





Improved Low-Depth Set-Multilinear Circuit Lower Bounds

Deepanshu Kush  

Department of Computer Science, University of Toronto, Canada

Shubhangi Saraf  

Department of Mathematics and Department of Computer Science, University of Toronto, Canada

Abstract

In this paper, we prove strengthened lower bounds for constant-depth set-multilinear formulas. More precisely, we show that over any field, there is an explicit polynomial f in VNP defined over n^2 variables, and of degree n , such that any product-depth Δ set-multilinear formula computing f has size at least $n^{\Omega(n^{1/\Delta}/\Delta)}$. The hard polynomial f comes from the class of Nisan-Wigderson (NW) design-based polynomials.

Our lower bounds improve upon the recent work of Limaye, Srinivasan and Tavenas (STOC 2022), where a lower bound of the form $(\log n)^{\Omega(\Delta n^{1/\Delta})}$ was shown for the size of product-depth Δ set-multilinear formulas computing the iterated matrix multiplication (IMM) polynomial of the same degree and over the same number of variables as f . Moreover, our lower bounds are novel for any $\Delta \geq 2$.

The precise quantitative expression in our lower bound is interesting also because the lower bounds we obtain are “sharp” in the sense that any asymptotic improvement would imply general set-multilinear circuit lower bounds via depth reduction results.

In the setting of general set-multilinear formulas, a lower bound of the form $n^{\Omega(\log n)}$ was already obtained by Raz (J. ACM 2009) for the more general model of multilinear formulas. The techniques of LST (which extend the techniques of the same authors in (FOCS 2021)) give a different route to set-multilinear formula lower bounds, and allow them to obtain a lower bound of the form $(\log n)^{\Omega(\log n)}$ for the size of general set-multilinear formulas computing the IMM polynomial. Our proof techniques are another variation on those of LST, and enable us to show an improved lower bound (matching that of Raz) of the form $n^{\Omega(\log n)}$, albeit for the same polynomial f in VNP (the NW polynomial). As observed by LST, if the same $n^{\Omega(\log n)}$ size lower bounds for unbounded-depth set-multilinear formulas could be obtained for the IMM polynomial, then using the self-reducibility of IMM and using hardness escalation results, this would imply super-polynomial lower bounds for general algebraic formulas.

2012 ACM Subject Classification Theory of computation \rightarrow Algebraic complexity theory

Keywords and phrases algebraic circuit complexity, complexity measure, set-multilinear formulas

Digital Object Identifier 10.4230/LIPIcs.CCC.2022.38

Related Version *Full Version:* <https://arxiv.org/abs/2205.00611>

Full Version: <https://ecc.weizmann.ac.il/report/2022/064/>

Funding *Shubhangi Saraf:* Research partially supported by a Sloan research fellowship and an NSERC Discovery Grant.

Acknowledgements We would like to thank Swastik Kopparty, Mrinal Kumar, and Ben Rossman for several helpful discussions.



© Deepanshu Kush and Shubhangi Saraf;
licensed under Creative Commons License CC-BY 4.0
37th Computational Complexity Conference (CCC 2022).

Editor: Shachar Lovett; Article No. 38; pp. 38:1–38:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



1 Introduction

1.1 Background

An *algebraic circuit* over a field \mathbb{F} for a multivariate polynomial $P(x_1, \dots, x_N)$ is a directed acyclic graph (DAG) whose internal vertices (called gates) are labeled as either $+$ (sum) or \times (product), and leaves (vertices of in-degree zero) are labeled by the variables x_i or constants from \mathbb{F} . A special output gate (the root of the DAG) represents the polynomial P . If the DAG happens to be a tree, such a resulting circuit is called an *algebraic formula*. The size of a circuit is the number of nodes in the DAG. We also consider the product-depth of the circuit, which is the maximum number of product gates on a root-to-leaf path.

An algebraic circuit is therefore a computational model, which solves the computational task of evaluating P on a given input (x_1, \dots, x_N) . The complexity of this model is measured by the size of the circuit, which serves as an indicator of the time complexity of computing the polynomial. The product-depth measures the degree to which this computation can be made parallel. As an algebraic circuit is supposed to construct a formal polynomial P , it is a *syntactic* model of computation. This is unlike a Boolean circuit, which is only required to model specific truth-table constraints. The problem of proving algebraic circuit lower bounds is therefore widely considered to be easier than its Boolean counterpart. Indeed, we know that proving $\text{VP} \neq \text{VNP}$, the algebraic analog of the P vs. NP problem, is implied by the latter separation, in the non-uniform setting ([3]). We refer the reader to [28] for a much more elaborate survey of this topic.

1.2 The LST breakthrough

Much like in the Boolean setting, the problem of showing lower bounds for *general* algebraic circuits (or even formulas) has remained elusive. However, some remarkable progress has been made very recently by Limaye, Srinivasan, and Tavenas ([20]) who in a spectacular breakthrough, showed the first super-polynomial lower bounds for algebraic circuits of *all* constant depths. Prior to their work, the best known lower bound ([14]) even for product-depth 1 (or $\Sigma\Pi\Sigma$ circuits) was only almost-cubic. This is in stark contrast with the Boolean setting, in which we have known strong constant-depth lower bounds for many decades [2, 8, 31, 9, 27, 29]. Constant-depth circuits are critical to the study of algebraic complexity theory, as unlike the Boolean setting, strong enough bounds against them are known to yield $\text{VP} \neq \text{VNP}$ ([1]). This helps put into perspective the importance of the work [20].

The crucial step in the proof of their result is to first establish super-polynomial lower bounds for a certain restricted class of (low-depth) algebraic circuits, namely *set-multilinear* circuits which we now define along with other important circuit models. A polynomial is said to be homogeneous if each monomial has the same total degree and *multilinear* if every variable occurs at most once in any monomial. Now, suppose that the underlying variable set is partitioned into d sets X_1, \dots, X_d . Then the polynomial is said to be *set-multilinear* with respect to this variable partition if each monomial in P has *exactly* one variable from each set. We also define different models of computation corresponding to these variants of polynomials classes. An algebraic formula (circuit) is set-multilinear with respect to a variable partition (X_1, \dots, X_d) if each internal node in the formula (circuit) computes a set-multilinear polynomial. Multilinear/homogeneous circuits and formulas are defined analogously.

Several well-studied and interesting polynomials happen to be set-multilinear. For example, the Determinant and the Permanent polynomials, the study of which is profoundly consequential to the field of algebraic complexity theory, are set-multilinear (with respect

to the column variables). Another well-studied polynomial, namely the Iterated Matrix Multiplication polynomial, is also set-multilinear. The polynomial $\text{IMM}_{n,d}$ is defined on $N = dn^2$ variables, where the variables are partitioned into d sets X_1, \dots, X_d of size n^2 , each of which is represented as an $n \times n$ matrix with distinct variable entries. The polynomial $\text{IMM}_{n,d}$ is defined to be the polynomial that is the $(1, 1)$ -th entry of the product matrix $X_1 X_2 \cdots X_d$. This polynomial has a simple divide-and-conquer-based set-multilinear formula of size $n^{O(\log d)}$, and more generally for every $\Delta \leq \log d$, a set-multilinear formula of product-depth Δ and size $n^{O(\Delta d^{1/\Delta})}$, and circuit¹ of size $n^{O(d^{1/\Delta})}$. Even without the set-multilinearity constraint, no significantly better upper bound is known. It is reasonable to conjecture that this simple upper bound is tight up to the constant in the exponent.

