Electronic Colloquium on Computational Complexity, Report No. 143 (2016)

An Almost Cubic Lower Bound for ΣΠΣ Circuits Computing a Polynomial in VP

Nikhil Balaji* Nutan Limaye† Srikanth Srinivasan‡

September 15, 2016

Abstract

In this note, we prove that there is an explicit polynomial in VP such that any ΣΠΣ arithmetic circuit computing it must have size at least $n^{3-o(1)}$. Up to $n^{o(1)}$ factors, this strengthens a recent result of Kayal, Saha and Tavenas (ICALP 2016) which gives a polynomial in VNP with the property that any ΣΠΣ arithmetic circuit computing it must have size $\tilde{\Omega}(n^3)$.

1 Introduction

Almost a decade ago, Agrawal and Vinay [AV08] showed that every low degree polynomial that can be computed by an arithmetic circuit of small size (which are exactly those polynomials in VP) can also be computed by a sub-exponential sized depth-4 circuit, which also has some nice structural properties. More efficient representations were given by Koiran [Koi12] and subsequently Tavenas [Tav15] gave an optimal construction. There are several nice aspects of these reduction to depth-4 circuits. Firstly, these results work for polynomials over any field. Secondly, the depth-4 circuit obtained thus is rather structured: it is a homogeneous ΣΠΣΠ circuit (see Section 2 for definitions). As already observed by Agrawal and Vinay, an important upshot of these *depth reduction* results is that they perhaps simplify the question of proving lower bounds for *arbitrary-depth* arithmetic circuits – that is, in order to separate VP from VNP, it suffices to prove strong exponential (the exact parameters are dictated by the strength of the parameters in the depth reduction) lower bounds for the aforementioned depth-4 circuits. Indeed, a series of papers building on the *shifted partial derivate method* of Kayal [Kay12], show exponential lower bounds (though not yet strong enough to show VP different from VNP) for restricted classes of depth-4 circuits [GKKS14, FLMS15, KS14, KLSS14, KS15].

Since any polynomial has a trivial representation as a depth-2 circuit as a sum of monomials (for which we can also prove strong lower bounds, trivially) and since already we have non-trivial depth reduction at depth-4, its natural to ask if depth-3 circuits, i.e. ΣΠΣ circuits, can simulate VP nontrivially. Here the state of affairs is more complicated. In an influential paper, Nisan and Wigderson [NW97] introduced the partial derivative method and proved an exponential lower bound for homogeneous ΣΠΣ circuits computing the elementary symmetric polynomial. They also note that by a simple interpolation idea attributed to Ben-Or, the elementary symmetric polynomials can be computed by a highly non-homogeneous ΣΠΣ circuit. As the elementary symmetric polynomials have small arithmetic circuits, any depth reduction to ΣΠΣ circuits cannot yield homogeneous circuits. A decade before the depth reduction theorems of Agrawal and Vinay, Grigoriev and Karpinski [GK98] proved that any depth-3 circuit computing the Determinant over a fixed finite

*Department of Computer Science and Engineering, IIT Bombay. nbalaji@cse.iitb.ac.in.
†Department of Computer Science and Engineering, IIT Bombay. nutan@cse.iitb.ac.in
‡Department of Mathematics, IIT Bombay. srikanth@math.iitb.ac.in.

ISSN 1433-8092

field requires exponential size. A natural question that arises from their work is to prove a similar lower bound for an explicit family of polynomials over characteristic zero. The best lower bound over large fields is the $\Omega(n^2)$ lower bound for the symmetric polynomial due to Shpilka and Wigderson [SW01]. So if at all there is a depth reduction to depth-3, it cannot work over arbitrary fields (since [GK98] precludes such bounds) and has to crucially exploit properties of the underlying large field. In a surprising result, Gupta, Kamath, Kayal and Saptharishi [GKKS16] proved such a non-trivial reduction to depth-3 over suitably large enough fields.

