Electronic Colloquium on Computational Complexity, Report No. 314 (2026)

Polynomial-size hitting sets for multilinear any-order read-once algebraic branching programs and for set-multilinear depth-3 circuits with constant-size variable sets

Amir Shpilka*

Abstract

In this paper we give polynomial-size hitting sets for multilinear any-order read-once oblivious algebraic branching programs and for depth-3 set-multilinear circuits where each variable set has constant size. Our results also extend to the border of both classes and hold over fields of any characteristic.

The proof combines the approach of Forbes, Saptharishi and Shpilka (STOC 2014) with the recent results of Bhardwaj and Saptharishi (2026) and of Bal and Ghosh (2026) that gave polynomial-size hitting sets for polynomials with small partial derivative dimension.

1 Introduction

Polynomial identity testing (PIT) is the problem of deciding whether a given algebraic circuit computes the zero polynomial. Efficient randomized algorithms follow from the polynomial identity lemma [Ore22, DL78, Sch80, Zip79], but obtaining a deterministic polynomial-time algorithm remains a major open problem. PIT has many algorithmic applications and is also closely connected to proving circuit lower bounds [HS80, KI04, Agr05]. We refer to the surveys [Sax09, SY10, Sax14, DG24] for background and applications of PIT.

PIT is studied in two settings. In the white-box setting, the algorithm is given the circuit as input. In the black-box setting, we can only evaluate the polynomial at points of our choice. Thus, giving a deterministic black-box PIT for a class of polynomials amounts to constructing an explicit hitting set. That is, a set of points such that every nonzero polynomial in the class is nonzero at one of them.

Due to its difficulty most research has focused on restricted models of algebraic computation. In particular, small depth circuits received a lot of attention. One reason is that depth reduction results imply that efficient PIT for depth-four, or even depth-three, circuits would yield PIT for general circuits [AV08, Koi12, Tav15, GKKS16]. Other natural models, such as sparse polynomials, read-once and bounded-read formulas, and the ROABPs discussed below, have also received much attention. For a more complete picture we refer to the surveys [Sax09, SY10, Sax14, DG24] mentioned above and to [KST26].

The main model that we consider is that of read-once oblivious algebraic branching programs (ROABPs for short). A width-w ROABP in the variable order $x_{\pi(1)}, \ldots, x_{\pi(n)}$ computes a polynomial of the form

$$f(x_1, \ldots, x_n) = \left(M_{\pi(1)}(x_{\pi(1)}) \cdots M_{\pi(n)}(x_{\pi(n)})\right)_{1,1}, \tag{1}$$

where each $M_i(x_i) \in \mathbb{F}[x_i]^{w \times w}$ is a square matrix of dimension w.

*Blavatnik School of Computer Science and AI, and School of Mathematical Sciences, Tel Aviv University. This research was funded by the European Union (ERC, EACTP, 101142020). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

ISSN 1433-8092

ROABPs are among the most studied algebraic models as they are the algebraic analog of read-once oblivious Boolean branching programs, a non-uniform model that captures space-bounded randomized algorithms. Thus, black-box PIT for ROABPs can be viewed as the algebraic analog of derandomizing RL. Indeed, both models also influenced one another. For example, the known-order hitting set of [FS13b] is the algebraic analog of Nisan's pseudorandom generator [Nis92]. Similarly, ideas from the algebraic setting have also influenced the Boolean setting: Gurjar and Volk [GV20] adapted the techniques of [AFS+18] to obtain pseudorandom generators for read-k oblivious Boolean branching programs that read their input in a known order. In addition to the connection to the Boolean setting, ROABPs are quite expressive. For example, they can compute sparse polynomials, sums of products of univariates and diagonal circuits [Sax08]. Moreover, black-box PIT for set-multilinear depth-three circuits and for non-commutative ABPs reduce to PIT for ROABPs [FS13b]. Finally, we note that Forbes and Shpilka [FS13a] showed that certain questions concerning derandomizing a special case of the Noether normalization lemma, which was studied by Mulmuley [Mul12], reduce to PIT for ROABPs.

Raz and Shpilka [RS05] gave a polynomial-time white-box PIT algorithm for ROABPs. Forbes and Shpilka [FS13b] gave quasipolynomial-size hitting sets, when the variable order is known. Later works extended this result to the unknown order case [FSS14, AGKS15], with further improvements in [GKS17, GG20]. We also mention that [GKST17] studied sums of ROABPs and that read-k oblivious ABPs were considered in [AFS+18]. Polynomial-size hitting sets are known only in special cases, such as constant width in a known order [GKS17] and width two [AGKS15].

In this paper we focus on *any-order ROABPs*. A polynomial has an any-order ROABP of width at most w if, for every permutation π, it can be computed by a width-w ROABP in that order. The work of Nisan [Nis91] on non-commutative formulas gave the following characterization for this model. A polynomial f has an any-order ROABP of width w if and only if for every partition of the variables into two sets, the matrix of coefficients of f with respect to this partition has rank at most w.

We note that even if a polynomial has a width-w any-order ROABP, it may have a different representation for each variable order π. Namely, the matrices M_i in (1) need not give a representation when we permute their order. A more restrictive model is that of *commutative ROABPs*, in which the matrices in the computation commute. In particular, permuting the matrices gives a ROABP that computes the polynomial in the permuted order. For more on the relation between any-order and commutative ROABPs see [RT22, BT24].

Despite its apparent restrictiveness, the class of any-order ROABPs captures sums of w products of univariate polynomials, and in particular w-sparse polynomials, and set-multilinear depth-3 circuits.