The lower bounds in [20] for general constant-depth algebraic circuits are shown in the following sequence of steps:

1. It is shown that general low-depth algebraic circuits can be transformed to set-multilinear algebraic circuits of low depth, and without much of a blow-up in size (as long as the degree is small). More precisely, any product-depth Δ circuit of size s computing a polynomial that is set-multilinear with respect to the partition (X_1, \dots, X_d) where each $|X_i| \leq n$, can be converted to a set-multilinear circuit² of product-depth 2Δ and size $\text{poly}(s) \cdot d^{O(d)}$. Such a “set-multilinearization” of general formulas of small degree was already shown before in [25] (which we describe soon in more detail); however, the main contribution of [20] here is to prove this *depth-preserving* version of it.
2. Strong lower bounds are then established for low-depth set-multilinear circuits (of small enough degree). More precisely, any set-multilinear circuit C computing $\text{IMM}_{n,d}$ (where $d = O(\log n)$) of product-depth Δ must have size at least $n^{d^{\exp(-O(\Delta))}}$. This combined with the first step yields the desired lower bound for general constant-depth circuits.

Given Raz’s set-multilinearization of formulas of small degree that we alluded to, and this description of the set-multilinear formula lower bounds from [20] when $d = O(\log n)$, it is evident the “small degree” regime is inherently interesting to study - as it provides an avenue, via “hardness escalation”, for tackling one of the grand challenges of algebraic complexity theory, namely proving super-polynomial lower bounds for general algebraic formulas. However, we shall now see that even the large degree regime can be equally consequential in this regard.

1.3 The large degree regime

Consider a polynomial P that is set-multilinear with respect to the variable partition (X_1, \dots, X_d) where each $|X_i| \leq n$. The main focus of this paper is to study set-multilinear circuit complexity in the regime where d and n are *polynomially* related (as opposed to say, the assumption $d = O(\log n)$ described above). We now provide some background and motivation for studying this regime.

In follow-up work [21], the same authors showed the first super-polynomial lower bound against unbounded-depth set-multilinear formulas computing $\text{IMM}_{n,n}$ ³. As is astutely described in [21], studying the set-multilinear formula complexity of IMM is extremely interesting and consequential even in the setting $d = n$ because of the following reasons:

¹ In this paper, when speaking of constant-depth models of computation at a high level, we shall often use the terms circuit and formula interchangeably as a product-depth Δ circuit of size s can be simulated by a product-depth Δ formula of size $s^{2\Delta}$.

² There is also an intermediate “homogenization” step which we skip describing here for the sake of brevity.

³ Note that for $\text{IMM}_{n,n}$, each X_i has size n^2 , not n . But the important thing for us here is that the degree, n , is polynomially related to this parameter.

- $\text{IMM}_{n,n}$ is a *self-reducible* polynomial i.e., it is possible to construct formulas for $\text{IMM}_{n,n}$ by recursively using formulas for $\text{IMM}_{n,d}$ (for any $d < n$). In particular, if we had formulas of size $n^{o(\log d)}$ for $\text{IMM}_{n,d}$ (for some $d < n$), this would imply formulas of size $n^{o(\log n)}$ for $\text{IMM}_{n,n}$. In other words, an optimal $n^{\Omega(\log n)}$ lower bound for $\text{IMM}_{n,n}$ implies $n^{\omega_d(1)}$ lower bounds for $\text{IMM}_{n,d}$ for any $d < n$.
- Raz in [25] showed that if an N -variate set-multilinear polynomial of degree d has an algebraic formula of size s , then it also has a set-multilinear formula of size $\text{poly}(s) \cdot (\log s)^d$. In particular, for a set-multilinear polynomial P of degree $d = O(\log N / \log \log N)$, it follows that P has a formula of size $\text{poly}(N)$ if and only if P has a set-multilinear formula of size $\text{poly}(N)$. Thus, having $N^{\omega_d(1)}$ set-multilinear formula size lower bounds for such a low degree would imply super-polynomial lower bounds for general formulas.

In particular, proving the optimal $n^{\Omega(\log n)}$ set-multilinear formula size lower bound for $\text{IMM}_{n,n}$ would have dramatic consequences. To this end, the authors in [21] are able to show a weaker bound of the form $(\log n)^{\Omega(\log n)}$ instead. Even though it is the case that “simply” improving the base of this exponent from $\log n$ to n yields general formula lower bounds, it seems that we are still far from achieving it. Indeed, as is observed in [21], we do not even have the optimal $n^{\Omega(\sqrt{n})}$ lower bound⁴ when product-depth $\Delta = 2$. Moreover, we do not know how to obtain a lower bound of the form $n^{\Omega(\sqrt{n})}$ for product-depth 2 set-multilinear circuits for *any* explicit polynomial of degree n and in $\text{poly}(n)$ variables. For product-depths $\Delta \leq \log n$, [21] shows a set-multilinear formula size lower bound of $(\log n)^{\Omega(\Delta n^{1/\Delta})}$ for $\text{IMM}_{n,n}$, which is in fact the best set-multilinear lower bound we know for any polynomial of degree n and in $\text{poly}(n)$ variables, and for any $\Delta \geq 2$. As far as we know, the previous best lower bound of $\exp(\Omega(n^{1/\Delta}))$, also for $\text{IMM}_{n,n}$, followed from the work of Nisan and Wigderson ([23]). It is therefore an interesting challenge to improve the base of this exponent from $\log n$ to n i.e., establish a near-optimal $n^{\Omega(n^{1/\Delta})}$ lower bound in the constant (or low) depth setting.

1.4 Our Results

In this paper, we obtain these “optimal” lower bounds, albeit not for $\text{IMM}_{n,n}$, but rather for another explicit polynomial in VNP. We show the following:

► **Theorem 1.** *Let N be a growing parameter and Δ be an integer such that $1 \leq \Delta \leq \log N / \log \log N$. There is an explicit polynomial P_N defined over $N = n^2$ variables with degree $d = n$ that is set-multilinear with respect to the variable partition $X = (X_1, \dots, X_d)$ where each $|X_i| = n$ and such that any set-multilinear formula of product-depth Δ computing $P_N(X)$ must have size at least $N^{\Omega(d^{1/\Delta}/\Delta)}$.*

Notice that obtaining this precise bound is interesting also when viewed through the lens of *depth reduction*. Tavenas ([30]), building on several prior works ([1, 16]), showed that any algebraic circuit of $\text{poly}(N)$ size computing a homogeneous N -variate polynomial of degree d can be converted to a homogeneous circuit of product-depth⁵ Δ of size $(Nd)^{O(d^{1/\Delta})}$. It easily follows from the proof that this depth reduction preserves syntactic restrictions. That is, if we start with a syntactically set-multilinear circuit, the resulting product-depth Δ circuit is also syntactically set-multilinear. Therefore, the precise bound in Theorem 1 is *sharp* in the

⁴ This is known for set-multilinear (and even multilinear) $\Sigma\Pi\Sigma\Pi$ circuits (see [7, 15]), but those are only special cases of general product-depth 2 circuits, which are $\Sigma\Pi\Sigma\Pi\Sigma$.

⁵ The result is stated in [30] for $\Sigma\Pi\Sigma\Pi$ circuits but the proof can be appropriately modified for larger product-depths.

sense that any asymptotic improvement in its exponent would imply super-polynomial set-multilinear circuit lower bounds, which would be quite a strong and interesting consequence. Another very intriguing direction is to consider the problem of *improved* depth reduction for set-multilinear circuits. If an asymptotic improvement in the exponent on the bound for general circuits from [30] could be shown to hold for set-multilinear circuits in the setting of Theorem 1 (i.e., when $N = d^2$), this would again imply super-polynomial set-multilinear circuit lower bounds. There is some evidence towards this possibility, as [17] shows such an improvement in a certain regime of parameters for multilinear circuits (see the discussion in Section 4 for more details).