In light of these results, it is important to understand the power of non-homogeneous depth-3 circuits over characteristic zero. In a recent result, Kayal, Saha and Tavenas [KST16] prove a near *cubic* lower bound for $\Sigma\Pi\Sigma$ circuits, improving on the $\Omega(n^2)$-lower bound of [SW01]. In particular, they designed a polynomial family $\{f_n\}_{n\in\mathbb{N}}$ such that for every $n \in \mathbb{N}$ the polynomial f_n is a $\Theta(n)$-variate polynomial of degree $\Theta(n)$ and any depth three circuit computing the polynomial must have size $\Omega\left(\frac{n^3}{\log^2 n}\right)$. Their polynomial family was computable in VNP.

Here, we design a polynomial family $\{h_n\}_{n\in\mathbb{N}}$ computable in VP such that for every $n \in \mathbb{N}$, the polynomial h_n is a $\Theta(n \operatorname{poly}\log n)$-variate polynomial of degree $\Theta(n \operatorname{poly}\log n)$ and any depth three circuit computing the polynomial must have size $n^3/2^{O(\log^{2/3} n)}$. Therefore, our result is a more direct strengthening of the result of [SW01] and up to $n^{o(1)}$ factors, we are able to prove the same lower bound as [KST16] for a polynomial in VP.

Our main result can be stated more formally as follows:

Theorem 1. *There exists an explicit polynomial family $\{h_n\}_{n\in\mathbb{N}}$ computable in VP on $\Theta(n \operatorname{poly}\log n)$ variables and of degree $\Theta(n \operatorname{poly}\log n)$ (in fact by depth 6 circuits of polynomial size) such that any depth 3 circuit computing it has size $n^3/2^{O(\log^{2/3} n)}$.*

Both our construction and proof heavily build upon the ideas in [KST16].

2 Preliminaries

We recall some notions related to arithmetic circuits. See the surveys of Shpilka and Yehudayoff [SY10] and Saptharishi [Sap] for a more thorough treatment. Throughout this section, $\mathbb{F}$ will denote an arbitrary field.

Polynomial rings. We work over the standard multivariate polynomial ring $\mathbb{F}[x_1, \ldots, x_N]$. For a parameter $\ell \in \mathbb{N}$, we use $\mathcal{M}_\ell$ to denote the set of monomials of degree at most ℓ. For a parameter $k \in \mathbb{N}$ and $f \in \mathbb{F}[x_1, \ldots, x_N]$, we use $\partial_k(f)$ to denote the set of all kth order partial derivatives of f.

Shifted partial derivatives. Let $f(x_1, \ldots, x_N) \in \mathbb{F}[x_1, \ldots, x_N]$ be an arbitrary polynomial. Given parameters $k, \ell \in \mathbb{N}$, we define $\langle \partial_k f \rangle_{\leq \ell}$ as follows.

$$\langle \partial_k f \rangle_{\leq \ell} = \{m \cdot g \mid m \in \mathcal{M}_\ell, g \in \partial_k(f)\}.$$

We will use $\dim(\langle \partial_k f \rangle_{\leq \ell})$ to denote the dimension of the $\mathbb{F}$-linear span of $\langle \partial_k f \rangle_{\leq \ell}$. This will be our measure of complexity of the polynomial f.

The following is easy to prove.

Fact 2. *Let f be as above and L an affine change of co-ordinates. Then for any $k, \ell \in \mathbb{N}$, $g = f \circ L$ satisfies $\dim(\langle \partial_k g \rangle_{\leq \ell}) \leq \dim(\langle \partial_k f \rangle_{\leq \ell})$. In particular, if L is invertible, then $\dim(\langle \partial_k g \rangle_{\leq \ell}) = \dim(\langle \partial_k f \rangle_{\leq \ell})$.*

It is usually more convenient to lower bound $\dim(\langle \partial_k f \rangle_{\leq \ell})$ by considering only a subset S of $\partial_k(f)$, where S is chosen so that the derivatives in S satisfy a nice distance property that we now define. Given two multilinear monomials m, m' over variables $x_1, \ldots, x_N$, we can think of them as subsets of $\{x_1, \ldots, x_N\}$ and define $m \Delta m'$ analogously. The following lemma is implicit in [FLMS15, CM14].