For any-order ROABPs we have stronger hitting sets than those known for general ROABPs. Forbes, Saptharishi and Shpilka [FSS14] gave hitting sets of size $\mathrm{poly}(n, w)^{O(\log\log w)}$ for multilinear any-order ROABPs.[1] Gurjar, Korwar and Saxena [GKS17] gave hitting sets of size $(ndw)^{O(\log\log w)}$, where d is an upper bound on the individual degrees. This was later improved by Guo and Gurjar [GG20] who constructed hitting sets of size $\mathrm{poly}(n, w^{\log\log w}, d)$, which is polynomial when $w \leqslant 2^{O(\log(nd)/\log\log(nd))}$.

Prior to this work, no hitting set of size polynomial in both n and w was known for any-order ROABPs, even for multilinear polynomials.

Another important model that we consider is that of set-multilinear depth-3 circuits. In this model the variable set is partitioned into disjoint sets $X = X_1 \sqcup \ldots \sqcup X_n$, and a circuit computes a polynomial of the form:

$$\sum_{i=1}^{s} \prod_{j=1}^{n} \ell_{i,j}(X_j),$$

where each $\ell_{i,j}$ is a homogeneous linear form in X_j. A polynomial computed by a set-multilinear circuit is also a tensor of rank at most s, and the smallest such s is its tensor rank. This makes depth-3 set-multilinear circuits the most natural model for computing tensors.

Depth-3 set-multilinear circuits also attracted much research. Many works focused on reconstructing such circuits from black-box access, see [BBB+00, KS06, BSV21, PSV24, BS25, SS26]. Raz and Shpilka [RS05]

[1] Forbes, Saptharishi and Shpilka [FSS14] used the term commutative ROABPs for the class we now call any-order ROABPs.

gave a polynomial-time white-box PIT algorithm for this model. In the black-box setting, Forbes and Shpilka [FS12] constructed hitting sets of essentially optimal size for $n = 2$, and of quasipolynomial size for tensors in general. Explicit quasipolynomial-size hitting sets also follow from [ASS13] and from the hitting sets for ROABPs. Indeed, replacing the variables of each X_j by distinct powers of a single new variable turns such a circuit into an any-order ROABP of width s, so when all the sets have size at most d the results of Guo and Gurjar [GG20] give hitting sets of size $\mathrm{poly}(n, d, s^{O(\log\log s)})$, which is polynomial when $s \leqslant 2^{O(\log(nd)/\log\log(nd))}$. We also note that polynomial-size hitting sets are known when the top fan-in s is constant [SS12, GKS17].

1.1 Our result

We construct the following hitting set.

Theorem 1.1. *There is an explicit construction of a set $\mathcal{H} \subseteq \mathbb{F}^n$ of size*

$$\mathrm{poly}(n, w)$$

that hits every nonzero multilinear polynomial in n variables that is computable by width-w ROABPs in every variable order.

We also give such a hitting set for depth-3 set-multilinear circuits whose variable sets have bounded size.

As mentioned above, any-order ROABPs can simulate this model, but since our result for any-order ROABPs only holds in the multilinear case and the simulation increases the individual degree we have to treat this case separately, although the proof is nearly identical.

Theorem 1.2. *Let $X = X_1 \sqcup \cdots \sqcup X_n$ be a partition of N variables into sets of size at most d each. There is an explicit construction of a set $\mathcal{H} \subseteq \mathbb{F}^N$ of size*

$$\mathrm{poly}(n, s) \cdot (d+1)^{O(\log(s+1))} = \mathrm{poly}(n, s^{O(\log(d+1))})$$

that hits every nonzero polynomial computed by a depth-3 set-multilinear circuit $\sum_{i=1}^{s} \prod_{j=1}^{n} \ell_{i,j}(X_j)$ with respect to this partition.

Both theorems hold over every field of size $|\mathbb{F}| \geqslant \mathrm{poly}(N, w)$, where N is the number of variables and $w = s$ in Theorem 1.2. If the size of $\mathbb{F}$ is too small then the evaluation points can be taken from an extension field of large enough polynomial size.

We remark again that prior to our work the best constructions for these models gave hitting sets of size $\mathrm{poly}(n, w^{\log\log w}, d)$, which is polynomial only for $w \leqslant 2^{O(\log(nd)/\log\log(nd))}$ [GG20].

Remark 1.3 (Extension to the border). We note that our proofs extend to the border as they only rely on two properties that are preserved under borders: the polynomial respects a variable partition to sets (that can be of size 1), that is, each of its monomials contains at most one variable from each set; and the polynomial has small partial derivative rank with respect to any partition of the variable sets. As these properties clearly hold in the border, our hitting sets also hit every polynomial in the closure of these classes.

1.2 Proof idea

We describe the proof of Theorem 1.1 and then explain the modification required for Theorem 1.2.

The construction is based on a combination of the technique of [FSS14] with a new observation that enables applying the hitting set of [BS26, BG26] for polynomials with a small space of partial (or Hasse) derivatives.

[FSS14] first constructs a small set of shifts so that the polynomial has a nonzero monomial of small support for one of the shifts. Specifically, [FSS14, Theorem 4.1 and Lemma 4.4] shows that shifting by an $O(\log w)$-wise independent monomial map gives low-support rank concentration for any-order ROABPs.

Since such maps can be obtained from the Klivans–Spielman generator [KS01], this gives an explicit set of $\mathrm{poly}(n, 2^{O(\log w)}) = \mathrm{poly}(n, w)$ shifts. Thus, for any multilinear any-order ROABP of width w, there is a shift in the set that after applying it the shifted polynomial has a nonzero monomial of support $r < \log w^2 + 1$. Since we will work with multilinear ROABPs, this gives a monomial of degree r (in the set-multilinear case we have to be slightly more careful as we explain later).

