

Low-Depth Algebraic Circuit Lower Bounds over Any Field

Michael A. Forbes 

Department of Computer Science, University of Illinois at Urbana-Champaign, IL, USA

Abstract

The recent breakthrough of Limaye, Srinivasan and Tavenas [15] (LST) gave the first super-polynomial lower bounds against low-depth algebraic circuits, for any field of zero (or sufficiently large) characteristic. It was an open question to extend this result to small-characteristic ([8, 9, 16]), which in particular is relevant for an approach to prove superpolynomial $AC^0[p]$ -Frege lower bounds ([9]).

In this work, we prove super-polynomial algebraic circuit lower bounds against low-depth algebraic circuits over any field, with the same parameters as LST (or even matching the improved parameters of Bhargav, Dutta, and Saxena [3]). We give two proofs. The first is *logical*, showing that even though the proof of LST naively fails in small characteristic, the proof is sufficiently algebraic that generic transfer results imply the result over characteristic zero implies the result over all fields. Motivated by this indirect proof, we then proceed to give a second *constructive* proof, replacing the field-dependent set-multilinearization result of LST with a set-multilinearization that works over any field, by using the Binet-Minc identity [17].

2012 ACM Subject Classification Theory of computation → Algebraic complexity theory

Keywords and phrases algebraic circuits, lower bounds, low-depth circuits, positive characteristic

Digital Object Identifier 10.4230/LIPIcs.CCC.2024.31

Funding *Michael A. Forbes*: Supported by NSF CAREER award 2047310.

1 Introduction

Algebraic complexity asks for the best methods to compute multivariate polynomials $f(x_1, \dots, x_n)$ from the primitive operations of addition and multiplication, using constants from the underlying field \mathbb{F} for free. These computations are organized as an *algebraic circuit*, a directed acyclic graph of these primitive operations from the input variables x_1, \dots, x_n . The complexity of this object is primarily governed by its *size*, which is the number of operations in the circuit. Another complexity measure is the *product depth*, which is the maximum number of product operations on any input to output path. In a recent breakthrough, Limaye, Srinivasan and Tavenas [15] (LST) gave the first superpolynomial lower bounds for explicit polynomials to be computed by low-product-depth algebraic circuits, in *large characteristic*.¹

The main motivation for our work was to better understand the breakthrough work of Limaye, Srinivasan and Tavenas [15] (LST), as it left several open problems. In particular, Limaye, Srinivasan and Tavenas noted in their survey [16] the challenge of obtaining their lower bound over arbitrary fields; a problem this paper resolves. We briefly review the motivation for this challenge.

Algebraic Circuit Lower Bounds

Obtaining circuit lower bounds over *any* field is a fundamental aspect of algebraic complexity theory, as both the large and small characteristic settings are fundamentally interesting. These two regimes are somewhat incomparable in their difficulty. However there is the general

¹ For ease of exposition, we will write “large characteristic” to mean “large (or zero) characteristic”.

intuition that lower bounds in small characteristic “should be” easier, as there are “fewer” efficient algorithms to rule out. Following this intuition, a small characteristic analogue of LST should be achievable, as done in this paper.

To justify this intuition, here are examples of efficient algebraic algorithms known to exist in large characteristic which currently lack analogues in small characteristic. First, consider the Newton identities, which relate the elementary symmetric polynomials to the power-sum polynomials, but only in large characteristic. These identities are crucial in various algebraic complexity settings. For example, for constructing non-trivially small depth-4 formulas for the elementary symmetric polynomials ([20]) (as used by Limaye, Srinivasan and Tavenas [15]). Another example is a folklore construction of a polylogarithmic-depth polynomial-size algebraic circuit for the determinant (essentially the Faddeev–LeVerrier algorithm, expressed as an algebraic circuit), which uses traces of matrix powers to compute the power-sums of eigenvalues, and then computes the product of the eigenvalues (the determinant) from these power-sums using a small formula constructed from the Newton identities.

Another example is Fischer’s identity [5], which shows how to compute the monomial $x_1 \dots x_n$ as a homogeneous sum of powers of linear forms, but only in large characteristic. This identity was crucial to the celebrated depth-reduction of algebraic circuits to depth-3 ([13]).

A more computational example is the fundamental result in algebraic complexity that small algebraic circuits can be factored efficiently ([14]); however in characteristic p the result only produces a factorization up to p -th powers. The current inability (despite recent progress of Andrews [2]) to take p -th roots of algebraic circuits in characteristic p in particular leads to weaker hardness versus randomness trade-offs.

On the other hand, there *are* efficient algorithms that arise only in small characteristic and as such the above intuition is not quite correct. In particular, in characteristic p we have the identity $(x + y)^p = x^p + y^p$. In characteristic 2, we can compute the permanent efficiently (as $\det = \text{perm}$ over such fields) and also have a tighter connection between boolean and algebraic circuits. One can also view $n \times n$ Hadamard matrices as examples of this phenomenon, as in large characteristic they are full rank (and as such, “hard”), but in characteristic two then can have rank $O(\log n)$ (and as such, “easy”). Additionally, lifted Reed-Solomon codes ([12]) and William’s algorithm for k -path ([22]) are other techniques that only work in small characteristic.

Proof Complexity

Aside from interest in algebraic circuit lower bounds in small characteristic for their own sake, there are significant applications of such lower bounds.

A long-standing open question in proof-complexity is to prove superpolynomial lower bounds over constant-depth reasoning using modular gates, in particular the $\text{AC}^0[p]$ -Frege system. This challenge is notable as the boolean circuit complexity version has been solved by Razborov [19] and Smolensky [21], but the techniques have thus far not been successfully exported into the proof-complexity setting. In the context of this challenge, Grochow and Pitassi [11] showed that their Ideal Proof System (IPS) can efficiently simulate $\text{AC}^0[p]$ -Frege when IPS is over a field of characteristic p , and hence in particular that superpolynomial lower bounds for constant-depth IPS refutations of CNFs in characteristic p would give superpolynomial lower bounds against $\text{AC}^0[p]$ -Frege.

Toward this goal, Govindasamy, Hakoniemi, and Tzameret [9] showed how the lower bound of Limaye, Srinivasan and Tavenas [15] can yield a superpolynomial lower bound against (multilinear) constant-depth IPS refutations for (a variant of) the subset-sum problem, in large characteristic. While the subset-sum problem they use is *easy* in small characteristic

and hence our lower bounds cannot as is improve their results, our results eliminate one barrier to progress toward constant-depth IPS lower bounds in characteristic p , and hence $AC^0[p]$ -Frege lower bounds.

Polynomial Identity Testing

For applications within algebraic complexity itself, Limaye, Srinivasan and Tavenas [15] showed how their lower bound yields a deterministic subexponential time polynomial identity testing (PIT) algorithm for constant-depth algebraic circuits, in large characteristic. The restriction on characteristic comes from two places, the lower bound itself, as well as the known relations between algebraic hardness and derandomization. By removing one of these restrictions, our result hence makes progress toward obtaining corresponding PIT algorithms in small characteristic.

2 Our results

The lower bound of LST has two components. The first is a new set-multilinearization result for algebraic circuits, showing that algebraic circuits can be non-trivially set-multilinearized in a particular regime of parameters. The second component is a new lower bound for set-multilinear formulas. While the second component holds over any field, the first requires large characteristic. As such, our work focuses on this first component, which we now discuss more in depth.

