such that X_v is τ -unbalanced for every gate $v \in \mathcal{L}$ (one could get slightly improved constants in the case $n = 4p$ by applying Theorem 9).

The proof now proceeds in the exact same manner as the proof of Lemma 5.2 in [31]. In Proposition 5.5 of [31], it is shown that one can write

$$f = \sum_{i \in [\ell]} g_i h_i + g,$$

where $\mathcal{L} = \{v_1, \dots, v_\ell\}$, h_i is the polynomial computed at v_i , and the set of variables appearing in g_i is disjoint from X_{v_i} .

In Claim 5.7 of [31], it is shown that for every $i \in [\ell]$, $\text{rank}_{Y,Z}(g_i h_i) \leq 2^{n/2-\tau}$. This uses the fact that X_{v_i} is τ -unbalanced, the upper bound in item 1 in Proposition 8, and item 3 in the same proposition.

In Proposition 5.8 of [31], it is shown (with the necessary change of parameters) that the degree of g is at most 200τ .

Thus, by the fact that $\tau = 3 \log n$, item 5 and item 2 of Proposition 8, it follows that for large enough n ,

$$\text{rank}_{Y,Z}(f) \leq \ell \cdot 2^{n/2-\tau} + 2^{\tau^3} < 2^{n/2-1}. \quad \blacktriangleleft$$

4.1 An explicit full-rank polynomial

In this section, for the sake of completeness, we give a construction of a polynomial which is full-rank under any partition of the variables.

► **Construction 18** (Full rank polynomial, [31]). *Let n be an even integer, and let $\mathcal{W} = \{\omega_1, \dots, \omega_n\}$ and $X = \{x_1, \dots, x_n\}$ be sets of variables. For a set $B \in \binom{[n]}{n/2}$, denote by $i_1 < \dots < i_{n/2}$ the elements of B in increasing order, and by $j_1 < \dots < j_{n/2}$ the elements of $[n] \setminus B$ in increasing order. Define $r_B = \prod_{\ell \in B} \omega_\ell$, and $g_B = \prod_{\ell \in [n/2]} (x_{i_\ell} + x_{j_\ell})$.*

Finally, define

$$f = \sum_{B \in \binom{[n]}{n/2}} r_B g_B.$$

► **Claim 19** ([31]). *For f from Construction 18, it holds that for every balanced partition of $X = Y \sqcup Z$, $\text{rank}_{Y,Z}(f) = 2^{n/2}$, where the rank is taken over $\mathbb{F}(\mathcal{W})$.*

We give a proof which is shorter and simpler than the one given in [31].

Proof of Claim 19. Fix a balanced partition $X = Y \sqcup Z$, and consider the matrix $M_{Y,Z}(f)$ where f is interpreted as a polynomial in $f \in (\mathcal{F}[\mathcal{W}])[X]$ (that is, the rows and columns of the matrix are indexed by X variables and its entries are polynomials in \mathcal{W}). We want to show that $\det(M_{Y,Z}(f)) \in \mathbb{F}[\mathcal{W}]$ is a non-zero polynomial. Fix $\omega_i = 1$ if $i \in Y$ and $\omega_i = 0$ otherwise. Under this restriction, $f = g_Y$. It is also not hard to see that $\det(M_{Y,Z}(g_Y)) \neq 0$, since this is a permutation matrix (this also follows from item 3 of Proposition 8). Thus, $\det(M_{Y,Z}(f))$ evaluates to a non-zero value under this setting of the variables \mathcal{W} , which implies it a non-zero polynomial. ◀

► **Corollary 20.** *Every syntactically multilinear circuit computing the polynomial f has size at least $\Omega(n^2/\log^2 n)$.*

The polynomial f in Construction 18 is in the class VNP of explicit polynomials, but it is not known whether there exists a polynomial size multilinear circuit for f .

Raz and Yehudayoff [32] constructed a full-rank polynomial $g \in \mathbb{F}[X, \mathcal{W}']$ that has a syntactically multilinear circuit of size $O(n^3)$. Their construction also uses a set of auxiliary variables \mathcal{W}' of size $O(n^3)$. Thus, if one measures the complexity as a function of $|X| \cup |\mathcal{W}'|$, the quadratic lower bound of Theorem 14 is meaningless, because a lower bound of $\Omega(n^3)$ holds trivially. However, we believe that since the rank is taken over $\mathbb{F}(\mathcal{W}')$, it is only fair to consider computations over $\mathbb{F}(\mathcal{W}')$, where any rational expression in the variables of \mathcal{W}' is merely a field constant. Thus, in this setting, an input gate can be labeled by an arbitrarily complex rational function in the variables of \mathcal{W}' , and the complexity is measured as a function of $|X|$ alone. In this model the lower bound of Theorem 14 is meaningful, and furthermore, this example shows that the partial derivative matrix technique cannot prove an $\omega(n^3)$ lower bound.