Lemma 3. *Let $f \in \mathbb{F}[x_1, \ldots, x_N]$ be homogeneous and $k, \ell \in \mathbb{N}$ be arbitrary. Assume that we have $S \subseteq \partial_k(f)$ such that each $g \in S$ is a multilinear monomial and further that $|g\Delta g'| \geq 2\tau$ for distinct $g, g' \in S$. Then,* $\dim(\langle \partial_k f \rangle_{\leq \ell}) \geq |S|\binom{N+\ell}{\ell} - |S|^2\binom{N+\ell-\tau}{\ell-\tau}$.

Subspace restrictions. Let $f \in \mathbb{F}[x_1, \ldots, x_N]$ and V be an affine subspace of $\mathbb{F}^N$ of dimension m. The restriction $g = f|_V \in \mathbb{F}[Y_1, \ldots, Y_m]$ of f to the subspace V is defined as follows. We can find an invertible affine change of co-ordinates L of the space $\mathbb{F}^N$ so that $V = L^{-1}(\mathbb{F}^m \times \{0\}^{N-m})$. We let g be the polynomial obtained by setting $x_i = 0$ for all $i > m$ in the polynomial $f \circ L$. While L is not unique, changing the map L only affects g by a change of co-ordinates. In particular, $\dim(\langle \partial_k g \rangle_{\leq \ell})$ remains unchanged for any k, ℓ.

Circuits. We consider depth-3 and depth-4 arithmetic circuits of the following form. The circuits are layered (i.e. each gate occurs in a layer $i \in \{0, 1, 2, 3, 4\}$) and gates at layer i take as inputs (any number of) gates from layer $i-1$ and feed into gates at layer $i+1$. The largest layer contains a unique output gate. The 0th layer contains gates that compute variables and constants. The remaining layers alternate between product ($\times$) and sum ($+$) gates.[1] The largest layer contains the unique output gate (which may be a $+$ or a $\times$ gate). All gates are allowed to have unbounded fan-in. The size of the circuit is the number of *wires* in the circuit.

Depth-3 and depth-4 circuits with an output $+$ gate are referred to as $\Sigma\Pi\Sigma$ and $\Sigma\Pi\Sigma\Pi$ circuits respectively. For parameters $D, t \in \mathbb{N}$, we use $\Sigma\Pi^{(D)}\Sigma\Pi^{(t)}$ to denote $\Sigma\Pi\Sigma\Pi$ circuits where the $\times$ gates at layers 3 and 1 have fan-ins bounded by D and t respectively. In this case, the parameter t is called the *bottom fan-in* of the circuit.

The following was shown by [GKKS14].

Lemma 4. *Let C be a $\Sigma\Pi^{(D)}\Sigma\Pi^{(t)}$ circuit of size s computing a polynomial $f \in \mathbb{F}[x_1, \ldots, x_N]$. Then,* $\dim(\langle \partial_k f \rangle_{\leq \ell}) \leq s \cdot \binom{D}{k}\binom{N+\ell+k(t-1)}{\ell+k(t-1)}$.

3 A hard polynomial in VP

In this section, we define an explicit family of polynomials in VP for which we will prove the $\Sigma\Pi\Sigma$ circuit lower bound. In Section 3.1, we define an explicit family of polynomials and lower bound the dimension of the space of shifted partial derivatives of these polynomials. This builds upon the ideas of Kumar and Saraf [KS15]. In Section 3.2, we will use these polynomials to define our family of hard polynomials. This is done in a simple generic way that ensures that each "large enough" subspace restriction of the newly defined polynomials contains a copy of the polynomials of high shifted partial derivative complexity.

Both these steps are analogous to the construction of [KST16].