Recall that for a partition of variables $\bar{x}_1 \sqcup \cdots \sqcup \bar{x}_d$, a *set-multilinear* monomial is one of the form $x_{1,i_1} \cdots x_{d,i_d}$. A set-multilinear polynomial is a linear combination of set-multilinear monomials. Many important polynomials are set-multilinear with respect to natural partitions of the variables. For example, the permanent of a matrix is set-multilinear with respect to the partition of the matrix into rows. For algebraic circuits computing set-multilinear polynomials it is natural to ask that the circuit respects the set-multilinear structure. In particular, we say that an algebraic circuit is (*syntactically*) *set-multilinear* if all product gates $f = f_1 \cdots f_k$ in the circuit have that the f_i are on disjoint parts of the variable partition. It follows then that set-multilinear circuits can only compute set-multilinear polynomials.

The first step of LST is a set-multilinearization result which transforms a low-depth algebraic circuit on n variables computing a degree d set-multilinear polynomial to one computing the same polynomial, where now the computation is itself set-multilinear. One pays for imposing this structure by increasing the circuit size. Crucially, the size only increases by a function of the degree d , *not* in the number of variables n .

► **Theorem 1** ([15]). *Let \mathbb{F} be a field of characteristic $\text{char}(\mathbb{F}) > d$ (or zero), where d is a parameter. Let f a set-multilinear polynomial of degree d , computed by a product-depth Δ circuit of size s . Then f is computed by a $(d^d \cdot s)^{O(1)}$ -size set-multilinear circuit of product-depth 2Δ .*

This result is proven in two steps. The first is to use an efficient low-depth homogenization transformation, that will double the product depth and increase the circuit size by $2^{O(\sqrt{d})}$. This construction uses that the characteristic is large, and is a generalization of the following result.

► **Theorem 2** ([20]). *Let \mathbb{F} be a field of characteristic $> d$, or zero. Then the elementary symmetric polynomial $\text{esym}_{n,d} = \sum_{S \in \binom{[n]}{d}} \prod_{i \in S} x_i$ has a homogeneous depth-4 sums of products of sums of powers $(\sum \prod \sum \wedge)$ formula of size $\text{poly}(n, 2^{\sqrt{d}})$.*

The standard connection between elementary symmetric polynomials and computing homogeneous parts allows the above theorem to homogenize depth-3 circuits into homogeneous depth-5 circuits with a $2^{O(\sqrt{d})}$ -blowup in circuit size. LST generalized Theorem 2 to *weighted* elementary symmetric polynomials to allow the idea to succeed in higher depths.

After converting product-depth Δ to homogeneous product-depth 2Δ , LST then convert the circuit to be set-multilinear while preserving the product-depth. This conversion will work over any field, and is a simple gate-simulation proof.

► **Theorem 3** ([15]). *Let \mathbb{F} be any field. Let f be a set-multilinear polynomial of degree d , computed by a homogeneous product-depth Δ circuit of size s . Then f is computed by a $\text{poly}(d^d, s)$ -size set-multilinear circuit of product-depth Δ .*

Combining the two steps of homogenization, then set-multilinearization, LST obtained the following.

► **Theorem 4** ([15]). *Let \mathbb{F} be a field of characteristic $\text{char}(\mathbb{F}) > d$ (or zero). Let f be a set-multilinear polynomial computed by a product-depth Δ circuit of size s . Then f can be computed by a $(d^d \cdot s)^{O(1)}$ -size set-multilinear circuit of product-depth 2Δ .*

The most natural method to obtaining the result of LST in arbitrary fields would be to give an efficient homogenization for low-depth circuits over all fields, and indeed this was posed as an open question in [8, 9, 16]. However, a barrier to this approach is that it is still open to develop an analogue of Theorem 2 in small characteristic fields.

Additionally, a recent work of Fournier, Limaye, Srinivasan, Tavenas [8] formalized this barrier, by showing that in small characteristic that a certain form of Newton identities cannot hold. This shows that while in large characteristic the Newton identities imply efficient low-depth homogenization for low-degree polynomials, there are provable barriers for obtaining an analogous result in small characteristic.

2.1 Lower bounds over any field, without explicit set-multilinearization

Our first result is to show that while we do not overcome this barrier, we never the less obtain the lower bound of LST over any field, which we state here incorporating the improved parameters from Bhargav, Dutta, and Saxena [3].

► **Theorem** (Main Theorem, Theorem 10). *Let \mathbb{F} be any field, and $d = o(\log n)$. Then the iterated matrix multiplication polynomial $\text{IMM}_{n,d} = (X_1 \cdots X_d)_{1,1}$ where X_i is an $n \times n$ symbolic matrix, requires*

$$n^{\Theta\left(d^{\frac{1}{F_{2\Delta+2} - 1}} / \Delta\right)}$$

size algebraic circuits of product depth Δ , where F_k is the k -th Fibonacci number (so $F_0 = 0$, $F_1 = F_2 = 1$, $F_3 = 2$, etc.).²

We observe that one can bypass the above barrier, and indeed the entire need for efficient low-depth homogenization, by arguing that the original techniques of LST already suffice for obtaining their result in low-characteristic, due to considerations from mathematical logic. That is, we study the *proof* of LST, and argue that the proof is sufficiently algebraic so that generic algebraic arguments imply the result holds over arbitrary fields.

² Bhargav, Dutta, and Saxena [3] use a slightly different indexing of Fibonacci numbers.

We abstract the methods of LST as follows, which is loosely inspired by Geometric Complexity Theory, as well as the theory of algebraic natural proofs ([7, 10]). Given a polynomial f of degree d in n variables, we can view the polynomial as a list of $N = \binom{n+d}{d}$ coefficients which we call the *coefficient vector* $\overline{\text{coeff}}(f)$. We then seek to construct a polynomial Q (or in fact a collection of polynomials Q_1, \dots, Q_M) in N variables such that $Q(\overline{\text{coeff}}(f)) = 0$ whenever f has a small low-depth algebraic circuit. At the same time, we want $Q(\overline{\text{coeff}}(P)) \neq 0$ (or in fact, $Q_i(\overline{\text{coeff}}(P)) \neq 0$ for some i) for some polynomial P . One can then conclude that P does not have a small low-depth algebraic circuit.

One can instantiate LST in this framework as follows. One can carefully rewrite the coefficients of polynomial f into a matrix C_f , and argue that the matrix C_f has low-rank. As matrix rank is characterized by the vanishing of determinant of submatrices, we can take the polynomials Q_i to be the relevant determinants. Any polynomial P whose associated matrix C_P has high-rank will have a non-vanishing determinant, and hence $Q_i(\overline{\text{coeff}}(P)) \neq 0$ for some i .

Many lower bounds techniques in algebraic complexity theory fall in the above “rank based” framework, and often in such proofs one proves that the matrix is high rank by arguing that there exists a large triangular submatrix whose diagonal entries are all 1. In such cases, the determinant of this submatrix is in fact 1, so the matrix is high-rank over *every* field. LST follows this approach, and as such this part of the framework does not depend on the characteristic.

Instead, the dependence on the characteristic comes into the argument that the rank of the matrices C_f is small. To show the rank is small over any field, we first note that the relevant determinants are in fact polynomials with *integer* coefficients, regardless of the actual field of consideration. Second, we note that small low-depth algebraic circuits have a *universal circuit* $U(\bar{x}, \bar{y})$ ([18]),³ such that $f = U(\bar{x}, \bar{\beta})$ for some constants $\bar{\beta}$ and $U(\bar{x}, \bar{y})$ has a small low-depth circuit. Further, U has *integer* coefficients. We then argue that viewing the field of computation as the characteristic zero field $\mathbb{Q}(\bar{y})$ (rational functions in the indeterminates \bar{y}), the matrix C_U must have low-rank, as the argument of LST applies. However, C_U having low-rank over $\mathbb{Q}(\bar{y})$ implies certain determinants of *integer* polynomials vanish, and thus also modulo p for any prime p . Hence C_U has low-rank over *any* field, and this will remain true even when we substitute $\bar{y} \leftarrow \bar{\beta}$. In particular, C_f has low rank as desired.