We now know that there is a monomial of degree $r = O(\log w)$, but we do not want to brute-force search over all sets of r variables. Instead, we hash the variables into r buckets (colors) using a perfect hash family [SS90] of polynomial size.

Assume that we have picked a hash function that maps the r variables of the monomial to distinct buckets (we may assume that we know r as it is an integer smaller than $O(\log w)$).

To every variable x_j that falls into bucket i we attach a new variable y_i. That is, we replace x_j by $y_i \cdot x_j$. The coefficient of $y_1 \cdots y_r$ then contains our nonzero monomial, and in particular it is nonzero. Denote this coefficient by $g(\mathbf{x})$.

The main new observation is that g has partial derivative dimension at most $2^r w \leqslant w^3$. Moreover, every evaluation of g is a linear combination of a small number of evaluations of f. Thus we have black-box access to g, and we can apply the hitting sets of [BS26, BG26] depending on the characteristic of the field.

To see why g has small partial derivative dimension, we observe that g is set-multilinear in the r sets $Y_1, \ldots, Y_r$, where Y_i contains all the variables that received color i.

Furthermore, since g is set-multilinear, differentiating g by a monomial m containing one variable from each Y_i, $i \in I$, gives the coefficient of m in g, when viewed as a polynomial in $\bigcup_{i \in I} Y_i$ over the remaining variables. It follows that for every set $I \subseteq [r]$, the dimension of the partial derivative space defined by multilinear derivatives whose variables come from $\cup_{i \in I} Y_i$ is at most the rank of (a submatrix of) the partial derivative matrix of (the shifted) f with respect to the partition to the variables that were mapped to the sets in I and the remaining ones. Thus, this dimension is at most w. Since there are at most 2^r such sets we get the $2^r w$ upper bound on the dimension of partial derivatives.

An important remark is that our proof only holds for multilinear polynomials. For higher individual degree, a nonzero monomial of small support need not contribute to the coefficient of the monomial $y_1 \cdots y_r$. Extracting a coefficient involving higher powers of the y_i preserves nonzeroness, but then our bound on the partial derivative dimension no longer holds. This is also the reason that Theorem 1.2 is not a corollary of Theorem 1.1: the natural transformation that replaces each set with powers of a corresponding new variable and turns the circuit into a sum of products of univariates, results in a polynomial with individual degrees higher than 1.

For the proof of Theorem 1.2 we first transform the set-multilinear circuit into an any-order ROABP by mapping each variable $x_{j,i} \in X_j$ to x_j^i. It is immediate that the resulting polynomial is computed by any-order ROABPs and has individual degree at most $\max_j |X_j|$. We then use the fact that for this model [FSS14] also gave an efficient construction of support-concentrating shifts. Once we find such a shift we translate it back to a shift of the original variables. We note that the low support property gives a low degree component in the original set-multilinear polynomial. Once we have that, a similar argument to the one above shows that its partial derivative dimension is small.

1.3 AI disclosure

I used Claude Opus 5.5 and ChatGPT astra to help prepare an initial draft, polish the writing and check the proofs. In addition, while preparing this paper, I asked ChatGPT to go over the list of submissions to STOC 2027 at https://sigact-stoc.github.io/stoc2027/ to find papers in algebraic complexity, and this is how I learned of [BG26].

2 Preliminaries

We give some basic definitions and known results that will be used in our proof.

For an exponent vector $\boldsymbol{\alpha} \in \mathbb{N}^n$, we define $\mathbf{x}^{\boldsymbol{\alpha}} = \prod_i x_i^{\alpha_i}$, and $|\boldsymbol{\alpha}| = \sum_i \alpha_i$. We also set $\mathrm{Supp}(\boldsymbol{\alpha}) = \{i : \alpha_i \neq 0\}$. A polynomial has individual degrees at most δ if every variable appears in degree at most δ. We write $[m]f$ for the coefficient of the monomial m in the polynomial f.

We next define the coefficient matrices of a polynomial with respect to a partition of its variables. These matrices are also often called the partial derivative matrices of the polynomial.

Definition 2.1 (Coefficient matrices)**.** For $f \in \mathbb{F}[\mathbf{x}]$ and $U \subseteq [n]$, let $C_U(f)$ have rows indexed by monomials in $\{x_i : i \in U\}$, columns indexed by monomials in the complementary set of variables $\{x_i : i \notin U\}$, and entry

$$C_U(f)_{m_1,m_2} = [m_1 m_2]f.$$

That is, for every two monomials m_1 and m_2 in the two variable sets, the entry in the matrix contains the coefficient of $m_1 m_2$ in f.

The characterization of ROABPs by Nisan [Nis91] says that f has width-w ROABPs in every order, if and only if

$$\operatorname{rank} C_U(f) \leqslant w \qquad \text{for every } U \subseteq [n]. \tag{2}$$

We note that shifting the variables preserves the any-order width bound.

Definition 2.2. Let $f = \sum_{\boldsymbol{\alpha}} v_{\boldsymbol{\alpha}} \mathbf{x}^{\boldsymbol{\alpha}}$, where the $v_{\boldsymbol{\alpha}}$ lie in a vector space over a field $\mathbb{F}$. The span of the $v_{\boldsymbol{\alpha}}$ is the *coefficient space* of f. We say that f is *degree-r concentrated* if the $v_{\boldsymbol{\alpha}}$ with $|\boldsymbol{\alpha}| \leqslant r$ span its coefficient space, and that f is *ℓ-support concentrated* if the $v_{\boldsymbol{\alpha}}$ with $|\mathrm{Supp}(\boldsymbol{\alpha})| < \ell$ span its coefficient space [ASS13, FSS14].