3.1 Polynomials with high shifted partial derivative complexity

We start with a variant of our construction. Let $k, t, n \in \mathbb{N}$ be arbitrary parameters (which we will fix later). For each $i \in [k]$, let X_i and $Y_{(t-1)(i-1)+1}, \ldots, Y_{(t-1)i}$ be t sets of n variables each. These tk sets of variables are pairwise disjoint. Let $X = \cup_{i\in[k]} X_i$ and $Y = \cup_{i\in[k], j\in[t-1]} Y_{(t-1)(i-1)+j}$.

We use $x_{i,p}$ and $y_{j,q}$ ($p, q \in [n]$) to denote the variables of X_i and Y_j respectively. Let $\mathtt{y}_{\mathtt{i},\mathtt{j}}$ denote the *monomial* $y_{(t-1)(i-1)+1,j} \cdot y_{(t-1)(i-1)+2,j} \cdots y_{(t-1)i,j}$.

$$\mathrm{SkewPGIP}'_{n,k,t}(X, Y) = \sum_{j_1, j_2, \ldots, j_k \in [n]} x_{1,j_1}\mathtt{y}_{1,j_1} \cdot x_{2,j_2}\mathtt{y}_{2,j_2} \cdots x_{k,j_k}\mathtt{y}_{k,j_k}$$

The polynomial is a product of a generalization of the inner product polynomial [NW97, KS15], it is multilinear and homogeneous. We call it skew because of the skewness in the number of X and Y variables. The polynomial can be computed by a $\Pi\Sigma\Pi$ circuit of size

[1] All our sum gates are allowed to compute arbitrary $\mathbb{F}$-linear combinations of their inputs.

$O(nkt)$ with bottom fan-in t. The total number of variables is nkt and the degree of the polynomial is kt.

We now define a version of the above polynomial such that it is on fewer variables. We reduce the number of variables by replacing each n-variable set of variables Y_j by a set Z_j of $C \log n$ fresh variables, where C is as defined below. We do this substitution while maintaining homogeneity and a certain distance property. Before describing the construction of the polynomial we state a standard combinatorial lemma which will be useful in the construction.

Lemma 5. *There exists constants $C_0 < C$ and an injective map $\phi : [n] \to 2^{[C \log n]}$ such that*

- *for each $j \in [n]$, let $\phi(j) = S_j \subseteq [C \log n]$ then $|S_j| = C_0 \log n$,*
- *for each $j \neq j' \in [n]$, $|S_j \Delta S_{j'}| \geq 2 \log n$.*

It is easy to see that such constants and map ϕ exist where the description of ϕ can be computed in time $n^{O(C)}$. The existence of such a map is implicit in, e.g., [KST16, Proposition 15].

Let $Z_i = \{z_{i,u} \mid u \in [C \log n]\}$. Given a ϕ as above, we will map each $y_{i,j}$ variable to a monomial in Z_i variables as follows: $y_{i,j} \longrightarrow \prod_{u \in \phi(j)} z_{i,u}$. With slight abuse of notation let us denote this monomial on Z_i variables by $\phi(y_{i,j})$. Similarly, let $\mathtt{z_{ij}}$ be the monomial $\phi(y_{(t-1)(i-1)+1,j}) \cdot \phi(y_{(t-1)(i-1)+2,j}) \cdots \phi(y_{(t-1)i,j})$.

Now we are ready to define our final polynomial.

$$\text{SkewPGIP}_{n,k,t}(X, Z) = \sum_{j_1, j_2, \ldots, j_k \in [n]} x_{1,j_1} \mathtt{z_{1,j_1}} \cdot x_{2,j_2} \mathtt{z_{2,j_2}} \cdots x_{k,j_k} \mathtt{z_{k,j_k}}$$

Note that the above is a polynomial on $N := n \cdot k + C \cdot \log n \cdot (t-1) \cdot k$ variables as opposed to nkt variables and its degree is $d := k + C_0 \cdot \log n \cdot (t-1) \cdot k$. Due to the first property of ϕ as stated in Lemma 5, the polynomial is multilinear and homogeneous. It is also computable by a polynomial sized $\Pi\Sigma\Pi$ circuit with bottom fan-in $O(t \log n)$.