The overall idea is the standard fact from mathematical logic that if you want to prove a polynomial identity $A(\bar{x}) = 0$ where A has integer coefficients, then proving this identity in characteristic zero implies the result over every field because zero over the integers is zero modulo every prime p . A well-known example of this is the Cayley-Hamilton theorem, which states that every $n \times n$ matrix A is a root of its characteristic polynomial $p_A(x) = \det(xI - A)$. Viewing the entries of A as symbolic this can be viewed as a polynomial identity with integer coefficients, so suffices to prove it in characteristic zero, even if you use techniques specific to characteristic zero. The Cayley-Hamilton theorem can be proven in characteristic zero in two steps. First, one argues the theorem is true for all diagonalizable matrices, which is simple. Second, one argues that diagonalizable matrices are topologically dense in \mathbb{C} (say, in the Euclidean topology, a notion highly tied to characteristic zero) amongst all matrices, and hence by continuity the identity must also vanish on all matrices as desired. As such, one proves the Cayley-Hamilton theorem over all fields using techniques highly specific to characteristic zero.⁴

³ In the actual proof we do not need the universal circuit machinery and appeal to simpler arguments.

⁴ One can avoid the use of characteristic zero here using better algebraic techniques, such as Jordan normal form.

2.2 Lower bounds over any field, via explicit set-multilinearization

The results from the previous section showed that we were able to obtain the result of LST over any field *without* actually efficiently converting algebraic circuits to be set-multilinear. This is very suggestive that such a transformation should be achievable. While the barrier from above regarding efficient small-characteristic homogenization still is present, we observe that we can bypass this barrier (again) by *combining* the homogenization and set-multilinearization steps used by LST into a single transformation.

► **Theorem (Corollary 27).** *Let \mathbb{F} be any field. Let the variables \bar{x} be partitioned into $\bar{x} = \bar{x}_1 \sqcup \cdots \sqcup \bar{x}_d$. Let f be a set-multilinear polynomial computed by a product-depth Δ circuit of size s . Then f can be computed by a $(d^d \cdot s)^{O(1)}$ -size set-multilinear circuit of product-depth 2Δ .*

This result is proven by a standard gate simulation argument. The new component is to replace the use of the Newton identities by LST (which only work in large characteristic), by the Binet-Minc identity [17], which is a non-trivial depth-4 set-multilinear identity for computing (rectangular) permanents over any field. This is a natural step, as the Binet-Minc identity can, in large characteristic, be used (see Corollary 19) to recover the efficient depth-4 homogeneous formula for the elementary symmetric polynomial (Theorem 2) whose lack of small characteristic analogue is a barrier for LST holding in small characteristic.

Replacing the set-multilinearization of LST with the above result allows the rest of the proof of LST (and the improvement of Bhargav, Dutta, and Saxena [3]) to work over any field, giving another proof of our main result from above.

We note that that Binet-Minc identity has also recently been used in algebraic complexity by Curticapean, Limaye and Srinivasan [4] for unrelated reasons.

The above Corollary 27 is a more *constructive* method for proving the LST result in small characteristic, while the logical method is much more indirect. However we present the logical approach because it was the first proof we discovered, which motivated the constructive proof, and also the logical approach may have applications in other situations.

2.3 Related Work

An intriguing aspect of the work of LST is that it does not use the somewhat recent notion of shifted partial derivatives, which has powered numerous advances in algebraic circuit lower bounds in the past decade. Motivated by this, Amireddy, Garg, Kayal, Saha, Thankey [1] were able to essentially re-establish LST using shifted partial derivatives. They gave a novel analysis of shifted partials using several “imbalance” ideas related to LST, but crucially their analysis avoided the need to discuss set-multilinear polynomials and as such is perhaps more flexible than the set-multilinear methods of LST.

One of their main lower bounds was that any homogeneous product-depth Δ formula computing $\text{IMM}_{n,d}$ requires size $\geq n^{\Omega(d^{2^{1-\Delta}}/\Delta)}$, when $d \leq O(\lg n)$, over any field. Recall that LST showed, over fields of large characteristic, that general circuits computing degree d polynomials can be homogenized with a doubling in product-depth, and a $2^{\Omega(\sqrt{d})}$ blow-up in circuit size. Invoking this transformation, it follows that the results of Amireddy, Garg, Kayal, Saha, Thankey [1] establish super-polynomial lower bounds for general low-depth circuits, over fields of large characteristic. Quantitatively, the resulting bounds are in between those of LST and those of Bhargav, Dutta, and Saxena [3].

While this paper does not provide a low-depth homogenization transformation akin to that of LST over arbitrary fields, the logical methods of section 3 straightforwardly extend to the setting of Amireddy, Garg, Kayal, Saha, Thankey [1]. That is, the rank-based lower bound they establish for $\text{IMM}_{n,d}$ holds over all fields. The rank-based upper bound shows that all

low-depth general circuits are simple in characteristic zero, using LST's homogenization and their shifted partial analysis. One then can, as in section 3, generically transfer these two rank bounds to arbitrary fields, hence obtaining the lower bound for computing $\text{IMM}_{n,d}$ via low-depth circuits over an arbitrary field. We omit the straightforward details, in particular because the resulting parameters are worse than what are obtained in this work because we invoke the improved parameters from the set-multilinear lower bounds of Bhargava, Dutta, and Saxena [3].

3 Lower bounds over any field, via mathematical logic

In this section, we prove that the LST lower bound holds over any field, using techniques from mathematical logic that transfer the result from large characteristic to all characteristic. The following is our key lemma that transfers algebraic statements between different fields, in particular showing that an integer matrix having low rank in characteristic zero implies it has low-rank over every field. For our actual needs, we will need to consider matrices with entries that are polynomials with integer coefficients.

► **Lemma 5.** *Let $M \in \mathbb{Z}[\bar{x}]^{m \times n}$ be a matrix with integer polynomial entries. Let \mathbb{F} be any field, and interpret $M \in \mathbb{F}[\bar{x}]^{m \times n}$ via the unique non-trivial ring homomorphism $\varphi : \mathbb{Z} \rightarrow \mathbb{F}$. Then for any $\bar{\alpha} \in \mathbb{F}^{\bar{x}}$,*

$$\text{rank}_{\mathbb{F}} M(\bar{\alpha}) \leq \text{rank}_{\mathbb{F}(\bar{x})} M(\bar{x}) \leq \text{rank}_{\mathbb{Q}(\bar{x})} M(\bar{x}).$$

Proof.

1) $\text{rank}_{\mathbb{F}(\bar{x})} M(\bar{x}) \leq \text{rank}_{\mathbb{Q}(\bar{x})} M(\bar{x})$: In both cases the matrix M is the same, we simply change the field of interpretation. Recall that a matrix is $\text{rank} \leq s$ iff all $(s+1) \times (s+1)$ submatrices have zero determinant. Let $r = \text{rank}_{\mathbb{Q}(\bar{x})} M(\bar{x})$ so that all $(r+1) \times (r+1)$ submatrices of $M[\bar{x}]$ have zero determinant in $\mathbb{Q}(\bar{x})$. As M has polynomial entries, and the determinant is a polynomial, it follows that under the homomorphism φ that these submatrices still have zero determinant. As such, all $(s+1) \times (s+1)$ submatrices of M have zero determinant when viewed as a matrix in $\mathbb{F}[\bar{x}]^{m \times n}$, so that $\text{rank}_{\mathbb{F}(\bar{x})} M(\bar{x}) \leq r$.