Theorem 2.3 ([FSS14, Theorem 4.1 and Lemma 4.4])**.** *Let $\ell = \lfloor \lg w^2 \rfloor + 1$ and $\delta \geqslant 1$. For[2] $m = O(\ell \log_{n\delta}(\delta+1))$ there is a $\mathrm{poly}(n, \delta, \ell)$-explicit polynomial map $\Gamma(\mathbf{t}, \mathbf{u})$ from $\mathbb{F}^m \times \mathbb{F}^m$ to $\mathbb{F}^n$, each of whose coordinates has individual degree $\mathrm{poly}(n, \delta)$, with the following property. For every polynomial $f \in \mathbb{F}[\mathbf{x}]$ of individual degree at most δ that has width-w ROABPs in every variable order, $f(\mathbf{x} + \Gamma(\mathbf{t}, \mathbf{u}))$ is ℓ-support concentrated over $\mathbb{F}(\mathbf{t}, \mathbf{u})$.*

Corollary 2.4. *Let ℓ be as in Theorem 2.3. There is an explicit set $S \subseteq \mathbb{F}^n$ of*

$$\mathrm{poly}(n, \delta) \cdot (\delta+1)^{O(\ell)} = \mathrm{poly}(n, \delta) \cdot (\delta+1)^{O(\log(w+1))}$$

shifts such that, for every nonzero polynomial f of individual degree at most δ that is computed by width-w ROABPs in every variable order, there exists $\mathbf{a} \in S$ for which $f(\mathbf{x} + \mathbf{a})$ has a nonzero coefficient of support less than ℓ. In particular, for polynomials with bounded individual degrees ($\delta = O(1)$) the set S has size $\mathrm{poly}(n, w)$.

We also require the well-known construction of perfect hash families of Schmidt and Siegel [SS90].

Theorem 2.5 (Perfect hash families [SS90])**.** *For $1 \leqslant e \leqslant n$ there is a family $\mathcal{P}_{n,e}$ of $2^{O(e)} \log^2 n$ maps $[n] \to [e]$, computable in time $\mathrm{poly}(n) \cdot 2^{O(e)}$, such that every e-element subset of $[n]$ is mapped injectively by some member of the family.*

The last tool that we need is hitting sets for polynomials with a small space of partial derivatives.

Definition 2.6 (Partial derivative dimension)**.** For $f \in \mathbb{F}[x_1, \ldots, x_N]$, define

$$\mathrm{PD}(f) = \mathrm{span}_{\mathbb{F}} \left\{ \frac{\partial^{|\boldsymbol{\alpha}|} f}{\partial x_1^{\alpha_1} \cdots \partial x_N^{\alpha_N}} : \boldsymbol{\alpha} \in \mathbb{N}^N \right\} \subseteq \mathbb{F}[\mathbf{x}].$$

The *partial derivative dimension* is $\dim_{\mathbb{F}} \mathrm{PD}(f)$.

Theorem 2.7 ([BS26, Theorem 3.1])**.** *Let $\mathbb{F}$ be a field of characteristic zero, $R \geqslant 1$, $B = NR^2 + R$. Set*

$$G_{N,R,i}(t) = \sum_{a=1}^{B} \frac{i^a}{a} t^a \qquad (1 \leqslant i \leqslant N),$$

to be the truncation of $-\log(1 - it)$. If $0 \neq f \in \mathbb{F}[x_1, \ldots, x_N]$ and $\dim \mathrm{PD}(f) \leqslant R$, then $f(G_{N,R}(t)) \not\equiv 0 \pmod{t^B}$. In particular, if also $\deg f \leqslant e$, then $f(G_{N,R}(\tau)) \neq 0$ for some τ in any set of $eB + 1$ elements of $\mathbb{F}$.

[2] We use $\delta + 1$ since we also apply the theorem for $\delta = 1$.

Over positive characteristic we work with the notion of Hasse derivative.

Definition 2.8 (Hasse derivative space)**.** For $f \in \mathbb{F}[x_1, \ldots, x_N]$ and $\boldsymbol{\gamma} \in \mathbb{N}^N$, the Hasse derivative $\partial_H^{(\boldsymbol{\gamma})} f$ is the coefficient of $\mathbf{y}^{\boldsymbol{\gamma}}$ in $f(\mathbf{x} + \mathbf{y}) \in \mathbb{F}[\mathbf{x}][\mathbf{y}]$. The *Hasse derivative space* $\partial_H^* f$ is the $\mathbb{F}$-linear span of all Hasse derivatives of f, including f itself.

A very recent result of Bal and Ghosh extends Theorem 2.7 to arbitrary characteristic using Hasse derivatives.

Theorem 2.9 ([BG26, Theorem 1.2])**.** *Let N, R, e be nonnegative integers and let $\mathbb{F}$ be a field of arbitrary characteristic with $|\mathbb{F}| \geqslant 2N^2R^4e^2$. There is an explicit polynomial map $G = (G_1, \ldots, G_N) : \mathbb{F} \times \mathbb{F} \to \mathbb{F}^N$, computable in time $\mathrm{poly}(N, R, e)$, with $\deg_t G_i = O(NR^2e)$ and $\deg_z G_i = O(NR^3e)$, such that for every $f \in \mathbb{F}[x_1, \ldots, x_N]$ of degree at most e with $\dim \partial_H^* f \leqslant R$,*

$$f \not\equiv 0 \iff f\big(G(z,t)\big) \not\equiv 0.$$

In particular, such polynomials have an explicit hitting set of size $O(N^2R^5e^4)$.