► **Remark 2.** The lower bound in Theorem 1 is actually $d^{\Omega(d^{1/\Delta}/\Delta)}$, where d is the degree of the underlying polynomial, and it holds as long as degree $d \leq n$ (the details are deferred to the proof of Theorem 9 in Section 3). Observe that for constant Δ this bound already nearly matches the bound $(\log n)^{\Omega(\Delta d^{1/\Delta})}$ in [21] (which was shown for $\text{IMM}_{n,d}$) when $d = (\log n)^{\Omega(1)}$ and exceeds it as soon as d becomes super-polylogarithmic in n . Moreover for $d < \log n / \log \log n$, both the bounds are trivial even for $\Delta = 1$.

We also remark that in several lower bounds for algebraic circuit classes in the past, the lower bound was initially shown for a polynomial in VNP and then with additional effort, was shown to also hold for a polynomial in VP (in particular, the IMM polynomial). A strong candidate for the choice of this polynomial family in VNP has been the Nisan-Wigderson (NW) design-based ([22]) family of polynomials. For instance, [13] showed a lower bound of $n^{\Omega(\sqrt{n})}$ for the top fan-in of a $\Sigma\Pi^{[O(\sqrt{n})]}\Sigma\Pi^{[\sqrt{n}]}$ circuit computing the NW polynomial, which was subsequently shown for IMM by [7]. Similarly, [11] showed an $n^{\Omega(\sqrt{d})}$ size lower bound for homogeneous depth-4 algebraic formulas for the NW polynomial, which was then shown for IMM later in [19]. Much like these examples, our hard polynomial family in Theorem 1 is also indeed the NW polynomial family, as we shall see in Section 3. Our motivation to study constant-depth set-multilinear formula complexity was to prove the optimal lower bounds for the IMM polynomial. Although we are presently able to show it only for the NW polynomial instead of IMM, we are hopeful that this is an important step in its direction.

In addition to our lower bound for bounded-depth set-multilinear formulas, we observe that the same proof technique also implies a lower bound of the form $n^{\Omega(\log n)}$ for unbounded-depth set-multilinear formulas. [21] showed a weaker bound of the form $(\log n)^{\Omega(\log n)}$ but for $\text{IMM}_{n,n}$.

► **Theorem 3.** *For a given integer N , there is an explicit polynomial P_N defined over $N = n^2$ variables with degree $d = n$ that is set-multilinear with respect to the variable partition $X = (X_1, \dots, X_d)$ where each $|X_i| = n$ such that any set-multilinear formula computing $P_N(X)$ must have size at least $N^{\Omega(\log N)}$.*

The hard polynomial in Theorem 3 is also the NW polynomial, which if “improved” to $\text{IMM}_{n,n}$, then as discussed, would yield super-polynomial general formula lower bounds. However, we note that in this case, our result is in some sense subsumed by the result of Raz ([24]) who showed an $n^{\Omega(\log n)}$ lower bound for the $n \times n$ permanent (or determinant) polynomial for unbounded-depth multilinear formulas.

1.5 Other Related Work

In the bounded-depth setting, other than the works [20, 21, 23] already mentioned, there have been several lower bounds for the class of low-depth *multilinear* circuits ([26, 5, 4, 12]). In the unbounded-depth setting, apart from the works [21, 24] already mentioned for set-multilinear formulas, there have also been several strong lower bounds of the form $n^{\Omega(\log n)}$ against

multilinear formulas ([6, 10, 15]). However, in both settings of depth, several of these works are not even applicable to the set-multilinear setting as the corresponding hard polynomial does not happen to be set-multilinear.

1.6 Proof overview

Our overall proof techniques are similar to that of many known lower bounds. We work with a measure that we show to be small for all polynomials computed by small enough set-multilinear formulas (appropriately so in the bounded and unbounded-depth settings) and large for the NW polynomial. These *partial derivative measures* were introduced by Nisan and Wigderson in [23], who used them to prove the constant-depth set-multilinear formula lower bounds we discussed earlier. [20, 21] use a particular variant of this measure and our measure is in turn inspired from these works.

Given a variable partition (X_1, \dots, X_d) , we label each set of variables X_i as “positive” or “negative” uniformly at random. Let \mathcal{P} and \mathcal{N} denote the set of positive and negative indices respectively, and let $\mathcal{M}^{\mathcal{P}}$ and $\mathcal{M}^{\mathcal{N}}$ denote the sets of all set-multilinear monomials over \mathcal{P} and \mathcal{N} respectively. For a polynomial that is set-multilinear over the given variable partition (X_1, \dots, X_d) , our measure then is simply the rank of the “partial derivative matrix” whose rows are indexed by the elements of $\mathcal{M}^{\mathcal{P}}$ and columns indexed by $\mathcal{M}^{\mathcal{N}}$, and the entry of this matrix corresponding to a row m_1 and a column m_2 is the coefficient of the monomial $m_1 \cdot m_2$ in the given polynomial.

In contrast, the measure used in [20] is deterministic and moreover, it is *asymmetric* with respect to the positive and negative variable sets, in the sense that while keeping the positive variable sets as is, it first reduces the size of the negative variable sets by arbitrarily setting a few of these variables to field constants, and then works with the resulting polynomial. On the other hand, [21] does use a randomized measure, but one that is still asymmetric, relying on randomly setting a few of the variables inside each set to constants. The way they control the discrepancy between the sizes of the positive and negative variable sets (which is indeed crucial for obtaining the claimed lower bounds) is by imposing a Martingale-like distribution. The lower bound of [23] also uses random restrictions to enable them to effectively “simplify” the circuit and upper bound its complexity. Our symmetric, randomized measure avoids random restrictions altogether, and though it is mainly inspired by the measure and the techniques from [20], it is also reminiscent of the measures used in [24, 26] to prove multilinear formula lower bounds.

2 Preliminaries

We begin by defining the hard polynomial of our main result (Theorem 1). As is done in previous lower bounds using the NW polynomials (for example, see [13]), we will identify the set of the first n integers as elements of \mathbb{F}_n via an arbitrary correspondence $\phi : [n] \rightarrow \mathbb{F}_n$. If $f(z) \in \mathbb{F}_n[z]$ is a univariate polynomial, then we abuse notation to let $f(i)$ denote the evaluation of f at the i -th field element via the above correspondence i.e., $f(i) := \phi^{-1}(f(\phi(i)))$. To simplify the exposition, in the following definition, we will omit the correspondence ϕ and identify a variable $x_{i,j}$ by the point $(\phi(i), \phi(j)) \in \mathbb{F}_n \times \mathbb{F}_n$.

► **Definition 4** (Nisan-Wigderson Polynomials). *For a prime power n , let \mathbb{F}_n be a field of size n . For an integer $d \leq n$ and the set X of nd variables $\{x_{i,j} : i \in [n], j \in [d]\}$, we define the degree d homogeneous polynomial $NW_{n,d}$ over any field as*

$$NW_{n,d}(X) = \sum_{\substack{f(z) \in \mathbb{F}_n[z] \\ \deg(f) < d/2}} \prod_{j \in [d]} x_{f(j),j}.$$

Next, we turn to the measure that we shall use to prove Theorems 1 and 3. For the purpose of setting it up, we follow the notation of [20] in the following definition. However, we do remark that we do not need it in its full generality as we will eventually work with a simpler, *symmetric* notion that was alluded to in Section 1. Nevertheless, employing the same notation has the advantage that the reader is quite possibly already familiar with it in the context of proving set-multilinear circuit lower bounds.