2) $\text{rank}_{\mathbb{F}} M(\bar{\alpha}) \leq \text{rank}_{\mathbb{F}(\bar{x})} M(\bar{x})$: Let $t = \text{rank}_{\mathbb{F}(\bar{x})} M(\bar{x})$, so that all $(t+1) \times (t+1)$ submatrices of M have determinant zero in $\mathbb{F}[\bar{x}]$. Define the ring homomorphism $\psi : \mathbb{F}[\bar{x}] \rightarrow \mathbb{F}$ by $\bar{x} \rightarrow \bar{\alpha}$. As above, it then follows that all $(t+1) \times (t+1)$ submatrices of M have zero determinant under this homomorphism. But as $\psi(M(\bar{x})) = M(\bar{\alpha})$, it then follows that all such submatrices of $M(\bar{\alpha})$ have zero determinant in \mathbb{F} , so $\text{rank}_{\mathbb{F}} M(\bar{\alpha}) \leq t$. ◀

► **Definition 6.** *Let \mathbb{F} a field. Let \bar{x} be a set of variables with a partition $\bar{x} = \bar{y}_1 \sqcup \cdots \sqcup \bar{y}_{d_y} \sqcup \bar{z}_1 \sqcup \cdots \sqcup \bar{z}_{d_z}$, where $d = d_y + d_z$. Let Y denote the set of all set-multilinear monomials with respect to the partition $\bar{y} = \bar{y}_1 \sqcup \cdots \sqcup \bar{y}_{d_y}$, and let Z denote the set of all set-multilinear monomials with respect to the partition $\bar{z} = \bar{z}_1 \sqcup \cdots \sqcup \bar{z}_{d_z}$.*

*Given a polynomial $f(\bar{x})$ which is set-multilinear with respect to the above partition, define the **coefficient matrix** $C_f \in \mathbb{F}^{Y \times Z}$ by*

$$(C_f)_{\bar{y}^{\bar{b}}, \bar{z}^{\bar{c}}} = \text{Coeff}_{\bar{y}^{\bar{b}} \bar{z}^{\bar{c}}}(f),$$

where $\bar{y}^{\bar{b}}, \bar{z}^{\bar{c}}$ are set-multilinear monomials, and Coeff takes the coefficient of $\bar{y}^{\bar{b}} \bar{z}^{\bar{c}}$ in f .

The **relative rank** of f is then defined as

$$\text{relrank}_{\mathbb{F}}(f) := \frac{\text{rank}_{\mathbb{F}}(C_f)}{\sqrt{|Y| \cdot |Z|}}.$$

where $\text{rank}_{\mathbb{F}}(C_f)$ is the matrix rank of C_f over the field \mathbb{F} .

The following set-multilinearization result is a slight extension of what is proven in LST, and follows from their methods as noted by [9].

► **Definition 7.** Let the variables \bar{x} be partitioned into $\bar{x} = \bar{x}_1 \sqcup \cdots \sqcup \bar{x}_d$. A monomial is *set-multilinear (with respect to the partition $\bar{x} = \bar{x}_1 \sqcup \cdots \sqcup \bar{x}_d$)* if it can be written as $\prod_{i=1}^d (\bar{x}_i)_{j_i}$ for some j_1, \dots, j_d .

Define the *set-multilinear projection (with respect to the partition $\bar{x} = \bar{x}_1 \sqcup \cdots \sqcup \bar{x}_d$)* to be the linear map $\pi_{\text{sm}} : \mathbb{F}[\bar{x}] \rightarrow \mathbb{F}[\bar{x}]$ which is identity on set-multilinear monomials, and zero on all other monomials.

► **Theorem 8 ([9, 15]).** Let \mathbb{F} be a field of characteristic $\text{char}(\mathbb{F}) > d$ (or zero), where d is a parameter. Let f be a product-depth Δ circuit of size s . Let the variables \bar{x} be partitioned into $\bar{x} = \bar{x}_1 \sqcup \cdots \sqcup \bar{x}_d$. Then the set-multilinear projection $\pi_{\text{sm}}(f)$ has a $(d^d \cdot s)^{O(1)}$ -size set-multilinear circuit of product-depth 2Δ .

We quote here the summary of the LST lower bound results that we need, incorporating the improved parameters from Bhargav, Dutta, and Saxena [3].

► **Theorem 9 ([3, 15]).** Let \mathbb{F} a field. Let \bar{x} be a set of n variables, and $d \leq o(\log n)$ a parameter. Then there exists a partition $\bar{x} = \bar{y}_1 \sqcup \cdots \sqcup \bar{y}_{d_y} \sqcup \bar{z}_1 \sqcup \cdots \sqcup \bar{z}_{d_z}$, only depending on n and d , where $d = d_y + d_z$, such that any $f(\bar{x})$ computed by a set-multilinear circuit of size s and product-depth Δ has relative rank bounded by

$$\text{relrank}(f) \leq s \cdot n^{-\Theta\left(d^{\frac{1}{\Delta+2}-1} / \Delta\right)}.$$

Further, there exists a set-multilinear polynomial P with $\{0, 1\}$ -coefficients such that

$$\text{relrank}(P) \geq \frac{1}{n^{\Theta(1)}},$$

and P can be computed via evaluating the iterated matrix multiplication polynomial $\text{IMM}_{n,d}$ to carefully chosen linear forms.

We now give our characteristic-free version of the above.

► **Theorem 10.** Let \mathbb{F} be any field, and $d = o(\log n)$. Then $\text{IMM}_{n,d}$ requires

$$n^{\Theta\left(d^{\frac{1}{2\Delta+2}-1} / \Delta\right)}$$

size algebraic circuits of product depth Δ .

Proof. Let \bar{x} be a set of n variables. From Theorem 9, there exists a partition $\bar{x} = \bar{y}_1 \sqcup \cdots \sqcup \bar{y}_{d_y} \sqcup \bar{z}_1 \sqcup \cdots \sqcup \bar{z}_{d_z}$, only depending on n and d (and *not* on \mathbb{F}), where $d = d_y + d_z$, such that any $f(\bar{x})$ computed by a set-multilinear circuit of size s and product-depth Δ has relative rank bounded by

$$\text{relrank}(f) \leq s \cdot n^{-\Theta\left(d^{\frac{1}{\Delta+2}-1} / \Delta\right)}.$$

Further, there exists a set-multilinear polynomial P with $\{0, 1\}$ -coefficients such that

$$\text{relrank}(P) \geq \frac{1}{n^{\Theta(1)}},$$

and P can be computed via evaluating the iterated matrix multiplication polynomial $\text{IMM}_{n,d}$ to carefully chosen linear forms. In particular, any algebraic circuit lower bound for P extends, up to $\text{poly}(n)$ factors, to $\text{IMM}_{n,d}$, so it suffices to prove the lower bound for P .

Suppose P (interpreted as a polynomial in $\mathbb{F}[\bar{x}]$) is computed by an algebraic circuit Φ over \mathbb{F} of size s and product-depth Δ . Create a new algebraic circuit Ψ by replacing each field constant used in Φ with a distinct variable, so Ψ is size- s product-depth Δ algebraic circuit over the original variables \bar{x} along with new variables \bar{w} . Denote $f(\bar{x}, \bar{w})$ to be the polynomial computed by Ψ . As Ψ is defined free from field constants, f can be viewed as an integer polynomial $f \in \mathbb{Z}[\bar{x}, \bar{w}]$. Finally, we can relate P and f by undoing the above, so that there are values $\bar{\gamma}$ from \mathbb{F} such that $P(\bar{x}) = f(\bar{x}, \bar{\gamma})$.