References

- 1 Noga Alon, E. E. Bergmann, Don Coppersmith, and Andrew M. Odlyzko. Balancing sets of vectors. *IEEE Trans. Information Theory*, 34(1):128–130, 1988. doi:10.1109/18.2610.
- 2 Noga Alon, Mrinal Kumar, and Ben Lee Volk. An almost quadratic lower bound for syntactically multilinear arithmetic circuits. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:124, 2017. URL: <https://eccc.weizmann.ac.il/report/2017/124>.
- 3 Noga Alon and Joel H. Spencer. *The Probabilistic Method*. Wiley Publishing, 4th edition, 2016. URL: <http://eu.wiley.com/WileyCDA/WileyTitle/productCd-1119061954.html>.
- 4 Richard P. Anstee, Lajos Rónyai, and Attila Sali. Shattering news. *Graphs and Combinatorics*, 18(1):59–73, 2002. doi:10.1007/s003730200003.
- 5 Walter Baur and Volker Strassen. The complexity of partial derivatives. *Theor. Comput. Sci.*, 22:317–330, 1983. doi:10.1016/0304-3975(83)90110-X.
- 6 Stuart J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Inf. Process. Lett.*, 18(3):147–150, 1984. doi:10.1016/0020-0190(84)90018-8.
- 7 Xi Chen, Neeraj Kayal, and Avi Wigderson. Partial derivatives in arithmetic complexity and beyond. *Foundations and Trends in Theoretical Computer Science*, 6(1-2):1–138, 2011. doi:10.1561/04000000043.
- 8 L. Csanky. Fast parallel matrix inversion algorithms. *SIAM J. Comput.*, 5(4):618–623, 1976. doi:10.1137/0205040.
- 9 H. Enomoto, Peter Frankl, N. Ito, and K. Nomura. Codes with given distances. *Graphs and Combinatorics*, 3(1):25–38, 1987. doi:10.1007/BF01788526.
- 10 Jeffrey B. Farr and Shuhong Gao. Computing gröbner bases for vanishing ideals of finite sets of points. In Marc P. C. Fossorier, Hideki Imai, Shu Lin, and Alain Poli, editors, *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes, 16th International Symposium, AAECC-16, Las Vegas, NV, USA, February 20-24, 2006, Proceedings*, volume 3857 of *Lecture Notes in Computer Science*, pages 118–127. Springer, 2006. doi:10.1007/11617983_11.
- 11 Hervé Fournier, Nutan Limaye, Guillaume Malod, and Srikanth Srinivasan. Lower bounds for depth 4 formulas computing iterated matrix multiplication. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 128–135. ACM, 2014. doi:10.1145/2591796.2591824.