► **Definition 5** (Relative Rank Measure of [20, 21]). *Let f be a polynomial that is set-multilinear with respect to the variable partition (X_1, X_2, \dots, X_d) where each set is of size n . Let $w = (w_1, w_2, \dots, w_d)$ be a tuple (or word) of non-zero real numbers such that $2^{|w_i|} \in [n]$ for all $i \in [d]$. For each $i \in [d]$, let $X_i(w)$ be the variable set obtained by removing arbitrary variables from the set X_i such that $|X_i(w)| = 2^{|w_i|}$, and let $\overline{X}(w)$ denote the tuple of sets of variables $(X_1(w), \dots, X_d(w))$. Corresponding to a word w , define $\mathcal{P}_w := \{i \mid w_i > 0\}$ and $\mathcal{N}_w := \{i \mid w_i < 0\}$. Let $\mathcal{M}_w^{\mathcal{P}}$ be the set of all set-multilinear monomials over a subset of the variable sets $X_1(w), X_2(w), \dots, X_d(w)$ indexed by \mathcal{P}_w , and similarly let $\mathcal{M}_w^{\mathcal{N}}$ be the set of all set-multilinear monomials over these variable sets indexed by \mathcal{N}_w .*

Define the ‘partial derivative matrix’ matrix $\mathcal{M}_w(f)$ whose rows are indexed by the elements of $\mathcal{M}_w^{\mathcal{P}}$ and columns indexed by the elements of $\mathcal{M}_w^{\mathcal{N}}$ as follows: the entry of this matrix corresponding to a row m_1 and a column m_2 is the coefficient of the monomial $m_1 \cdot m_2$ in f . We define

$$\text{relrk}_w(f) := \frac{\text{rank}(\mathcal{M}_w(f))}{\sqrt{|\mathcal{M}_w^{\mathcal{P}}| \cdot |\mathcal{M}_w^{\mathcal{N}}|}} = \frac{\text{rank}(\mathcal{M}_w(f))}{2^{\frac{1}{2} \sum_{i \in [d]} |w_i|}}.$$

► **Definition 6.** *For any tuple $w = (w_1, \dots, w_t)$ and a subset $S \subseteq [t]$, we shall refer to the sum $\sum_{i \in S} w_i$ by w_S . And by $w|_S$, we will refer to the tuple obtained by considering only the elements of w that are indexed by S . We denote by $\mathbb{F}_{\text{sm}}[\mathcal{T}]$ the set of set-multilinear polynomials over the tuple of sets of variables \mathcal{T} .*

The following is a simple result that establishes various useful properties of the relative rank measure.

▷ **Claim 7** ([20]).

1. (Imbalance) Say $f \in \mathbb{F}_{\text{sm}}[\overline{X}(w)]$. Then, $\text{relrk}_w(f) \leq 2^{-|w_{[d]}|/2}$.
2. (Sub-additivity) If $f, g \in \mathbb{F}_{\text{sm}}[\overline{X}(w)]$, then $\text{relrk}_w(f + g) \leq \text{relrk}_w(f) + \text{relrk}_w(g)$.
3. (Multiplicativity) Say $f = f_1 f_2 \cdots f_t$ and assume that for each $i \in [t]$, $f_i \in \mathbb{F}_{\text{sm}}[\overline{X}(w|_{S_i})]$, where (S_1, \dots, S_t) is a partition of $[d]$. Then

$$\text{relrk}_w(f) = \prod_{i \in [t]} \text{relrk}_{w|_{S_i}}(f_i).$$

3 Main Result

We are now ready to prove our main result. We start by showing that the *symmetric* relative rank is large for the NW polynomial.

▷ **Claim 8.** For an integer $n = 2^k$ and $d \leq n$, let $w \in \{k, -k\}^d$ with $w_{[d]} = 0$. Then $\text{relrk}_w(NW_{n,d}) = 1$ i.e., $\mathcal{M}_w(NW_{n,d})$ has full rank.

Proof. Fix $n = 2^k$ and d , so that we can also write NW for $NW_{n,d}$, and let $n' = d/2$. The condition on w implies that $|\mathcal{P}_w| = |\mathcal{N}_w| = n'$. Observe that $\mathcal{M}_w(NW)$ is a square matrix of dimension $|\mathcal{M}_w^{\mathcal{P}}| = |\mathcal{M}_w^{\mathcal{N}}| = n^{n'}$. Consider a row of $\mathcal{M}_w(NW)$ indexed by a monomial

$m_1 = x_{i_1, j_1} \cdots x_{i_{n'}, j_{n'}} \in \mathcal{M}_w^P$. m_1 can be thought of as a map from $S = \{j_1, \dots, j_{n'}\}$ to \mathbb{F}_n which sends j_ℓ to i_ℓ for each $\ell \in [n']$. Next, by interpolating the pairs $(j_1, i_1), \dots, (j_{n'}, i_{n'})$, we know that there exists a unique polynomial $f(z) \in \mathbb{F}_n(z)$ of degree $< n'$ for which $f(j_\ell) = i_\ell$ for each $\ell \in [n']$. As a consequence, there is a unique “extension” of the monomial $x_{i_1, j_1} \cdots x_{i_{n'}, j_{n'}}$ that appears as a term in NW , which is precisely $m_1 \cdot \prod_{j \in \mathcal{N}_w} x_{f(j), j}$. Therefore, all but one of the entries in the row corresponding to m_1 must be zero, and the remaining entry must be 1. Applying the same argument to the columns of $\mathcal{M}_w(NW)$, we deduce that $\mathcal{M}_w(NW)$ is a permutation matrix, and so has full rank. \triangleleft

The following is a more precise and general version of Theorem 1 that is stated in Section 1. We also incorporate Remark 2 here and show our lower bound for any degree $d \leq n$. Theorem 1 follows from the special case $d = n$.

► Theorem 9. *For an integer $n = 2^k$, let \mathbb{F}_n be a field of size n . Let d, Δ be integers such that $d \leq n$ is large enough⁶ and $\Delta \leq \log d / \log \log d$. Let X_i denote the set of n variables $\{x_{i,j} : j \in [d]\}$ and X be the tuple (X_1, \dots, X_d) . Then, any set-multilinear formula family of product-depth Δ computing $NW_{n,d}(X)$ must have size at least $d^{\Omega(d^{1/\Delta}/\Delta)}$.*

Proof. We show that the symmetric relative rank of low-depth set-multilinear formulas is small with high probability in the lemma below, and then combine it with Claim 8 above to prove the desired bound.

► Lemma 10. *Let C be a set-multilinear formula of product-depth $1 \leq \Delta \leq \log d / \log \log d$ of size at most s which computes a polynomial that is set-multilinear with respect to the partition (X_1, \dots, X_d) where each $|X_i| = n$. Let $w \in \{k, -k\}^d$ be chosen uniformly at random. Then, we have*

$$\text{relrk}_w(C) \leq s \cdot 2^{-\frac{kd^{1/\Delta}}{20}}$$

with probability at least $1 - s \cdot d^{-\frac{d^{1/\Delta}}{12\Delta}}$.

Proof. We prove the statement by induction on Δ .

If $\Delta = 1$, then $C = C_1 + \dots + C_t$ where each C_i is a product of linear forms. So, for all $i \in [t]$, by Claim 7,

$$\text{relrk}_w(C_i) = \prod_{j=1}^d 2^{-\frac{1}{2}|w_j|} = 2^{-\frac{kd}{2}}$$

where in the last step, we used the observation that regardless of the choice of w , $|w_j| = k$ for all $j \in [n]$. Hence, by the sub-additivity of relrk_w , with probability 1, we have

$$\text{relrk}_w(C) \leq s \cdot 2^{-\frac{kd}{2}} \leq s \cdot 2^{-\frac{kd}{20}}.$$

Next, we assume the statement is true for all formulas of product-depth $\leq \Delta$. Let C be a formula of product-depth $\Delta + 1$. So, C is of the form $C = C_1 + \dots + C_t$. Following an overall proof strategy similar to the one in [20], we say that a sub-formula C_i of size s_i is of type 1 if one of its factors has degree at least $T_\Delta = d^{\frac{\Delta}{\Delta+1}}$, otherwise we say it is of type 2.