We will need a lower bound on $\dim(\langle \partial_k \text{SkewPGIP}_{n,k,t}(X, Z) \rangle_{\leq \ell})$, proved below.

Lemma 6. $\dim(\langle \partial_k \text{SkewPGIP}_{n,k,t}(X, Z) \rangle_{\leq \ell}) \geq n^k \binom{N+\ell}{\ell} - n^{2k} \binom{N+\ell-\tau}{\ell-\tau}$, *where* $\tau = (t-1) \log n$.

Proof. We use Lemma 3 with a judicious choice of S. The proof is motivated by [KS15].

Let $J = (j_1, j_2, \ldots, j_k)$ denote a k tuple where each $j_i \in [n]$. Let $\partial_J(\cdot)$ denote the partial derivative with respect to the variables $x_{1,j_1}, x_{2,j_2} \ldots x_{k,j_k}$. Let $S = \{\partial_J(\text{SkewPGIP}_{n,k,t}(X, Z)) \mid J \in [n]^k\}$. Clearly, $|S| = n^k$.

From the definition of the polynomial it follows that for any fixed $J \in [n]^k$, $\partial_J(\text{SkewPGIP}_{n,k,t}(X, Z))$ is the multilinear monomial $\mathtt{z_{1,j_1}} \cdot \mathtt{z_{2,j_2}} \cdots \mathtt{z_{k,j_k}}$. Let us denote this monomial by m_J. For $J \neq J'$, there exists a $1 \leq i \leq k$ such that $j_i \neq j'_i$. And using the second condition in Lemma 5, $|m_J \Delta m'_J| \geq 2(t-1) \log n$.

Applying Lemma 3 now finishes the proof. □

3.2 An explicit polynomial hard under subspace restrictions

Now we will design a modified polynomial in VP which has a high shifted partial derivative measure even under subspace restrictions.

Let $g(V) = g(v_1, \ldots v_N)$ be any polynomial on N variables. We will now define another polynomial, $f_N(U, W)$ where $|U| := C'N \log N$ and $W = \{w_1, \ldots, w_{2N}\}$ with the following property. For every subset $A \subseteq [2N]$ of size N, there exists a 0-1 assignment a to the variables in U such that $f_N(U, W)|_{u \leftarrow a} = g(W_A) := g(w_{i_1}, \ldots, w_{i_N})$, where $A = \{i_1, \ldots, i_N\}$.

By Lemma 5, we know that there exists constants $C'_0 < C'$ and an injective map $\phi : [2N] \to 2^{[C' \log N]}$ such that for any $1 \leq i \leq [2N]$, $\phi(i) = T_i$ and $|T_i| = C'_0 \log N$. (Here, we will not need the second property of the map.)

Let $U_1, U_2, \ldots, U_N$ be pairwise disjoint sets of $C' \log N$ variables each and let $U = U_1 \cup U_2 \cup \ldots \cup U_N$. For each $1 \leq i \leq N$, we denote the elements of U_i by $\{u_{i,1}, u_{i,2}, \ldots, u_{i,C' \log N}\}$.

We define $\tilde{f}_N(U, W)$ on $C'N \log N + 2N$ variables as follows:

$$\tilde{f}_N(U,W) = g\left(\sum_{i=1}^{2N} w_i \prod_{j \in \phi(i)} u_{1,j}, \sum_{i=1}^{2N} w_i \prod_{j \in \phi(i)} u_{2,j}, \ldots, \sum_{i=1}^{2N} w_i \prod_{j \in \phi(i)} u_{N,j}\right)$$

Proposition 7. *Let $g(V) = g(v_1, \ldots v_N)$ be any polynomial on N variables. For any set $A \subseteq [2N]$ there exists a setting of variables in the set U such that $\tilde{f}_N(U, W) = g(W_A)$.*

Proof. Let $A = \{a_1, a_2, \ldots, a_N\}$. Set $u_{i,j} = 1$ if and only if $j \in \phi(a_i)$. It is easy to see that this encoding proves the proposition above. □

Note that the above construction of $\tilde{f}_N$ uses the polynomial g in a black-box way. Also, $\tilde{f}_N$ adds depth 2 to the depth of the circuit computing g and increases its size by a small polynomial in N.