Note that $f(\bar{x}, \bar{w})$ may not be set-multilinear, or even of particularly low-degree. However, it does follow that $\pi_{\text{sm}}(P(\bar{x})) = \pi_{\text{sm}}(f(\bar{x}, \bar{\gamma}))$ as the equality $P(\bar{x}) = f(\bar{x}, \bar{\gamma})$ is coefficient-wise, so applying π_{sm} to each side of equation either keeps the coefficient of P and f the same (if the monomial is set-multilinear) or makes both coefficients zero (if the monomial is not set-multilinear). As P is set-multilinear, we have $P = \pi_{\text{sm}}(P)$, so hence $P = \pi_{\text{sm}}(f(\bar{x}, \bar{\gamma}))$.

Now view Ψ as a circuit with constants over the field $\mathbb{Q}(\bar{w})$, so that $f \in \mathbb{Q}(\bar{w})[\bar{x}]$. It follows from Theorem 8 that $\pi_{\text{sm}}(f)$ has a $(d^d \cdot s)^{O(1)}$ -size set-multilinear circuit of product-depth 2Δ , and as such,

$$\text{relrank}_{\mathbb{Q}(\bar{w})}(\pi_{\text{sm}}(f(\bar{x}, \bar{w}))) \leq (d^d \cdot s)^{O(1)} \cdot n^{-\Theta\left(d^{\frac{1}{F_{2\Delta+2}^{-1}}}/\Delta\right)}.$$

Using that relative rank is just the (scaled) rank of a matrix, we can invoke Lemma 5, to see that

$$\text{relrank}_{\mathbb{F}}(\pi_{\text{sm}}(f(\bar{x}, \bar{\gamma}))) \leq \text{relrank}_{\mathbb{Q}(\bar{w})}(\pi_{\text{sm}}(f(\bar{x}, \bar{w}))).$$

As $P = \pi_{\text{sm}}(f(\bar{x}, \bar{\gamma}))$, we thus obtain that

$$\frac{1}{n^{\Theta(1)}} \leq \text{relrank}_{\mathbb{F}}(P) \leq (d^d \cdot s)^{O(1)} \cdot n^{-\Theta\left(d^{\frac{1}{F_{2\Delta+2}^{-1}}}/\Delta\right)}$$

which yields the desired lower bound for s (using that $d = o(\log n)$). ◀

4 Lower bounds over any field, constructively

In this section we give a constructive proof that any small low-depth algebraic circuit can be non-trivially set-multilinearized, over any field. As mentioned, by replacing the field-dependent set-multilinearization of LST with our set-multilinearization, this gives another proof of our main theorem (Theorem 10). The starting point for our construction is the *rectangular permanent*.

► **Definition 11.** Let X be an $n \times m$ symbolic matrix of variables, with $n \leq m$, where $X_{i,j} = x_{i,j}$ are distinct variables. Define the (**rectangular permanent**) $\text{perm}(X) \in \mathbb{Z}[(x_{i,j})_{i \in [n], j \in [m]}]$ by

$$\text{perm}_{n \times m}(X) = \sum_{\sigma: [n] \rightarrow [m]} x_{1, \sigma(1)} \cdots x_{n, \sigma(n)},$$

that is, the sum runs over all injective maps σ from $[n]$ to $[m]$.

Note in particular that $\text{perm}_{n \times m}$ is a set-multilinear polynomial when we partition the matrix into its rows.

It is sometimes helpful to view the rectangular permanent as a sum of square permanents.

31:10 Low-Depth Algebraic Circuit Lower Bounds over Any Field

► **Lemma 12.** *Let X be an $n \times m$ symbolic matrix of variables, with $n \leq m$, where $X_{i,j} = x_{i,j}$ are distinct variables. Then,*

$$\text{perm}_{n \times m}(X) = \sum_{S \in \binom{[m]}{n}} \text{perm}_{n \times n}(X|_{[n] \times S}).$$

The above lemma shows the rectangular permanent is computable by $\text{poly}(m^n)$ size algebraic circuits. However, the following identity gives a non-trivially better algorithm.

► **Theorem 13** (Binet-Minc Identity [17]). *Let X be an $n \times m$ symbolic matrix of variables, with $n \leq m$, then the permanent can be computed by*

$$\text{perm}_{n \times m}(X) = \sum_{\mathcal{F} \in \mathcal{P}_n} (-1)^{n-|\mathcal{F}|} \prod_{S \in \mathcal{F}} (|S| - 1)! \sum_{j=1}^m \prod_{i \in S} x_{i,j}$$

where \mathcal{P}_n is the set of all partitions of $[n]$ into (non-empty) sets, and $|\mathcal{F}|$ is the number of parts in the partition \mathcal{F} of $[n]$.

To understand the complexity of this expression it is helpful to have the following definition.

► **Definition 14.** *Define the n -th Bell number B_n to be the number of ways to partition $[n]$ into (non-empty) sets.*

We will use the following asymptotic estimate of Bell number size.

► **Fact 15** ([6]). $B_n = \Theta\left(\frac{n}{\ln n}\right)^n$.

The following lemma is easy to prove from the definition of Bell numbers.

► **Lemma 16.** $B_n \cdot B_m \leq B_{n+m}$.

The Binet-Minc identity immediately implies the following algebraic circuit for the permanent.

► **Corollary 17.** *The rectangular permanent $\text{perm}_{n \times m}$ has a $\sum^{B_n} \prod^n \sum^m \prod^n$ formula of size $\text{poly}(m, B_n)$, where the super-scripts are upper bounds on the respective fan-ins of the formula. Further, this formula is set-multilinear (and hence homogeneous) with respect to the partition of the $n \times m$ variables into rows.*

Note that for $m = n$ that this formula has complexity $\Theta\left(\frac{n}{\ln n}\right)^n$, whereas Ryser's formula is also set-multilinear but has size $\text{poly}(2^n)$ (and is depth-3).

We now note that when the rows of the matrix are identical, the Binet-Minc identity yields a small homogeneous depth-4 formula for the elementary symmetric polynomials, in large characteristic. The resulting formula has the same parameters as the argument of Shpilka and Wigderson [20], who proved it using the Newton identities. In particular, this relation is analogous to how Ryser's formula for the permanent, when applied to a matrix with identical rows, yields Fischer's depth-3 powering formula for the monomial.

► **Lemma 18.** *Let \bar{x} be n variables, and let Y be an $d \times n$ symbolic matrix of variables, with $d \leq n$, where $Y_{i,j} = x_j$. Then,*

$$\text{perm}_{d \times n}(Y) = d! \text{esym}_{n,d}(\bar{x})$$

Proof. $d = n$: This is immediate, as each monomial in the permanent now becomes $x_1 \cdots x_n$, and there are $n!$ many copies of this monomial.

$d < n$: Via Lemma 12,

$$\begin{aligned} \text{perm}_{d \times n}(Y) &= \sum_{S \in \binom{[n]}{d}} \frac{\text{perm}_{d \times d}(Y|_{[d] \times S})}{=d! \prod_{i \in S} x_i} \\ &= d! \text{esym}_{n,d}(\bar{x}). \end{aligned}$$

We now re-analyze the complexity of the above.

► **Corollary 19.** *Let \mathbb{F} be a field of characteristic $> d$ (or zero). The degree- d elementary symmetric polynomial in n variables $\text{esym}_{n,d}$ has a homogeneous $\sum^{2^{O(\sqrt{d})}} \prod^d \sum^n \wedge^d$ formula of size $\text{poly}(n, 2^{\sqrt{d}})$.*

Proof.