⁶ We only need d to be larger than some absolute constant.

Suppose $C_i = C_{i,1} \cdots C_{i,t_i}$ is of type 1 with, say, $C_{i,1}$ having degree at least T_Δ . Let $w^{i,1}$ be the corresponding word i.e., $w^{i,1} = w|_{S_1}$ if $C_{i,1}$ is set-multilinear with respect to $S_1 \subsetneq [d]$. If it has size $s_{i,1}$, then since it has product-depth at most Δ , it follows by induction that

$$\text{relrk}_w(C_i) \leq \text{relrk}_{w^{i,1}}(C_{i,1}) \leq s_{i,1} \cdot 2^{-\frac{kT_\Delta^{1/\Delta}}{20}} \leq s_i \cdot 2^{-\frac{kd^{1/(\Delta+1)}}{20}}$$

with probability at least

$$1 - s_{i,1} \cdot T_\Delta^{-\frac{T_\Delta^{1/\Delta}}{12\Delta}} \geq 1 - s_i \cdot d^{-\frac{d^{1/(\Delta+1)}}{12\Delta} \cdot \frac{\Delta}{\Delta+1}} = 1 - s_i \cdot d^{-\frac{d^{1/(\Delta+1)}}{12(\Delta+1)}}.$$

Now suppose that $C_i = C_{i,1} \cdots C_{i,t_i}$ is of type 2 i.e., each factor $C_{i,j}$ has degree $< T_\Delta$. Note that this forces $t_i > d/T_\Delta = d^{\frac{1}{\Delta+1}}$. As the formula is set-multilinear, (S_1, \dots, S_{t_i}) form a partition of $[d]$ where each $C_{i,j}$ is set-multilinear with respect to $(X_\ell)_{\ell \in S_j}$ and C_i is set-multilinear with respect to $(X_\ell)_{\ell \in S}$. Let $w^{i,1}, \dots, w^{i,t_i}$ be the corresponding decomposition, whose respective sums are denoted simply by $w_{S_1}, \dots, w_{S_{t_i}}$.

From the properties of relrk_w (Claim 7), we have

$$\text{relrk}_w(C_i) = \prod_{j=1}^{t_i} \text{relrk}_{w^{i,j}}(C_{i,j}) \leq \prod_{j=1}^{t_i} 2^{-\frac{1}{2}|w_{S_j}|} = 2^{-\frac{1}{2} \sum_{j=1}^{t_i} |w_{S_j}|},$$

from which we observe that the task of upper bounding $\text{relrk}_w(C)$ can be reduced to the task of lower bounding the sum $\sum_{j=1}^{t_i} |w_{S_j}|$, which is established in the following claim. For the sake of convenience, the choice of the alphabet for w below is scaled down to $\{-1, 1\}$.

▷ **Claim 11.** For large enough d , suppose (S_1, \dots, S_ℓ) is a partition of $[d]$ such that each $|S_j| < T_\Delta = d^{\frac{\Delta}{\Delta+1}}$. Then, we have

$$\mathbb{P}_{w \sim \{-1, 1\}^d} \left[\sum_{j=1}^{\ell} |w_{S_j}| < \frac{d^{1/(\Delta+1)}}{10} \right] \leq d^{-\frac{d^{1/(\Delta+1)}}{12}}.$$

Proof. We first show that without loss of generality, we may assume that each S_j has size “roughly” T_Δ . To see this, we apply the following *clubbing* procedure to the sets in the partition (S_1, \dots, S_ℓ) :

- Start with the given partition (S_1, \dots, S_ℓ) . At each step in the procedure, we shall “club” two of the sets in the partition according to the following rule.
- If there are two distinct sets S' and S'' in the current partition each of size $< T_\Delta/2$, we remove both of them and add their union $S' \cup S''$ to the partition.
- If the rule above is no longer applicable, then we have at most one set in the current partition of size $< T_\Delta/2$. If there is none, then we halt the procedure here. Otherwise, we union this set with any one of the other sets and then halt.

After the clubbing procedure, we are left with a partition $(S'_1, \dots, S'_{\ell'})$ of $[d]$ such that $\frac{T_\Delta}{2} \leq |S'_j| \leq \frac{3T_\Delta}{2}$ for each $j \in [\ell']$, also implying that $\frac{2d^{1/(\Delta+1)}}{3} \leq \ell' \leq 2d^{1/(\Delta+1)}$. Through a repeated use of the triangle inequality, we see that $\sum_{j=1}^{\ell'} |w_{S'_j}| \leq \sum_{j=1}^{\ell} |w_{S_j}|$. Therefore, upper bounding the latter sum is a “smaller” event than upper bounding the former sum. Hence, it suffices to prove the statement of the claim with the assumption that $\frac{T_\Delta}{2} \leq |S_j| \leq \frac{3T_\Delta}{2}$ for each $j \in [\ell]$ (we henceforth drop the primed notation).

38:10 Improved Low-Depth Set-Multilinear Circuit Lower Bounds

Now, in the event that the sum $\sum_{j=1}^{\ell} |w_{S_j}|$ is at most $\frac{d^{1/(\Delta+1)}}{10}$, since $\ell \geq \frac{2d^{1/(\Delta+1)}}{3}$, it follows that for at least half of the sets S_j , $w_{S_j} = 0$ (as $\frac{2}{3} - \frac{1}{10} = \frac{17}{30} > \frac{1}{2}$). By Stirling's approximation, it follows that for a fixed j , the probability

$$\mathbb{P}_{w \sim \{-1,1\}^d} [w_{S_j} = 0] \leq \sqrt{\frac{2}{\pi |S_j|}} \leq \sqrt{\frac{4}{\pi T_{\Delta}}} = \sqrt{\frac{4}{\pi}} \cdot \frac{1}{d^{\frac{1}{2(\Delta+1)}}} < \frac{2}{d^{1/3}},$$

where in the final step, we used $\Delta \geq 2$. Therefore, the probability that this happens for $\ell/2$ distinct j is bounded by

$$\binom{\ell}{\ell/2} \cdot \left(\frac{2}{d^{1/3}}\right)^{\ell/2} < 2^{\ell} \cdot \left(\frac{2}{d^{1/3}}\right)^{\ell/2} = \left(\frac{2\sqrt{2}}{d^{1/6}}\right)^{\ell} \leq \left(\frac{2}{d^{1/9}}\right)^{\ell} < d^{-\frac{d^{1/(\Delta+1)}}{12}},$$

where we used the bound $\ell \geq \frac{2d^{1/(\Delta+1)}}{3}$. \triangleleft

The claim above and the preceding calculation immediately implies that for a sub-formula C_i of type 2,

$$\text{relrk}_w(C_i) \leq s_i \cdot 2^{-\frac{kd^{1/(\Delta+1)}}{20}}$$

with probability at least $1 - d^{-\frac{d^{1/(\Delta+1)}}{12}} \geq 1 - s_i \cdot d^{-\frac{d^{1/(\Delta+1)}}{12(\Delta+1)}}$.