Remark 2.10. If g is multilinear then $\partial_H^{(\boldsymbol{\gamma})} g = 0$ unless $\boldsymbol{\gamma} \in \{0,1\}^N$, and for $\boldsymbol{\gamma} \in \{0,1\}^N$ the Hasse derivative $\partial_H^{(\boldsymbol{\gamma})} g$ equals the partial derivative $\partial^{\boldsymbol{\gamma}} g$. Hence $\dim \partial_H^* g = \dim \mathrm{PD}(g)$ for multilinear g, over any field.

We shall only apply these results to multilinear polynomials so we can speak of partial derivatives regardless of the characteristic of the field.

3 The hitting sets

The proofs of Theorems 1.1 and 1.2 are very similar and so we state a general setting and then prove each.

Let $X = X_1 \sqcup \cdots \sqcup X_n$ be a partition of the N variables into disjoint sets of variables. We refer to these sets as *blocks*. For a set I of blocks we write $X_I = \bigcup_{j \in I} X_j$, and we write $C_I(f)$ for the partial derivative matrix of f with rows indexed by monomials in X_I and columns indexed by monomials in the remaining variables. Our argument holds for all $f \in \mathbb{F}[X]$ that satisfy

1. Multilinear set-partitioned: f is multilinear and every monomial in it contains at most one variable from each block.
2. Coefficient matrices of small rank: $\mathrm{rank}\, C_I(f) \leqslant w$ for every set I of blocks.

It is immediate to verify that both properties continue to hold if we shift the variables. We denote the class of polynomials satisfying these properties by $\mathcal{C}$.

As described in Section 1.2 we first find a shift of the variables so that $f(X + \mathbf{a})$ has a nonzero monomial of small support. In our case, we will work with multilinear polynomials so small support implies small degree. The next step is hashing the blocks into different buckets (colors) in a way that maps the blocks of that low degree monomial injectively. We then multiply color class j (i.e., all the variables that were mapped to j) with a new variable y_j.

So, fix a shift $\mathbf{a} \in \mathbb{F}^N$, an integer $r \geqslant 1$ and a map $h : [n] \to [r]$, which colors the blocks. Let $y_1, \ldots, y_r$ be new variables. We define

$$g_{\mathbf{a},h}(X) = [y_1 y_2 \cdots y_r]\, f\big(\mathbf{a} + (y_{h(j)} X_j)_{j=1}^n\big), \tag{3}$$

that is, every variable $x \in X_j$ is replaced by $a_x + y_{h(j)} x$, and we take the coefficient of $y_1 \cdots y_r$.

We now need to access the coefficient of $y_1 \cdots y_r$. We obtain this using standard interpolation. Let z be a new variable. Replace each $y_j x_i$ with $z x_i y_j$ and then interpolate to get the degree-r term in z. This term will contain a multilinear monomial in y of degree r and we will then interpolate (using inclusion-exclusion as in Ryser's formula) to get access to its coefficient. We combine both steps into one interpolation lemma (Lemma 3.2).

We also need the following notation. Denote the shifted polynomial by $F(X)$ and write $F(X) = f(X+\mathbf{a}) = \sum_m c_m m$, where $c_m \in \mathbb{F}$. For each monomial m we define $J(m)$ to be the set of blocks that the variables of m belong to. Since

$$f(\mathbf{a} + (y_{h(j)} X_j)_{j=1}^n) = F((y_{h(j)} X_j)_{j=1}^n)$$

and F is set-partitioned, every monomial m of F is multiplied by $\prod_{j \in J(m)} y_{h(j)}$. Hence

$$g_{\mathbf{a},h} = \sum_{m:\ h \text{ maps } J(m) \text{ bijectively onto } [r]} c_m\, m. \tag{4}$$

In particular, $g_{\mathbf{a},h}$ is homogeneous of degree r, and every monomial of $g_{\mathbf{a},h}$ contains exactly one variable from $X_{h^{-1}(i)}$ for every color $i \in [r]$.

For $I \subseteq [r]$ and $\mathbf{b} \in \mathbb{F}^N$, we next define the projection of $\mathbf{b}$ to the blocks colored I, that we denote by $\mathbf{b}^{h,I} \in \mathbb{F}^N$, as follows:

$$b_x^{h,I} = \begin{cases} b_x, & x \in X_j \text{ with } h(j) \in I, \\ 0, & \text{otherwise.} \end{cases} \tag{5}$$

In other words, $\mathbf{b}^{h,I}$ equals $\mathbf{b}$ on the variables that were mapped by h to a color in the set I and zero on any other variable. It is immediate from the definition that for any monomial m, $m(\mathbf{b}^{h,I})$ equals $m(\mathbf{b})$ if all the blocks in $J(m)$ have colors in I. If this is not the case then $m(\mathbf{b}^{h,I}) = 0$.

Lemma 3.1. *If $F = f(X + \mathbf{a})$ has a nonzero monomial m of degree r and h is injective on $J(m)$, then $g_{\mathbf{a},h} \neq 0$.*

Proof. Since f (and also F) is set-partitioned, $|J(m)| = \deg m = r$, so h maps $J(m)$ bijectively onto $[r]$. By (4), m appears in $g_{\mathbf{a},h}$ with its coefficient $c_m \neq 0$. □

Lemma 3.2. *For every $\mathbf{b} \in \mathbb{F}^N$,*

$$g_{\mathbf{a},h}(\mathbf{b}) = [z^r] \sum_{I \subseteq [r]} (-1)^{r-|I|} f\big(\mathbf{a} + z\mathbf{b}^{h,I}\big). \tag{6}$$

Thus, for every set $\Lambda \subseteq \mathbb{F}$ of $n+1$ elements, if $g_{\mathbf{a},h}(\mathbf{b}) \neq 0$ then $f(\mathbf{a} + \lambda \mathbf{b}^{h,I}) \neq 0$ for some $\lambda \in \Lambda$ and $I \subseteq [r]$.