Construction, correctness: Apply the Binet-Minc identity (Theorem 13) to the $d \times n$ matrix Y where $Y_{i,j} = x_j$. By Lemma 18 this computes $d! \text{esym}_{n,d}$, and we can divide by $d!$ in \mathbb{F} .

Complexity: We now analyze the complexity of the above.

$$\begin{aligned} \text{perm}_{d \times n}(Y) &= \sum_{\mathcal{F} \in \mathcal{P}_d} (-1)^{d-|\mathcal{F}|} \prod_{S \in \mathcal{F}} (|S| - 1)! \sum_{j=1}^n \prod_{i \in S} \underbrace{Y_{i,j}}_{=x_j^{|S|}} \\ &= \sum_{\mathcal{F} \in \mathcal{P}_d} \underbrace{(-1)^{d-|\mathcal{F}|} \prod_{S \in \mathcal{F}} (|S| - 1)! \sum_{j=1}^n x_j^{|S|}}_{f_{\mathcal{F}}} \end{aligned}$$

noting that $f_{\mathcal{F}}$ only depends on the sizes of the how the partition \mathcal{F} refines the set $[n]$, we can group together the identical summands $f_{\mathcal{F}}$ based on how they partition the integer n ,

$$= \sum_{\bar{\lambda} \vdash d} (-1)^{d-|\bar{\lambda}|} N_{\bar{\lambda}} \prod_{\lambda \in \bar{\lambda}} (\lambda - 1)! \sum_{j=1}^n x_j^{\lambda}$$

where the summation runs over all integer partitions $\bar{\lambda}$ of d , $|\bar{\lambda}|$ is the number of parts in the partition $\bar{\lambda}$, and $N_{\bar{\lambda}} \in \mathbb{Z}$ is the number of set partitions of $[n]$ whose set sizes equal the integer partition $\bar{\lambda}$.

Now note that the above formula is homogeneous (as the Binet-Minc identity is), and the fan-ins of the formula are as desired because in particular the number of integer partitions of d is $2^{O(\sqrt{d})}$ ([6]). Finally, observe that the bottom-most product gate is a powering operation.

The above implies that Binet-Minc can recover that depth-3 circuits can be efficient homogenized into depth-5 circuits. As Binet-Minc holds over *all* fields and is additionally set-multilinear, this suggests that we can perhaps go directly to set-multilinearization, bypassing homogenization as an intermediate step. To do so, we need a slightly more general variant of the permanent.

► **Definition 20.** *Let X be an $n \times m$ symbolic matrix of variables, with $n \leq m$, where $X_{i,j} = x_{i,j}$ are distinct variables. Let $k \leq n$. Define the **k -surjective (rectangular) permanent** $\text{perm}_{n \times m; k}(X) \in \mathbb{Z}[(x_{i,j})_{i \in [n], j \in [m]}]$ by*

$$\text{perm}_{n \times m; k}(X) = \sum_{\sigma: [n] \hookrightarrow [m], \text{im}(\sigma) \supseteq [k]} x_{1, \sigma(1)} \cdots x_{n, \sigma(n)},$$

that is, the sum runs over all injective maps σ from $[n]$ to $[m]$ that contain $[k]$ in their image.

31:12 Low-Depth Algebraic Circuit Lower Bounds over Any Field

Note in particular that $\text{perm}_{n \times m; k}$ is a set-multilinear polynomial when we partition the matrix into its rows.

We can compute the surjective permanent by standard permanents.

► **Lemma 21.** *Let X be an $n \times m$ symbolic matrix of variables, with $n \leq m$, where $X_{i,j} = x_{i,j}$ are distinct variables. Let $k \leq n$. Then, the k -surjective permanent can be written as*

$$\text{perm}_{n \times m; k}(X) = \sum_{S \in \binom{[n]}{k}} \text{perm}_{k \times k}(X|_{S \times [k]}) \cdot \text{perm}_{(n-k) \times (m-k)}(X|_{([n] \setminus S) \times ([m] \setminus [k])}).$$

Further, this expression is set-multilinear with respect to partitioning X by its rows.

Proof. This follows from the observation that each map $\sigma : [n] \hookrightarrow [m]$ with $\text{im}(\sigma) \supseteq [k]$ uniquely decomposes into a bijection $\tau : S \leftrightarrow [k]$ for some $S \in \binom{[n]}{k}$, and an injection $\nu : ([n] \setminus S) \hookrightarrow ([m] \setminus [k])$.

The claim about set-multilinearity then follows from noting that the multiplication $\text{perm}_{k \times k}(X|_{S \times [k]}) \cdot \text{perm}_{(n-k) \times (m-k)}(X|_{([n] \setminus S) \times ([m] \setminus [k])})$ is a product of two polynomials who only use disjoint rows of X . ◀

We now analyze the complexity of computing the surjective permanent, by reduction to standard permanents, and then applying the Binet-Minc identity.

► **Corollary 22.** *Let $k \leq n \leq m$. The $n \times m$ k -surjective permanent has a $\text{poly}(m, \Theta(\frac{n}{\ln n})^n)$ -size depth-4 formula that is set-multilinear with respect to rows.*

Proof. Via the above lemma, and the formula complexity of the Binet-Minc identity,

$$\text{perm}_{n \times m; k}(X) = \sum_{S \in \binom{[n]}{k}} \frac{\text{perm}_{k \times k}(X|_{S \times [k]})}{\sum^{B_k} \prod^k \sum^k \prod^k} \cdot \frac{\text{perm}_{(n-k) \times (m-k)}(X|_{([n] \setminus S) \times ([m] \setminus [k])})}{\sum^{B_{n-k}} \prod^{n-k} \sum^{m-k} \prod^{n-k}}$$

distributing the multiplication past the addition,

$$= \frac{\sum_{\binom{[n]}{k}} \sum^{B_k \cdot B_{n-k}} \left(\prod^k \sum^k \prod^k \right) \cdot \left(\prod^{n-k} \sum^{m-k} \prod^{n-k} \right)}{\sum^{\binom{[n]}{k}} \prod^{B_k \cdot B_{n-k}} \prod^n \sum^m \prod^n}$$

from which the size bound follows by noting that $\binom{[n]}{k} B_k B_{n-k} \leq 2^n B_n \leq \Theta(\frac{n}{\ln n})^n$ via Lemma 16 and Fact 15.

The set-multilinearity of this formula follows from the fact that the decomposition used here from the above lemma is set-multilinear, that Binet-Minc is set-multilinear, and that our use of the distributive law preserves set-multilinearity. ◀

We now proceed to give a non-trivial set-multilinearization for low-depth algebraic circuits. To do so, it will be helpful to have more notation for extracting various set-multilinear components of polynomials.

► **Definition 23.** *Let the variables \bar{x} be partitioned into $\bar{x} = \bar{x}_1 \sqcup \dots \sqcup \bar{x}_d$.*

*Let $S \subseteq [d]$. A monomial is **S -set-multilinear** if it can be written as $\prod_{i \in S} (\bar{x}_i)_{j_i}$ for some $(j_i)_{i \in S}$. Define the **S -set-multilinear projection** to be the linear map $\pi_{\text{sm}, S}$ which is identity on S -set-multilinear monomials, and zero on all other monomials.*

*The set of S -set-multilinear monomials for some $S \subseteq [d]$ are called **at most set-multilinear monomials**. Define the **non-set-multilinear projection** to be the linear map $\pi_{\neg \text{sm}}$ which is zero on at-most set-multilinear monomials, and identity on all other monomials.*

In particular, for $S = [d]$, $\pi_{\text{sm}}(f)$ and $\pi_{\text{sm},S}(f)$ are the same. For $S = \emptyset$, $\pi_{\text{sm},S}(f)$ is the constant part of f , $\pi_{\text{sm},\emptyset}(f) = f(\bar{0})$. More generally, we can decompose f into its set-multilinear parts as follows.