- 12 Peter Frankl and Vojtěch Rödl. Forbidden intersections. *Trans. Amer. Math. Soc.*, 300(1):259–286, 1987. doi:10.2307/2000598.
- 13 Dima Grigoriev and Marek Karpinski. An exponential lower bound for depth 3 arithmetic circuits. In Jeffrey Scott Vitter, editor, *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 577–582. ACM, 1998. doi:10.1145/276698.276872.
- 14 Dima Grigoriev and Alexander A. Razborov. Exponential lower bounds for depth 3 arithmetic circuits in algebras of functions over finite fields. *Appl. Algebra Eng. Commun. Comput.*, 10(6):465–487, 2000. doi:10.1007/s002009900021.
- 15 Ankit Gupta, Pritish Kamath, Neeraj Kayal, and Ramprasad Satharishi. Approaching the chasm at depth four. *J. ACM*, 61(6):33:1–33:16, 2014. doi:10.1145/2629541.
- 16 Gábor Hegedűs. Balancing sets of vectors. *Studia Sci. Math. Hungar.*, 47(3):333–349, 2010. doi:10.1556/SScMath.2009.1134.
- 17 Gábor Hegedűs and Lajos Rónyai. Gröbner bases for complete uniform families. *J. Algebraic Combin.*, 17(2):171–180, 2003. doi:10.1023/A:1022934815185.
- 18 Maurice J. Jansen. Lower bounds for syntactically multilinear algebraic branching programs. In Edward Ochmanski and Jerzy Tyszkiewicz, editors, *Mathematical Foundations of Computer Science 2008, 33rd International Symposium, MFCS 2008, Torun, Poland, August 25-29, 2008, Proceedings*, volume 5162 of *Lecture Notes in Computer Science*, pages 407–418. Springer, 2008. doi:10.1007/978-3-540-85238-4_33.
- 19 K. Kalorkoti. A lower bound for the formula size of rational functions. *SIAM J. Comput.*, 14(3):678–687, 1985. doi:10.1137/0214050.
- 20 Neeraj Kayal, Nutan Limaye, Chandan Saha, and Srikanth Srinivasan. An exponential lower bound for homogeneous depth four arithmetic formulas. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 61–70. IEEE Computer Society, 2014. doi:10.1109/FOCS.2014.15.
- 21 Donald E. Knuth. Efficient balanced codes. *IEEE Trans. Information Theory*, 32(1):51–53, 1986. doi:10.1109/TIT.1986.1057136.
- 22 Mrinal Kumar. A quadratic lower bound for homogeneous algebraic branching programs. In Ryan O’Donnell, editor, *32nd Computational Complexity Conference, CCC 2017, July 6-9, 2017, Riga, Latvia*, volume 79 of *LIPICs*, pages 19:1–19:16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017. doi:10.4230/LIPICs.CCC.2017.19.
- 23 Mrinal Kumar and Ramprasad Satharishi. An exponential lower bound for homogeneous depth-5 circuits over finite fields. In Ryan O’Donnell, editor, *32nd Computational Complexity Conference, CCC 2017, July 6-9, 2017, Riga, Latvia*, volume 79 of *LIPICs*, pages 31:1–31:30. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017. doi:10.4230/LIPICs.CCC.2017.31.
- 24 Mrinal Kumar and Shubhangi Saraf. On the power of homogeneous depth 4 arithmetic circuits. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 364–373. IEEE Computer Society, 2014. doi:10.1109/FOCS.2014.46.
- 25 Meena Mahajan and V. Vinay. Determinant: Combinatorics, algorithms, and complexity. *Chicago J. Theor. Comput. Sci.*, 1997, 1997. Preliminary version in the *8th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 1997)*. URL: <http://cjtcs.cs.uchicago.edu/articles/1997/5/contents.html>.
- 26 Noam Nisan. Lower bounds for non-commutative computation (extended abstract). In Cris Koutsougeras and Jeffrey Scott Vitter, editors, *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing, May 5-8, 1991, New Orleans, Louisiana, USA*, pages 410–418. ACM, 1991. doi:10.1145/103418.103462.

- 27 Noam Nisan and Avi Wigderson. Lower bounds on arithmetic circuits via partial derivatives. *Computational Complexity*, 6(3):217–234, 1997. doi:10.1007/BF01294256.
- 28 Ran Raz. Separation of multilinear circuit and formula size. *Theory of Computing*, 2(6):121–135, 2006. doi:10.4086/toc.2006.v002a006.
- 29 Ran Raz. Multi-linear formulas for permanent and determinant are of super-polynomial size. *J. ACM*, 56(2):8:1–8:17, 2009. doi:10.1145/1502793.1502797.
- 30 Ran Raz. Elusive functions and lower bounds for arithmetic circuits. *Theory of Computing*, 6(1):135–177, 2010. doi:10.4086/toc.2010.v006a007.
- 31 Ran Raz, Amir Shpilka, and Amir Yehudayoff. A lower bound for the size of syntactically multilinear arithmetic circuits. *SIAM J. Comput.*, 38(4):1624–1647, 2008. doi:10.1137/070707932.
- 32 Ran Raz and Amir Yehudayoff. Balancing syntactically multilinear arithmetic circuits. *Computational Complexity*, 17(4):515–535, 2008. doi:10.1007/s00037-008-0254-0.
- 33 Ran Raz and Amir Yehudayoff. Lower bounds and separations for constant depth multilinear circuits. *Computational Complexity*, 18(2):171–207, 2009. doi:10.1007/s00037-009-0270-8.
- 34 Ramprasad Saptharishi. A survey of lower bounds in arithmetic circuit complexity. Github survey, <https://github.com/dasarpmar/lowerbounds-survey/>, 2016. URL: <https://github.com/dasarpmar/lowerbounds-survey/releases/>.
- 35 Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends in Theoretical Computer Science*, 5(3-4):207–388, 2010. doi:10.1561/04000000039.
- 36 Matthew Skala. Hypergeometric tail inequalities: ending the insanity. *arXiv preprint arXiv:1311.5939*, 2013. URL: <https://arxiv.org/abs/1311.5939>.
- 37 Volker Strassen. Die berechnungskomplexität von elementarsymmetrischen funktionen und von interpolationskoeffizienten. *Numerische Mathematik*, 20(3):238–251, 1973. doi:10.1007/BF01436566.