Proof. By definition, $f(\mathbf{a} + z\mathbf{b}^{h,I}) = F(z\mathbf{b}^{h,I}) = \sum_m c_m z^{\deg m} m(\mathbf{b}^{h,I})$. Thus, the coefficient of z^r contains all monomials of degree r.

Since for every monomial m we have $m(\mathbf{b}^{h,I}) = m(\mathbf{b})$ if all the blocks in $J(m)$ have colors in I (i.e., $h(J(m)) \subseteq I$) and $m(\mathbf{b}^{h,I}) = 0$ otherwise, the right-hand side of (6) can be written as

$$\begin{aligned} [z^r] \sum_{I \subseteq [r]} (-1)^{r-|I|} f\big(\mathbf{a} + z\mathbf{b}^{h,I}\big) &= [z^r] \sum_{I \subseteq [r]} (-1)^{r-|I|} \sum_m c_m z^{\deg m} m(\mathbf{b}^{h,I}) \\ &= \sum_{m:\ \deg m = r} c_m \sum_{I \subseteq [r]} (-1)^{r-|I|} m(\mathbf{b}^{h,I}) \\ &= \sum_{m:\ \deg m = r} c_m\, m(\mathbf{b}) \sum_{I:\ h(J(m)) \subseteq I \subseteq [r]} (-1)^{r-|I|}. \end{aligned}$$

Thus, the coefficient of $c_m m(\mathbf{b})$ in this sum equals

$$\sum_{I:\ h(J(m)) \subseteq I \subseteq [r]} (-1)^{r-|I|} = \begin{cases} 1, & h(J(m)) = [r], \\ 0, & \text{otherwise.} \end{cases}$$

Consequently, only monomials of degree r satisfying $h(J(m)) = [r]$ survive. Therefore (6) follows from the definition of $g_{\mathbf{a},h}$ (recall (4)).

To conclude the proof we note that we can access the coefficient of z^r by simple interpolation, since F is set-partitioned and hence $\sum_I (-1)^{r-|I|} f(\mathbf{a} + z\mathbf{b}^{h,I})$ has degree at most n in z. By assumption, this coefficient is nonzero and therefore there must be a point $\lambda \in \Lambda$ and a set $I \subseteq [r]$ on which $f(\mathbf{a} + \lambda \mathbf{b}^{h,I})$ is nonzero. □

We next prove that the partial derivative dimension of $g_{a,h}$ is at most $\mathrm{poly}(w)$.

Lemma 3.3. *Let $f \in \mathcal{C}$. Then, in the notation above,* $\dim \mathrm{PD}(g_{a,h}) \leqslant 2^r w$.

Proof. To simplify notation denote $g = g_{a,h}$ and, for a color i, let $Y_i = X_{h^{-1}(i)}$ be the set of variables whose block has color i. By (4), every monomial of g contains exactly one variable from each Y_i. Hence a partial derivative $\partial^\alpha g$ is zero unless x^α is a product of distinct variables from distinct sets Y_i. For such α, let $I = h(J(x^\alpha)) \subseteq [r]$ be the set of colors that variables of x^α received. In particular, x^α is a monomial in $Y_I = \bigcup_{i \in I} Y_i$. Since g is multilinear,

$$\partial^\alpha g = \sum_{m_2} \left([x^\alpha m_2] g\right) m_2,$$

where m_2 ranges over monomials in the variables outside Y_I. Thus $\partial^\alpha g$ is the polynomial whose coefficient vector is the row of $C_{h^{-1}(I)}(g)$ indexed by x^α. It follows that the space of partial derivatives according to monomials whose set of colors is I is contained in the row-span of $C_{h^{-1}(I)}(g)$. Hence, it has dimension at most $\operatorname{rank} C_{h^{-1}(I)}(g)$. Since $C_{h^{-1}(I)}(g)$ is a submatrix of the coefficient matrix of $f(X + a)$ (after deleting its zero rows and columns), with respect to the same partition, we get from our assumption that $\operatorname{rank} C_{h^{-1}(I)}(g) \leqslant w$. Going over all 2^r possible subsets I gives the claimed bound. □

We are ready to give the hitting set.

Construction. Let $\ell = \lfloor \lg w^2 \rfloor + 1$. Let $\mathcal{A} \subseteq \mathbb{F}^N$ be a set of shifts with the following property: for every nonzero $f \in \mathcal{C}$, there is $a \in \mathcal{A}$ such that $f(X + a)$ has a nonzero monomial of support less than ℓ. Finally, for $1 \leqslant r < \ell$ let $\mathcal{P}_r = \mathcal{P}_{n,r}$ be the perfect hash family of Theorem 2.5.

Denote by $R_r = 2^r w$ the upper bound on the partial derivative dimension of $g_{a,h}$ (Lemma 3.3). Additionally, denote $B_r = NR_r^2 + R_r$ and set

$$G_r = G_{N,R_r}.$$

Fix sets $\Lambda, T_r \subseteq \mathbb{F}$ with $|\Lambda| = n + 1$ and $|T_r| = rB_r + 1$. Define

$$\mathcal{H} = \mathcal{A} \cup \bigcup_{r=1}^{\ell-1} \left\{ a + \lambda \big(G_r(\tau)\big)^{h,I} : a \in \mathcal{A},\ h \in \mathcal{P}_r,\ I \subseteq [r],\ \tau \in T_r,\ \lambda \in \Lambda \right\}, \tag{7}$$

where we have used the notation of (5).

Over fields of positive characteristic, we replace the set $\{G_r(\tau) : \tau \in T_r\}$ in (7) by the hitting set of Theorem 2.9 with parameters N, $R = R_r$ and $e = r$. For this we assume that $|\mathbb{F}| \geqslant 2N^2R_r^4r^2 = \mathrm{poly}(N, w)$ (otherwise we take the points from an extension field), and the new set has size $O(N^2R_r^5r^4) = \mathrm{poly}(N, w)$ as well.