► **Lemma 24.** *Let the variables \bar{x} be partitioned into $\bar{x} = \bar{x}_1 \sqcup \cdots \sqcup \bar{x}_d$. Then for any polynomial f ,*

$$f = \pi_{\neg\text{sm}}(f) + f(\bar{0}) + \sum_{\emptyset \neq S \subseteq [d]} \pi_{\text{sm},S}(f).$$

It immediately follows that we can simulate an addition (or even a linear combination) of polynomials by instead adding the constituent set-multilinear parts.

► **Lemma 25.** *Let f be any field. Let the variables \bar{x} be partitioned into $\bar{x} = \bar{x}_1 \sqcup \cdots \sqcup \bar{x}_d$. Let f_1, \dots, f_m be polynomials, with $f = \alpha_1 f_1 + \cdots + \alpha_m f_m$ for $\alpha_1, \dots, \alpha_m \in \mathbb{F}$. Let $S \subseteq [d]$. Then $\pi_{\text{sm},S}(f)$ can be computed by a depth-1 $\text{poly}(m, 2^d)$ -size set-multilinear \sum -circuit given $\{\pi_{\text{sm},S}(f_j)\}_{S \subseteq [d], j \in [m]}$ as inputs.*

Less trivially is the simulation of a multiplication, for which we use the formula for the surjective permanent (Corollary 22).

► **Lemma 26.** *Let the variables \bar{x} be partitioned into $\bar{x} = \bar{x}_1 \sqcup \cdots \sqcup \bar{x}_d$. Let f_1, \dots, f_m be polynomials, with $f = f_1 \cdots f_m$. Let $S \subseteq [d]$ be of size ℓ . Then $\pi_{\text{sm},S}(f)$ can be computed by a depth-4 $\text{poly}(m, \Theta(\frac{\ell}{\ln \ell})^\ell, 2^d)$ -size set-multilinear $\sum \Pi \sum \Pi$ -circuit given $\{\pi_{\text{sm},S}(f_j)\}_{S \subseteq [d], j \in [m]}$ as inputs.*

Proof. The number of inputs to the circuit is $m2^d$. It remains to bound the number of gates by $\text{poly}(m, \Theta(\frac{\ell}{\ln \ell})^\ell)$.

Rearrange the f_i as needed so that $f_1(\bar{0}) = \cdots = f_k(\bar{0}) = 0$ and $f_{k+1}(\bar{0}), \dots, f_m(\bar{0}) \neq 0$, for some $k \leq m$. Thus, we can normalize the computation via

$$f = \prod_{i>k} f_i(\bar{0}) \cdot \prod_{i \in [k]} f_i \cdot \prod_{i>k} \frac{f_i}{f_i(\bar{0})},$$

and thus define

$$g_i = \begin{cases} f_i & i \leq k \\ f_i/f_i(\bar{0}) & i > k \end{cases},$$

and $g = \prod_i g_i$. It then follows that $f = \prod_{i>k} f_i(\bar{0}) \cdot g$, $g_{k+1}(\bar{0}) = \cdots = g_m(\bar{0}) = 1$, and $\pi_{\text{sm},S}(f) = \prod_{i>k} f_i(\bar{0}) \cdot \pi_{\text{sm},S}(g)$. As $\prod_{i>k} f_i(\bar{0})$ is a non-zero constant, it suffices to prove the claim for $\pi_{\text{sm},S}(g)$.

Now write g as

$$g = \prod_{i=1}^k \left(\pi_{\neg\text{sm}}(g_i) + \sum_{\emptyset \neq S_i \subseteq [d]} \pi_{\text{sm},S_i}(g_i) \right) \cdot \prod_{i>k} \left(\pi_{\neg\text{sm}}(g_i) + 1 + \sum_{\emptyset \neq S_i \subseteq [d]} \pi_{\text{sm},S_i}(g_i) \right).$$

In expanding the above product, the only *at-most* set-multilinear terms are of the form

$$\prod_j \pi_{\text{sm},S_{i_j}}(g_{i_j}),$$

where the $S_{i_j} \subseteq [d]$ are disjoint (as otherwise we create non-multilinear terms), and the indices $\{i_j\}_j$ are distinct (as we take exactly one term from each g_i [possibly the term 1]).

31:14 Low-Depth Algebraic Circuit Lower Bounds over Any Field

Further, as $g_i(\bar{0}) = 0$ for $i \leq k$, we must have $\{i_j\}_j \supseteq [k]$. Hence, by collecting like terms, we see that

$$\pi_{\text{sm},S}(g) = \sum_{\mathcal{F} \in \mathcal{P}_S} \sum_{\sigma: \mathcal{F} \hookrightarrow [m]; \text{im}(\sigma) \supseteq [k]} \prod_{T \in \mathcal{F}} \pi_{\text{sm},T}(g_{\sigma(T)})$$

where \mathcal{P}_S is the collection of set partitions of S , which we can rewrite in terms of k -surjective permanents as

$$= \sum_{\mathcal{F} \in \mathcal{P}_S} \text{perm}_{|\mathcal{F}| \times m; k}(A_{\mathcal{F}}),$$

where for a partition \mathcal{F} of S the matrix $A_{\mathcal{F}}$ is $|\mathcal{F}| \times m$ size matrix defined by

$$(A_{\mathcal{F}})_{T,j} = \pi_{\text{sm},T}(g_j),$$

for $T \in \mathcal{F}$ and $j \in m$. Note that the matrix $A_{\mathcal{F}}$ has rows that access disjoint parts of the set-multilinear partition of \bar{x} , and so as the k -surjective permanent is set-multilinear with respect to rows, the computation $\text{perm}_{|\mathcal{F}| \times m; k}(A_{\mathcal{F}})$ is set-multilinear with respect to the partition of \bar{x} . Further, a set-multilinear computation of this permanent will be set-multilinear with respect to \bar{x} . Hence, we can use the set-multilinear computation from Corollary 22 for the surjective permanent, which in this case yields a depth-4 set-multilinear formula of size $\text{poly}(m, \Theta(\frac{|\mathcal{F}|}{\ln |\mathcal{F}|})^{|\mathcal{F}|})$.

Computing $\pi_{\text{sm},S}(g)$ is then a sum of $B_{|S|}$ many such $\text{perm}_{|\mathcal{F}| \times m; k}(A_{\mathcal{F}})$ terms, which does not increase depth as we collapse two sequential layers of addition gates, and this preserves set-multilinearity. The resulting size of the expression is $\text{poly}(m, \Theta(\frac{|\mathcal{F}|}{\ln |\mathcal{F}|})^{|\mathcal{F}|}, B_{|S|}, 2^d) \leq \text{poly}(m, \Theta(\frac{\ell}{\ln \ell})^{\ell}, 2^d)$, as $|S| = \ell$, $|\mathcal{F}| \leq |S|$. ◀

We now conclude with our set-multilinearization result over any field by gate-simulation.

► **Corollary 27.** *Let \mathbb{F} be any field. Let the variables \bar{x} be partitioned into $\bar{x} = \bar{x}_1 \sqcup \dots \sqcup \bar{x}_d$. Suppose $f \in \mathbb{F}[\bar{x}]$ can be computed by a size s product-depth Δ algebraic circuit. Then the set-multilinear projection $\pi_{\text{sm}}(f) \in \mathbb{F}[\bar{x}]$ can be computed by a size $\text{poly}(s, \Theta(\frac{d}{\ln d})^d)$ -size product-depth 2Δ circuit.*

Proof. By gate simulation. Let Φ be the hypothesized circuit computing f . For each node v in Φ , split v into its at-most set-multilinear parts $\pi_{\text{sm},S}(v)$ for each $S \subseteq [d]$.