Finally, by instantiating the above construction with g being the $\text{SkewPGIP}_{n,k,t}$ polynomial, we obtain the polynomial f_N for which we prove our lower bound. Let N be equal to $m_1 + m_2$, where $m_1 := n \cdot k$ and $m_2 := C \cdot \log n \cdot (t-1) \cdot k$. Let $V = X \cup Z$ and let us also assume a natural ordering on the V variables, say $V = \{v_1, \ldots, v_N\}$. Define

$$f_N(U,W) := \text{SkewPGIP}_{n,k,t}\left(\sum_{i=1}^{2N} w_i \prod_{j \in \phi(i)} u_{1,j}, \sum_{i=1}^{2N} w_i \prod_{j \in \phi(i)} u_{2,j}, \ldots, \sum_{i=1}^{2N} w_i \prod_{j \in \phi(i)} u_{N,j}\right).$$

Remark 8. *It follows quite straightforwardly from the above construction that f_N actually has a depth-6 homogeneous circuit of polynomial size.*

4 Setting of parameters and the lower bound

In this section, we prove the $\Sigma\Pi\Sigma$ circuit lower bound. As in [KST16], we first prove a $\Sigma\Pi\Sigma\Pi$ circuit lower bound with certain constraints on the fan-ins of the gates. We then use a restriction argument [SW01, KST16] to prove the main result.

Fact 9. *Let N, ℓ, τ be positive integers with $\tau \leq \ell$. Then, we have*

$$\binom{N+\ell}{\ell} \cdot \left(\frac{N+\ell+\tau}{\ell+\tau}\right)^{\tau} \leq \binom{N+\ell+\tau}{\ell+\tau} \leq \binom{N+\ell}{\ell} \cdot \left(\frac{N+\ell}{\ell}\right)^{\tau}.$$

Lemma 10. *Let $n, k, D, \epsilon, t \in \mathbb{N}$ be such that $\epsilon = 1/\log^{1/3} n$, $k = \epsilon \log n/2$, $N = n \cdot k + C \cdot \log n \cdot (t-1) \cdot k$ and $D = N^{1-\epsilon}$. [2] Then any $\Sigma\Pi^{(D)}\Sigma\Pi^{(t)}$ computing $\text{SkewPGIP}_{n,k,t}(X, Z)$ requires size $n^{\Omega(\log^{1/3} n)}$.*

Proof. Let s denote the size of a $\Sigma\Pi^{(D)}\Sigma\Pi^{(t)}$ circuit computing $\text{SkewPGIP}_{n,k,t}(X, Z)$. By Lemma 4, we then know that for any parameter $\ell \in \mathbb{N}$, we have

$$\dim(\langle \partial_k \text{SkewPGIP}_{n,k,t}(X,Z) \rangle_{\leq \ell}) \leq s \cdot \binom{D}{k}\binom{N+\ell+k(t-1)}{\ell+k(t-1)}. \tag{1}$$

We choose our parameter ℓ so as to obtain a lower bound on $\dim(\langle \partial_k \text{SkewPGIP}_{n,k,t}(X,Z) \rangle_{\leq \ell})$, which will then yield a lower bound on s.

By Lemma 6 we know that

$$\dim(\langle \partial_k \text{SkewPGIP}_{n,k,t}(X,Z) \rangle_{\leq \ell}) \geq n^k \binom{N+\ell}{\ell} - n^{2k}\binom{N+\ell-\tau}{\ell-\tau},$$

[2] To get a superpolynomial lower bound, the constant 1/3 in the definition of ϵ can be replaced by anything strictly smaller than 1/2.

where $\tau = (t-1)\log n$. We choose ℓ so that $\binom{N+\ell}{\ell}^\tau = \frac{n^k}{2}$. By Fact 9, this implies that $\dim(\langle \partial_k \text{SkewPGIP}_{n,k,t}(X,Z)\rangle_{\leq \ell}) \geq \frac{1}{2} n^k \binom{N+\ell}{\ell}$.