Next, by a union bound over $i \in [t]$ and the sub-additivity property of relrk_w , it follows that

$$\text{relrk}_w(C) \leq \text{relrk}_w(C_1) + \dots + \text{relrk}_w(C_t) \leq s_1 \cdot 2^{-\frac{kd^{1/(\Delta+1)}}{20}} + \dots + s_t \cdot 2^{-\frac{kd^{1/(\Delta+1)}}{20}} = s \cdot 2^{-\frac{kd^{1/(\Delta+1)}}{20}}$$

with probability at least $1 - s \cdot d^{-\frac{d^{1/(\Delta+1)}}{12(\Delta+1)}}$, which concludes the proof of the lemma. \blacktriangleleft

Returning to the proof of the theorem, let C be a set-multilinear formula of product depth Δ of size s computing $NW_{n,d}(X)$. Suppose $s < d^{\frac{d^{1/\Delta}}{24\Delta}}$. Then, by Lemma 10, with probability at least $1 - d^{-\frac{d^{1/\Delta}}{24\Delta}}$,

$$\text{relrk}_w(C) \leq s \cdot 2^{-\frac{kd^{1/\Delta}}{20}}.$$

But now, we can condition on the event that $w_{[d]} = 0$ (which occurs with probability $\Theta(\frac{1}{\sqrt{d}})$) to establish the existence of a word $w \in \{-k, k\}^d$ with $w_{[d]} = 0$ such that w satisfies $\text{relrk}_w(C) \leq s \cdot 2^{-\frac{kd^{1/\Delta}}{20}}$. This is because of the asymptotic bound $\frac{1}{\sqrt{d}} \gg d^{-\frac{d^{1/\Delta}}{24\Delta}}$, which follows from the given constraints on the parameters d, Δ . Therefore, by Claim 8,

$$s \geq 2^{\frac{kd^{1/\Delta}}{20}} \cdot \text{relrk}_w(C) = n^{\frac{d^{1/\Delta}}{20}}$$

which contradicts the assumption that $s < d^{\frac{d^{1/\Delta}}{24\Delta}}$. Thus, we conclude that $s \geq d^{\frac{d^{1/\Delta}}{24\Delta}} = d^{\Omega(d^{1/\Delta}/\Delta)}$. \blacktriangleleft

Next, we show the supplementary result (Theorem 3) mentioned in Section 1, stated more precisely below.

► Theorem 12. *For an integer $n = 2^k$, let \mathbb{F}_n be a field of size n and suppose $d \leq n$ is large enough. Let X_i denote the set of n variables $\{x_{i,j} : j \in [n]\}$ and X be the tuple (X_1, \dots, X_d) . Then, any set-multilinear formula family computing $NW_{n,d}(X)$ must have size at least $d^{\Omega(\log d)}$.*

Proof. We first need the following structural result, whose proof can be immediately extrapolated from [28] (see Lemma 13.3), where it is shown for multilinear and homogeneous formulas.

► **Lemma 13** (Product Lemma). *Assume that F is a formula with at most s leaves, and is set-multilinear with respect to the set partition (X_1, \dots, X_d) . Then, we can write*

$$F = \sum_{i=1}^s \prod_{j=1}^{\ell} F_{i,j}$$

where $\ell \geq \log_3 d$ and for each $i \in [s]$, the product $F_i = \prod_{j=1}^{\ell} F_{i,j}$ is also set-multilinear. Furthermore, the degrees of $F_{i,j}$ satisfy the following geometric decay property:

$$\left(\frac{1}{3}\right)^j d \leq \deg(F_{i,j}) \leq \left(\frac{2}{3}\right)^j d, \text{ and } \deg(F_{i,\ell}) = 1.$$

► **Lemma 14.** *Let F be a set-multilinear formula of size at most s which computes a polynomial that is set-multilinear with respect to the partition (X_1, \dots, X_d) where each $|X_i| = n$. Let $w \in \{k, -k\}^d$ be chosen uniformly at random. Then, we have*

$$\text{relrk}_w(C) \leq s \cdot 2^{-\frac{k \log d}{20}}$$

with probability at least $1 - s \cdot d^{-\frac{\log d}{60}}$.

Proof. We begin by writing F in the form that is given by Lemma 13. Now, because of the geometric decay of the degrees of $F_{i,j}$, we observe that for each $i \in [s]$, at least for the first $\frac{3\ell}{4}$ many values of j , $\deg(F_{i,j}) \geq d^{1/4}$. In other words, at least a *constant* fraction of the $F_{i,j}$ s have their degrees at least *polynomially large* in d . This observation will be instrumental in establishing the following claim, which is akin to Claim 11 used in the proof of Lemma 10.

▷ **Claim 15.** For large enough d , suppose (S_1, \dots, S_ℓ) is a partition of $[d]$ such that $(\frac{1}{3})^j d \leq |S_j| \leq (\frac{2}{3})^j d$ for all $j \in [\ell]$, and $|S_\ell| = 1$. Then, we have

$$\mathbb{P}_{w \sim \{-1,1\}^d} \left[\sum_{j=1}^{\ell} |w_{S_j}| < \frac{\log d}{10} \right] \leq d^{-\frac{\log d}{60}}.$$

Proof. Consider the given event that $\frac{\log d}{10}$ exceeds the sum $\sum_{j=1}^{\ell} |w_{S_j}|$. As $\ell \geq \frac{\log d}{\log 3} > \frac{5 \log d}{8}$, it follows that for at least half of the sets S_j , $w_{S_j} = 0$ (since $\frac{5}{8} - \frac{1}{10} = \frac{21}{40} > \frac{1}{2}$). By the observation above, it also follows that at least for $\frac{\ell}{4}$ many of the *first* $\frac{3\ell}{4}$ values of j , $w_{S_j} = 0$. But for a fixed such j , since $|S_j| \geq d^{1/4}$, the probability

$$\mathbb{P}_{w \sim \{-1,1\}^d} [w_{S_j} = 0] \leq \sqrt{\frac{2}{\pi |S_j|}} < \frac{1}{\sqrt{|S_j|}} \leq \frac{1}{d^{1/8}},$$

Therefore, the probability that this happens for $\ell/4$ distinct j amongst the first $\frac{3\ell}{4}$ values of j is bounded by

$$\binom{3\ell/4}{\ell/4} \cdot \left(\frac{1}{d^{1/8}}\right)^{\ell/4} < 2^{3\ell/4} \cdot \left(\frac{1}{d^{1/8}}\right)^{\ell/4} < \left(\frac{2}{d^{1/32}}\right)^{\ell} < d^{-\frac{\log d}{60}}. \quad \triangleleft$$

38:12 Improved Low-Depth Set-Multilinear Circuit Lower Bounds

By sub-additivity of relrk_w (Claim 7), we have

$$\text{relrk}_w(F) \leq \text{relrk}_w(F_1) + \cdots + \text{relrk}_w(F_s). \quad (1)$$

So, fix an $i \in [s]$. As the formula is set-multilinear, let (S_1, \dots, S_ℓ) be the partition of $[d]$ such that each $F_{i,j}$ is set-multilinear with respect to $(X_t)_{t \in S_j}$. Let $w^{i,1}, \dots, w^{i,\ell}$ be the corresponding decomposition, whose respective sums are denoted by $w_{S_1}, \dots, w_{S_\ell}$. Then, by Claim 15,

$$\text{relrk}_w(F_i) = \prod_{j=1}^{\ell} \text{relrk}_{w^{i,j}}(F_{i,j}) \leq \prod_{j=1}^{\ell} 2^{-\frac{1}{2}|w_{S_j}|} = 2^{-\frac{1}{2} \sum_{j=1}^{\ell} |w_{S_j}|} \leq 2^{-\frac{k \log d}{20}}$$

with probability at least $1 - d^{-\frac{\log d}{60}}$. Therefore, by a union bound over $i \in [s]$ and (1), we conclude that

$$\text{relrk}_w(F) \leq s \cdot 2^{-\frac{k \log d}{20}}$$

with probability at least $1 - s \cdot d^{-\frac{\log d}{60}}$. \blacktriangleleft