If $v = \alpha_1 v_1 + \dots + \alpha_m v_m$, then we can express the at-most set-multilinear parts of v in terms of the v_i using a $\text{poly}(m, 2^d)$ -size product-depth 0 set-multilinear circuit by Lemma 25. If $v = v_1 \times \dots \times v_m$, then we can express the at-most set-multilinear parts of v in terms of the v_i using a $\text{poly}(m, \Theta(\frac{d}{\ln d})^d)$ -size product-depth 2 set-multilinear circuit by Lemma 26.

Correctness of the computation follows by induction on the circuit.

The overall size of the circuit has increased from s to $\text{poly}(s, \Theta(\frac{d}{\ln d})^d)$ by counting the size of the additional local gadgets of the gate simulation. Simulation of addition gates adds no product depth. Each product gate in the original circuit is turned into a product-depth 2 circuit in the gate simulation, hence the overall product-depth has at most doubled. ◀

References

- 1 Prashanth Amireddy, Ankit Garg, Neeraj Kayal, Chandan Saha, and Bhargav Thankey. Low-depth arithmetic circuit lower bounds: Bypassing set-multilinearization. In *Proceedings of the 50th International Colloquium on Automata, Languages and Programming (ICALP 2023)*, volume 261 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 12:1–12:20, 2023. Full version in the Electronic Colloquium on Computational Complexity (ECCC), Technical Report TR22-151. doi:10.4230/LIPICS.ICALP.2023.12.

- 2 Robert Andrews. Algebraic hardness versus randomness in low characteristic. In *Proceedings of the 35th Annual Computational Complexity Conference (CCC 2020)*, volume 169 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 37:1–37:32, 2020. Full version at [arXiv:2005.10885](https://arxiv.org/abs/2005.10885). doi:10.4230/LIPICS.CCC.2020.37.
- 3 C. S. Bhargav, Sagnik Dutta, and Nitin Saxena. Improved lower bound, and proof barrier, for constant depth algebraic circuits. In *Proceedings of the 47th International Symposium on the Mathematical Foundations of Computer Science (MFCS 2022)*, volume 241 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 18:1–18:16, 2022. doi:10.4230/LIPICS.MFCS.2022.18.
- 4 Radu Curticapean, Nutan Limaye, and Srikanth Srinivasan. On the VNP-hardness of some monomial symmetric polynomials. In *42nd International Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2022)*, volume 250 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 16:1–16:14, 2022. Full version in the Electronic Colloquium on Computational Complexity (ECCC), Technical Report TR22-139. doi:10.4230/LIPICS.FSTTCS.2022.16.
- 5 Ismor Fischer. Sums of like powers of multivariate linear forms. *Mathematics Magazine*, 67(1):59–61, 1994. URL: <http://www.jstor.org/stable/2690560>.
- 6 Philippe Flajolet and Robert Sedgewick. *Analytic combinatorics*. Cambridge University Press, 2009. doi:10.1017/CB09780511801655.
- 7 Michael A. Forbes, Amir Shpilka, and Ben Lee Volk. Succinct hitting sets and barriers to proving algebraic circuits lower bounds. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC 2017)*, pages 653–664, 2017. Full version at [arXiv:1701.05328](https://arxiv.org/abs/1701.05328). doi:10.1145/3055399.3055496.
- 8 Hervé Fournier, Nutan Limaye, Srikanth Srinivasan, and Sébastien Tavenas. On the power of homogeneous algebraic formulas. *Electronic Colloquium on Computational Complexity (ECCC)*, TR23-191, 2023. URL: <https://eccc.weizmann.ac.il/report/2023/191>.
- 9 Nashlen Govindasamy, Tuomas Hakoniemi, and Iddo Zameret. Simple hard instances for low-depth algebraic proofs. In *Preliminary version in the 63rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2022)*, pages 188–199, 2022. Full version at [arXiv:2205.07175](https://arxiv.org/abs/2205.07175). doi:10.1109/FOCS54457.2022.00025.
- 10 Joshua A. Grochow, Mrinal Kumar, Michael E. Saks, and Shubhangi Saraf. Towards an algebraic natural proofs barrier via polynomial identity testing. *arXiv*, 1701.01717, 2017. URL: <http://arxiv.org/abs/1701.01717>.
- 11 Joshua A. Grochow and Toniann Pitassi. Circuit complexity, proof complexity, and polynomial identity testing. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2014)*, pages 110–119, 2014. Full version at [arXiv:1404.3820](https://arxiv.org/abs/1404.3820). doi:10.1109/FOCS.2014.20.
- 12 Alan Guo, Swastik Kopparty, and Madhu Sudan. New affine-invariant codes from lifting. In *Proceedings of Innovations in Theoretical Computer Science (ITCS 2013)*, pages 529–540, 2013. Full version at [arXiv:1208.5413](https://arxiv.org/abs/1208.5413). doi:10.1145/2422436.2422494.
- 13 Ankit Gupta, Pritish Kamath, Neeraj Kayal, and Ramprasad Satharishi. Arithmetic circuits: A chasm at depth three. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2013)*, pages 578–587, 2013. Full version in the Electronic Colloquium on Computational Complexity (ECCC), Technical Report TR13-026. doi:10.1109/FOCS.2013.68.
- 14 Erich L. Kaltofen. Factorization of polynomials given by straight-line programs. In Silvio Micali, editor, *Randomness and Computation*, volume 5 of *Advances in Computing Research*, pages 375–412. JAI Press, Inc., Greenwich, CT, USA, 1989. URL: http://www.math.ncsu.edu/~kaltopen/bibliography/89/Ka89_slpfac.pdf.
- 15 Nutan Limaye, Srikanth Srinivasan, and Sébastien Tavenas. Superpolynomial lower bounds against low-depth algebraic circuits. In *Preliminary version in the 62nd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2021)*, pages 804–814, 2022. Full version in the Electronic Colloquium on Computational Complexity (ECCC), Technical Report TR21-081. doi:10.1109/FOCS52979.2021.00083.

31:16 Low-Depth Algebraic Circuit Lower Bounds over Any Field

- 16 Nutan Limaye, Srikanth Srinivasan, and Sébastien Tavenas. Guest column: Lower bounds against constant-depth algebraic circuits. *SIGACT News*, 53(2):40–62, 2022. doi:10.1145/3544979.3544989.
- 17 Henryk Minc. Evaluation of permanents. *Proc. Edinburgh Math. Soc. (2)*, 22(1):27–32, 1979. doi:10.1017/S0013091500027760.
- 18 Ran Raz. Elusive functions and lower bounds for arithmetic circuits. *Theory of Computing*, 6(1):135–177, 2010. 40th Annual ACM Symposium on Theory of Computing (STOC 2008). doi:10.4086/TOC.2010.V006A007.
- 19 Alexander A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Matematicheskie Zametki*, 41(4):598–607, April 1987.
- 20 Amir Shpilka and Avi Wigderson. Depth-3 arithmetic circuits over fields of characteristic zero. *Computational Complexity*, 10(1):1–27, 2001. Preliminary version in the 14th Annual IEEE Conference on Computational Complexity (CCC 1999). doi:10.1007/PL00001609.
- 21 Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing (STOC 1987)*, pages 77–82, 1987. doi:10.1145/28395.28404.
- 22 Ryan Williams. Finding paths of length k in $O^*(2^k)$ time. *Inf. Process. Lett.*, 109(6):315–318, 2009. doi:10.1016/J.IPL.2008.11.004.