By (1), this implies

$$s \cdot \binom{D}{k} \cdot \binom{N+\ell+k(t-1)}{\ell+k(t-1)} \geq \frac{1}{2} \cdot n^k \cdot \binom{N+\ell}{\ell}$$

$$\therefore \quad s \geq n^{\epsilon k} \cdot \left(\frac{\ell}{N+\ell}\right)^{k(t-1)} \qquad \text{using Fact 9 and } D = N^{1-\epsilon}.$$

As we have set ℓ so that $\binom{N+\ell}{\ell}^\tau = \binom{N+\ell}{\ell}^{(t-1)\log n} = \frac{n^k}{2}$, we get that

$$s \geq \frac{1}{2} \cdot n^{\epsilon k - (k^2/\log n)} = n^{\Omega(\log^{1/3} n)}$$

where the final equality follows by our choice of ϵ and k. □

Remark 11. *Note that the above proof works for any value of t. Note also that the parameter t considered as the bottom fan-in of the circuit in the above lemma is the same as the parameter t used in defining the polynomial* $\text{SkewPGIP}_{n,k,t}$. *As long as $t \leq \text{poly}(n)$, $n^{\Omega(\log^{1/3} n)}$ is superpolynomial in N.*

4.1 Putting it together

In this section we prove Theorem 1.

Let $n \in \mathbb{N}$ be arbitrary. Let the parameters k, t, ϵ, N be chosen as in Lemma 10. Further we set $t = N/k \log n > N^{1-\epsilon}$. Note that this implies that $N = \Theta(nk)$.

The polynomial family that we prove this lower bound for is $\{h_n\}_{n\in\mathbb{N}}$, where $h_n := f_N(U,W)$ for f_N is as defined in Section 3.2. Let m denote the number of variables in $f_N(U,W)$ and Δ denote the degree of the polynomial. Note that $m = O(N \log N)$ and $\Delta = O(N \log N)$ and hence $\Delta = \Theta(m)$.

Now we prove that any $\Sigma\Pi\Sigma$ circuit computing h_n must have size $\Omega(n^3/2^{O(\log^{2/3} n)})$.

Proof. Let $\mathcal{C}$ denote the $\Sigma\Pi\Sigma$ circuit computing $f_N(U,W)$. Let $\mathcal{L}$ denote the set of (affine) linear functions computed by + gates at layer 1. We say that $L \in \mathcal{L}$ is *W-relevant* if it depends on some variable $w \in W$ (i.e., the coefficient of w in the linear function L is non-zero). The *W-degree* of a $\times$ gate at layer 2 is the number of W-relevant gates that feed into it.

Let a $\times$ gate in the circuit be called *heavy* if the W-degree of the gate is $> N^{2(1-\epsilon)}$ and *light* otherwise. Suppose there are $\geq N$ heavy $\times$ gates in the circuit then $\mathcal{C}$ has at least $N^{3-2\epsilon}$ wires and hence we are done.

Otherwise, the number of heavy $\times$ gates in the circuit is at most N. Then, as in [SW01, KST16], using at most N different affine restrictions we eliminate all the heavy gates as follows. As long as there is a heavy gate in the circuit, we pick one such gate arbitrarily. Call this gate P.

We choose an arbitrary W-relevant L feeding into P. Since L is W-relevant, we can write $L = \alpha w_i + L'(U, W \setminus \{w_i\})$ for $\alpha \neq 0$ and some $i \in [2N]$. We consider the restriction where $L(U,W) = 0$, which is equivalent to setting $w_i = -\frac{1}{\alpha} L'$. Substituting this into the circuit $\mathcal{C}$ eliminates the variable w_i from the circuit and moreover sets the product gate P to 0. Note that this may cause further simplifications, such as eliminating other heavy gates by either setting them to 0 or making them light.