Returning to the proof of the theorem, let F be a set-multilinear formula of size s computing $NW_{n,d}$. Suppose $s < d^{\frac{\log d}{120}}$. Then, by Lemma 14, with probability at least $1 - d^{-\frac{\log d}{120}}$,

$$\text{relrk}_w(F) \leq s \cdot 2^{-\frac{k \log d}{20}}.$$

But now, we can condition on the event that $w_{[d]} = 0$ (which occurs with probability $\Theta(\frac{1}{\sqrt{d}})$) to establish the existence of a word $w \in \{-k, k\}^d$ with $w_{[d]} = 0$ such that w satisfies $\text{relrk}_w(F) \leq s \cdot 2^{-\frac{k \log d}{20}}$. This is because of the trivial asymptotic bound $\frac{1}{\sqrt{d}} \gg d^{-\frac{\log d}{120}}$. Therefore, again by Claim 8,

$$s \geq 2^{\frac{k \log d}{20}} \cdot \text{relrk}_w(F) = n^{\frac{\log d}{20}}$$

which contradicts the assumption that $s < d^{\frac{\log d}{120}}$. Thus, we conclude that $s \geq d^{\frac{\log d}{120}} = d^{\Omega(\log d)}$. \blacktriangleleft

4 Discussion and Open Problems

We conclude by mentioning some interesting directions for future work.

- The most interesting and natural question is to make the hard polynomial in our main result $\text{IMM}_{n,n}$. This would imply super-polynomial algebraic formula lower bounds. As far as we know, it is conceivable that our complexity measure could be used to prove the lower bound for the $\text{IMM}_{n,n}$ polynomial. While the relative rank of $\text{IMM}_{n,n}$ itself is low, there might be a suitable “restriction” of it such that for a randomly chosen $w \in \{-k, k\}^n$, with reasonably high probability the restriction has large rank. This could then be used to prove the lower bound for $\text{IMM}_{n,n}$ (using Lemma 10 or Lemma 14). The result from [20] also showed its lower bound for the IMM polynomial by first analyzing a suitable restriction of IMM (although unfortunately that very same restriction idea does not work for us; please see the discussion in the appendix). Perhaps an intermediate question is to make the hard polynomial computationally simpler, for instance to find any hard polynomial that lies in VP.

- Another interesting question is to prove an improved depth hierarchy theorem for constant-depth set-multilinear formulas. [20] shows a depth hierarchy theorem for low-depth set-multilinear formulas. However, since their lower bounds only hold for small degrees, the depth hierarchy theorem in [20] only gives a quasi-polynomial separation of successive product-depths. It would be very interesting to obtain exponential separations (which for instance have been shown for low-depth multilinear circuits in [4]) using our measure.
- Another interesting direction could be to obtain lower bounds for general set-multilinear circuits via improved depth reduction results. The work of Kumar, Oliveira, and Saptharishi ([17]) provides some insight in this context, which shows an improved depth reduction to product-depth Δ with a size blow-up of $N^{O(\Delta \cdot (N/\log N)^{1/\Delta})}$ for multilinear circuits (regardless of degree). If a similar improvement (or any asymptotic improvement in the exponent) on the bound for general circuits from [30] could be shown to hold for set-multilinear circuits in the setting of Theorem 1 or Theorem 9 (i.e., when $N \geq d^2$), then combined with our lower bounds, this would imply super-polynomial set-multilinear circuit lower bounds. We should note that [7] rules out the possibility of obtaining a stronger reduction to depth-4, or $\Sigma\Pi\Sigma\Pi$ circuits, as it shows an $n^{\Omega(\sqrt{n})}$ size lower bound for set-multilinear depth-4 circuits computing $\text{IMM}_{n,n}$, which of course has small polynomial-sized set-multilinear circuits. Nevertheless, there is still the possibility of obtaining improved depth reduction statements for product-depths 2 (which as noted earlier, is $\Sigma\Pi\Sigma\Pi\Sigma$ and hence more general than depth-4) or higher, and combining it with our Theorem 1 to obtain unbounded-depth set-multilinear circuit lower bounds. [18] shows a quasi-polynomial separation between the strength of homogeneous $\Sigma\Pi\Sigma\Pi$ and $\Sigma\Pi\Sigma\Pi\Sigma$ circuits, which could be considered as some evidence towards the validity of this possibility.

References

- 1 Manindra Agrawal and V. Vinay. Arithmetic circuits: A chasm at depth four. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 67–75. IEEE Computer Society, 2008. doi:10.1109/FOCS.2008.32.
- 2 Miklós Ajtai. \sum^1_1 -formulae on finite structures. *Ann. Pure Appl. Log.*, 24(1):1–48, 1983. doi:10.1016/0168-0072(83)90038-6.
- 3 Peter Bürgisser. Cook’s versus valiant’s hypothesis. *Theor. Comput. Sci.*, 235(1):71–88, 2000. doi:10.1016/S0304-3975(99)00183-8.
- 4 Suryajith Chillara, Christian Engels, Nutan Limaye, and Srikanth Srinivasan. A near-optimal depth-hierarchy theorem for small-depth multilinear circuits. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 934–945. IEEE Computer Society, 2018. doi:10.1109/FOCS.2018.00092.
- 5 Suryajith Chillara, Nutan Limaye, and Srikanth Srinivasan. Small-depth multilinear formula lower bounds for iterated matrix multiplication with applications. *SIAM J. Comput.*, 48(1):70–92, 2019. doi:10.1137/18M1191567.
- 6 Zeev Dvir, Guillaume Malod, Sylvain Perifel, and Amir Yehudayoff. Separating multilinear branching programs and formulas. In Howard J. Karloff and Toniann Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 615–624. ACM, 2012. doi:10.1145/2213977.2214034.
- 7 Hervé Fournier, Nutan Limaye, Guillaume Malod, and Srikanth Srinivasan. Lower bounds for depth-4 formulas computing iterated matrix multiplication. *SIAM J. Comput.*, 44(5):1173–1201, 2015. doi:10.1137/140990280.