Lemma 3.4. *Let $\mathcal{A}$ have the property above for $\mathcal{C}$. Then $\mathcal{H}$ hits every nonzero $f \in \mathcal{C}$. Moreover,*

$$|\mathcal{H}| \leqslant |\mathcal{A}| \cdot \mathrm{poly}(n, N, w).$$

Proof. Choose $a \in \mathcal{A}$ such that $F = f(X + a)$ has a nonzero monomial m of support $r < \ell$. By multilinearity (which follows from being set-partitioned) we also have $\deg(m) = r$. If $r = 0$, then $f(a) = F(0) \neq 0$ and $a \in \mathcal{H}$. Otherwise, since the polynomial is set-partitioned, m contains variables from exactly r blocks, and by Theorem 2.5 there is $h \in \mathcal{P}_r$ that is injective on $J(m)$. It follows that $g_{a,h}$ is a nonzero homogeneous polynomial of degree r (Lemma 3.1), and from Lemma 3.3 we get that $\dim \mathrm{PD}(g_{a,h}) \leqslant R_r$.

Over fields of characteristic zero, Theorem 2.7 implies that there is $\tau \in T_r$ with $g_{a,h}(G_r(\tau)) \neq 0$. By Lemma 3.2 with $b = G_r(\tau)$, there are $\lambda \in \Lambda$ and $I \subseteq [r]$ with $f(a + \lambda(G_r(\tau))^{h,I}) \neq 0$, and $a + \lambda(G_r(\tau))^{h,I}$ is an evaluation point from $\mathcal{H}$.

Over fields of positive characteristic the argument is the same: since $g_{a,h}$ is multilinear, $\dim \partial^*_H g_{a,h} = \dim \mathrm{PD}(g_{a,h}) \leqslant R_r$ (recall Remark 2.10), so by Theorem 2.9 some point b of the corresponding set satisfies $g_{a,h}(b) \neq 0$, and we apply Lemma 3.2 to this b.

For upper bounding the size, we note that since $r \leqslant \ell - 1$, $2^r \leqslant w^2$, so $R_r \leqslant w^3$. Hence, $B_r \leqslant Nw^6 + w^3$ and $|T_r| \leqslant \ell B_r + 1$. By Theorem 2.5 we have $|\mathcal{P}_r| \leqslant 2^{O(r)} \log^2 n = \mathrm{poly}(w) \log^2 n$. We therefore have

$$|\mathcal{H}| \leqslant |\mathcal{A}|\Big(1 + (n+1)\sum_{r=1}^{\ell-1} |\mathcal{P}_r| 2^r |T_r|\Big) = |\mathcal{A}| \cdot \mathrm{poly}(n, N, w). \qquad \square$$

3.1 Proofs of the main theorems

Proof of Theorem 1.1. In this case all our blocks are singletons and so $N = n$. By (2) a polynomial with width-w ROABPs in every order belongs to the class $\mathcal{C}$.

Let $\mathcal{A} = \mathcal{S}$ be the set of Corollary 2.4 with $\delta = 1$. Then, $|\mathcal{A}| = \mathrm{poly}(n, w)$. By Lemma 3.4, $\mathcal{H}$ is a hitting set of size $\mathrm{poly}(n, w)$. $\square$

The proof of Theorem 1.2 first translates the polynomial into one that can be computed by any-order ROABPs and then shows that a good shift for the translated polynomial yields one for the original polynomial.

Proof of Theorem 1.2. Let $X = X_1 \sqcup \ldots \sqcup X_n$ and $f = \sum_{i=1}^{s} \prod_{j=1}^{n} \ell_{i,j}(X_j)$. For every set I of blocks,

$$\prod_j \ell_{i,j}(X_j) = \Big(\prod_{j \in I} \ell_{i,j}(X_j)\Big) \cdot \Big(\prod_{j \notin I} \ell_{i,j}(X_j)\Big),$$

so each of the s products contributes a matrix of rank at most one to $C_I(f)$, so the rank of every partial derivative matrix is at most s. Thus, set-multilinear polynomials are members of the class $\mathcal{C}$.

Denote $X_j = \{x_{j,1}, \ldots, x_{j,d_j}\}$ with $d_j \leqslant d$. To apply Corollary 2.4 we transform the polynomial so that it is computed by any-order ROABPs by replacing each $x_{j,i}$ with x_j^i, the ith power of a new variable x_j.

The resulting polynomial, which we denote by $q(\mathbf{x})$, has individual degree at most d and width-s ROABPs in every order, as it is a sum of s products of univariate polynomials. Let $\mathcal{S} \subseteq \mathbb{F}^n$ be the set of Corollary 2.4 with $\delta = d$ and $w = s$. This gives a shift $\mathbf{b} \in \mathcal{S}$ such that $q(\mathbf{x} + \mathbf{b})$ has a low support monomial (but not necessarily of low degree). We next translate this shift into a shift of the original polynomial f.

Define the corresponding set of substitutions for the original variables by

$$\mathcal{A} = \{(b_j^t)_{j \in [n], t \in [d_j]} : \mathbf{b} \in \mathcal{S}\} \subseteq \mathbb{F}^N.$$

Define

$$\rho_{j,t}(x_j) = (x_j + b_j)^t - b_j^t.$$

Then

$$q(\mathbf{x} + \mathbf{b}) = f\big(\mathbf{a} + \rho(\mathbf{x})\big) = F(\rho(\mathbf{x})),$$

where $\mathbf{a} = (b_j^t)_{j,t} \in \mathcal{A}$ is the point corresponding to $\mathbf{b}$ and $F = f(X + \mathbf{a})$.