After at most N such restrictions (effectively writing $N' \leq N$ variables $w_{i_1}, \ldots, w_{i_{N'}}$ as linear combinations of the other variables), the circuit $\mathcal{C}$ simplifies to a circuit $\mathcal{C}'$ such that each $\times$ gate in $\mathcal{C}'$ has W-degree at most $N^{2(1-\epsilon)}$. Let $A \subseteq [2N]$ be a set of size N such that the variables w_i ($i \in A$) are not restricted by the above process. By Proposition 7, we can set the variables in U to values in $\mathbb{F}$ so that the circuit $\mathcal{C}''$ now computes $\text{SkewPGIP}_{n,k,t}(W_A)$.

Having restricted the variables in U, this also ensures that the formal degree of each $\times$ gate in $\mathcal{C}''$ is at most $N^{2(1-\epsilon)}$.

Now, as in [KST16], we note that any such $\Sigma\Pi\Sigma$ circuit can then be converted into a $\Sigma\Pi^{(N^{1-\epsilon})}\Sigma\Pi^{(N^{1-\epsilon})}$ circuit $\mathcal{C}'''$ of the same size as $\mathcal{C}''$ (up to a constant multiplicative factor) as follows. We write each product gate in $\mathcal{C}''$ as a product of $N^{1-\epsilon}$ polynomials, each of which is in turn a product of at most $N^{1-\epsilon}$ linear forms.

But then, using Lemma 10 we know any $\Sigma\Pi^{(N^{1-\epsilon})}\Sigma\Pi^{(N^{1-\epsilon})}$ circuit computing $\mathrm{SkewPGIP}_{n,k,t}(W_A)$ must have superpolynomial size. This proves Theorem 1. □

References

[AV08] 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, 2008.

[CM14] Suryajith Chillara and Partha Mukhopadhyay. Depth-4 lower bounds, determinantal complexity: A unified approach. In *31st International Symposium on Theoretical Aspects of Computer Science (STACS 2014), STACS 2014, March 5-8, 2014, Lyon, France*, pages 239–250, 2014.

[FLMS15] 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.

[GK98] Dima Grigoriev and Marek Karpinski. An exponential lower bound for depth 3 arithmetic circuits. In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 577–582, 1998.

[GKKS14] Ankit Gupta, Pritish Kamath, Neeraj Kayal, and Ramprasad Saptharishi. Approaching the chasm at depth four. *J. ACM*, 61(6):33:1–33:16, 2014.

[GKKS16] Ankit Gupta, Pritish Kamath, Neeraj Kayal, and Ramprasad Saptharishi. Arithmetic circuits: A chasm at depth 3. *SIAM J. Comput.*, 45(3):1064–1079, 2016.

[Kay12] Neeraj Kayal. An exponential lower bound for the sum of powers of bounded degree polynomials. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:81, 2012.

[KLSS14] 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, 2014.

[Koi12] Pascal Koiran. Arithmetic circuits: The chasm at depth four gets wider. *Theor. Comput. Sci.*, 448:56–65, 2012.

[KS14] 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, 2014.

[KS15] Mrinal Kumar and Shubhangi Saraf. The limits of depth reduction for arithmetic formulas: It's all about the top fan-in. *SIAM J. Comput.*, 44(6):1601–1625, 2015.

[KST16] Neeraj Kayal, Chandan Saha, and Sébastien Tavenas. An almost cubic lower bound for depth three arithmetic circuits. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, pages 33:1–33:15, 2016.

[NW97] Noam Nisan and Avi Wigderson. Lower bounds on arithmetic circuits via partial derivatives. *Computational Complexity*, 6(3):217–234, 1997.

[Sap] Ramprasad Saptharishi. A survey of lower bounds in arithmetic circuit complexity. Github Survey.

[SW01] Amir Shpilka and Avi Wigderson. Depth-3 arithmetic circuits over fields of characteristic zero. *Computational Complexity*, 10(1):1–27, 2001.

[SY10] 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.

[Tav15] Sébastien Tavenas. Improved bounds for reduction to depth 4 and depth 3. *Inf. Comput.*, 240:2–11, 2015.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de