- 8 Merrick L. Furst, James B. Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Math. Syst. Theory*, 17(1):13–27, 1984. doi:10.1007/BF01744431.
- 9 Johan Håstad. Almost optimal lower bounds for small depth circuits. In Juris Hartmanis, editor, *Proceedings of the 18th Annual ACM Symposium on Theory of Computing, May 28-30, 1986, Berkeley, California, USA*, pages 6–20. ACM, 1986. doi:10.1145/12130.12132.
- 10 Pavel Hrubes and Amir Yehudayoff. Homogeneous formulas and symmetric polynomials. *Comput. Complex.*, 20(3):559–578, 2011. doi:10.1007/s00037-011-0007-3.
- 11 Neeraj Kayal, Nutan Limaye, Chandan Saha, and Srikanth Srinivasan. An exponential lower bound for homogeneous depth four arithmetic formulas. *SIAM J. Comput.*, 46(1):307–335, 2017. doi:10.1137/151002423.
- 12 Neeraj Kayal, Vineet Nair, and Chandan Saha. Separation between read-once oblivious algebraic branching programs (roabps) and multilinear depth-three circuits. *ACM Trans. Comput. Theory*, 12(1):2:1–2:27, 2020. doi:10.1145/3369928.
- 13 Neeraj Kayal, Chandan Saha, and Ramprasad Saptharishi. A super-polynomial lower bound for regular arithmetic formulas. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 146–153. ACM, 2014. doi:10.1145/2591796.2591847.
- 14 Neeraj Kayal, Chandan Saha, and Sébastien Tavenas. An almost cubic lower bound for depth three arithmetic circuits. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPIcs*, pages 33:1–33:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.ICALP.2016.33.
- 15 Neeraj Kayal, Chandan Saha, and Sébastien Tavenas. On the size of homogeneous and of depth-four formulas with low individual degree. *Theory Comput.*, 14(1):1–46, 2018. doi:10.4086/toc.2018.v014a016.
- 16 Pascal Koiran. Arithmetic circuits: The chasm at depth four gets wider. *Theor. Comput. Sci.*, 448:56–65, 2012. doi:10.1016/j.tcs.2012.03.041.
- 17 Mrinal Kumar, Rafael Oliveira, and Ramprasad Saptharishi. Towards optimal depth reductions for syntactically multilinear circuits. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, volume 132 of *LIPIcs*, pages 78:1–78:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ICALP.2019.78.
- 18 Mrinal Kumar and Ramprasad Saptharishi. Finer separations between shallow arithmetic circuits. In Akash Lal, S. Akshay, Saket Saurabh, and Sandeep Sen, editors, *36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2016, December 13-15, 2016, Chennai, India*, volume 65 of *LIPIcs*, pages 38:1–38:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.FSTTCS.2016.38.
- 19 Mrinal Kumar and Shubhangi Saraf. On the power of homogeneous depth 4 arithmetic circuits. *SIAM J. Comput.*, 46(1):336–387, 2017. doi:10.1137/140999335.
- 20 Nutan Limaye, Srikanth Srinivasan, and Sébastien Tavenas. Superpolynomial lower bounds against low-depth algebraic circuits. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 804–814. IEEE, 2021. doi:10.1109/FOCS52979.2021.00083.
- 21 Nutan Limaye, Srikanth Srinivasan, and Sébastien Tavenas. Set-multilinear and non-commutative formula lower bounds for iterated matrix multiplication. *To appear in STOC*, 2022. URL: <https://eccc.weizmann.ac.il/report/2021/094>.
- 22 Noam Nisan and Avi Wigderson. Hardness vs randomness. *J. Comput. Syst. Sci.*, 49(2):149–167, 1994. doi:10.1016/S0022-0000(05)80043-1.
- 23 Noam Nisan and Avi Wigderson. Lower bounds on arithmetic circuits via partial derivatives. *Comput. Complex.*, 6(3):217–234, 1997. doi:10.1007/BF01294256.

- 24 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.
- 25 Ran Raz. Tensor-rank and lower bounds for arithmetic formulas. *J. ACM*, 60(6):40:1–40:15, 2013. doi:10.1145/2535928.
- 26 Ran Raz and Amir Yehudayoff. Lower bounds and separations for constant depth multilinear circuits. *Comput. Complex.*, 18(2):171–207, 2009. doi:10.1007/s00037-009-0270-8.
- 27 Alexander A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematical notes of the Academy of Sciences of the USSR*, 41:333–338, 1987.
- 28 Ramprasad Satharishi. A survey of lower bounds in arithmetic circuit complexity. *GitHub Survey*, 2015. URL: <https://github.com/dasarpmar/lowerbounds-survey>.
- 29 Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In Alfred V. Aho, editor, *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 77–82. ACM, 1987. doi:10.1145/28395.28404.
- 30 Sébastien Tavenas. Improved bounds for reduction to depth 4 and depth 3. *Inf. Comput.*, 240:2–11, 2015. doi:10.1016/j.ic.2014.09.004.
- 31 Andrew Chi-Chih Yao. Separating the polynomial-time hierarchy by oracles (preliminary version). In *26th Annual Symposium on Foundations of Computer Science, Portland, Oregon, USA, 21-23 October 1985*, pages 1–10. IEEE Computer Society, 1985. doi:10.1109/SFCS.1985.49.

A Word Polynomials from [20, 21] and Our Measure

Both [20, 21] show their set-multilinear formula lower bounds for $\text{IMM}_{n,d}$ by showing that small enough set-multilinear formulas have low relative rank and that a certain “restriction” of $\text{IMM}_{n,d}$ has large relative rank. This restriction possesses the desirable property that if there is a small low-depth set-multilinear circuit computing $\text{IMM}_{n,d}$, then there is one for this restriction as well. It is then natural to wonder if we can use these same restrictions for our *symmetric* measure and deduce strong lower bounds for IMM (in order to show super-polynomial general formula lower bounds as discussed), in addition to obtaining them for the NW polynomial. Unfortunately, it is straightforward to show that this is not possible, as we shall now see.

► **Definition 16** (Word Polynomials of [20, 21]). *Let $w \in \mathbb{R}^d$ be any word with non-zero entries. Say $X(w) = (X_1, \dots, X_d)$ where each X_i has size $2^{|w_i|}$; we assume that the variables of X_i are labelled by strings in $\{0, 1\}^{|w_i|}$.*

Given any monomial $m \in \mathbb{F}_{\text{sm}}[\overline{X}(w)]$, let m_+ denote the corresponding “positive” monomial from \mathcal{M}_w^P and m_- the corresponding “negative” monomial from \mathcal{M}_w^N . As each variable of $\overline{X}(w)$ is labelled by a Boolean string and each set-multilinear monomial over any subset of $\overline{X}(w)$ is associated with a string of variables, we can associate any such monomial m' with a Boolean string $\sigma(m')$. More precisely, if $j_1 < \dots < j_t$ and $m' = x_{\sigma_1}^{(j_1)} x_{\sigma_1}^{(j_1)} \dots x_{\sigma_t}^{(j_t)}$ with $x_{\sigma_i}^{(j_i)} \in X_{j_i}$ and $\sigma_i \in \{0, 1\}^{|w_{j_i}|}$ for each $i \in [t]$, then $\sigma(m')$ is defined to be $\sigma_1 \dots \sigma_t$. We will write $\sigma(m_+) \sim \sigma(m_-)$ when the shorter one is a prefix of the other one. The polynomial P_w is defined as follows

$$P_w = \sum_{\substack{m \in \mathbb{F}[\overline{X}(w)], \\ \sigma(m_+) \sim \sigma(m_-)}} m.$$

Clearly, the matrices $\mathcal{M}_w(P_w)$ are full-rank (i.e., have rank equal to either the number of rows or the number of columns, whichever is smaller). So, $\text{relrk}_w(P_w) = 2^{-|w_{[d]}|/2}$.

38:16 Improved Low-Depth Set-Multilinear Circuit Lower Bounds

In our measure, $w \in \{k, -k\}^d$ with $w_{[d]} = 0$ i.e., there is an equal number of positive and negative variable sets and an equal number of variables $n = 2^k$ in each set. Thus, in the sum above, $\sigma(m_+) \sim \sigma(m_-)$ gets replaced with $\sigma(m_+) = \sigma(m_-)$. The sum is indexed over all Boolean strings of length $kd/2$, and so there are $n^{d/2}$ terms in all. Moreover, there is a canonical bijection between the positive and negative variables: since $|\mathcal{P}_w| = |\mathcal{N}_w| = d/2$, if an element $j \in \mathcal{P}_w$ is the k -th largest element in \mathcal{P}_w , it corresponds to the k -th largest element $j' \in \mathcal{N}_w$ such that $x_{i,j}$ appears in a monomial of P_w if and only if so does $x_{i,j'}$. Let $\phi : \mathcal{P}_w \rightarrow \mathcal{N}_w$ denote this correspondence. Then, we see that

$$P_w = \prod_{j \in \mathcal{P}_w} \sum_{i=1}^n x_{i,j} \cdot x_{i,\phi(j)},$$

implying that P_w actually has small depth-3 set-multilinear formulas.