As a polynomial in x_j, $\rho_{j,t}(x_j)$ is divisible by x_j. Therefore, after substituting ρ into a monomial m of F, we obtain a polynomial all of whose monomials have support exactly $J(m)$. Hence a coefficient of $q(\mathbf{x} + \mathbf{b})$ of support J is a linear combination of coefficients of monomials of F of degree $|J|$.

By Corollary 2.4, for some $\mathbf{b} \in \mathcal{S}$, $q(\mathbf{x} + \mathbf{b})$ has a nonzero coefficient of support less than $\ell = \lfloor \lg s^2 \rfloor + 1$. Consequently, one of these coefficients of F is nonzero. In particular, $f(X + \mathbf{a})$ has a nonzero monomial of degree less than ℓ for the corresponding $\mathbf{a} \in \mathcal{A}$. Thus $\mathcal{A}$ has the required property, and $|\mathcal{A}| \leqslant |\mathcal{S}| = \mathrm{poly}(n, d) \cdot (d+1)^{O(\log(s+1))}$. By Lemma 3.4, with $w = s$ and $N \leqslant nd$, the set $\mathcal{H}$ hits f and has size

$$\mathrm{poly}(n, d) \cdot (d+1)^{O(\log(s+1))} \cdot \mathrm{poly}(n, nd, s) = \mathrm{poly}(n, s) \cdot (d+1)^{O(\log(s+1))}. \qquad \square$$

We note that in the proof above q is used only to find the shift and Lemma 3.4 is applied to f itself, which is in $\mathcal{C}$.

Remark 3.5. In the proof of Theorem 1.2 we justified q having an any-order ROABP by saying that it is a sum of products of univariates. This is not necessarily preserved in the border so we give a more detailed argument here that does hold for the border. Observe that the substitution replacing $x_{j,i}$ with x_j^i maps the monomials of f to distinct monomials. Hence, the coefficient matrices of q are the same as those of f up to permuting rows and columns. It follows that $\operatorname{rank} C_I(q) \leqslant s$ for every partition of the variables, and by Nisan's characterization [Nis91] q has an any-order ROABP of width at most s. This rank based argument clearly holds in the border.

4 Open problems

We conclude with some open problems.

1. *Higher individual degree.* The main question left open by this work is to construct hitting sets of size $\operatorname{poly}(n, w, d)$ for any-order ROABPs of individual degree d. From [FSS14] we still get a set of size $\operatorname{poly}(n) \cdot d^{O(\log w)}$ that contains a shift that yields a nonzero coefficient of support $O(\log w)$. However, a monomial of small support need not survive when we extract the coefficient of $y_1 \cdots y_r$. We can instead extract a coefficient involving higher powers of the y_i, but our bound on the partial derivative dimension no longer applies. In particular, the case of individual degree two is open. Solving this for arbitrary individual degree will also yield hitting sets for arbitrary set-multilinear depth-3 circuits.

2. *Set-multilinear depth-3 circuits with large sets.* Can the factor $(d+1)^{O(\log s)}$ in Theorem 1.2 be removed? Equivalently, construct explicit hitting sets of size $\operatorname{poly}(n, d, s)$ for tensors of rank at most s with n factors of dimension d.

3. *Orbits.* Hitting sets for affine orbits of any-order ROABPs were constructed in [ST24, BG22]. Can the approach of this paper give polynomial-size hitting sets for orbits of multilinear any-order ROABPs, or set-multilinear depth-3 circuits? We note that an affine transformation does not preserve multilinearity (nor set-multilinearity), but perhaps existing ideas could already yield such a hitting set.

Bibliography

[AFS+18] Matthew Anderson, Michael A. Forbes, Ramprasad Saptharishi, Amir Shpilka, and Ben Lee Volk. Identity testing and lower bounds for read-k oblivious algebraic branching programs. *ACM Transactions on Computation Theory*, 10(1):3:1–3:30, 2018. Preliminary version in CCC 2016. 2

[AGKS15] Manindra Agrawal, Rohit Gurjar, Arpita Korwar, and Nitin Saxena. Hitting-sets for ROABP and sum of set-multilinear circuits. *SIAM Journal on Computing*, 44(3):669–697, 2015. 2

[Agr05] Manindra Agrawal. Proving lower bounds via pseudo-random generators. In *Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2005)*, volume 3821 of *Lecture Notes in Computer Science*, pages 92–105. Springer, 2005. 1

[ASS13] Manindra Agrawal, Chandan Saha, and Nitin Saxena. Quasi-polynomial hitting-set for set-depth-Δ formulas. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC 2013)*, pages 321–330, 2013. 3, 5

[AV08] Manindra Agrawal and V. Vinay. Arithmetic circuits: A chasm at depth four. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2008)*, pages 67–75, 2008. 1

[BBB+00] Amos Beimel, Francesco Bergadano, Nader H. Bshouty, Eyal Kushilevitz, and Stefano Varricchio. Learning functions represented as multiplicity automata. *Journal of the ACM*, 47(3):506–530, 2000. 2

[BG22] Vishwas Bhargava and Sumanta Ghosh. Improved hitting set for orbit of ROABPs. *Computational Complexity*, 31(2):Article 15, 2022. Preliminary version in APPROX/RANDOM 2021, LIPIcs 207, doi:10.4230/LIPIcs.APPROX/RANDOM.2021.30. 10

[BG26] Aranya Kumar Bal and Sumanta Ghosh. Hitting sets for polynomials with small Hasse derivative space over arbitrary characteristic fields. Manuscript, 2026. STOC 2027 submission no. 1410, https://github.com/sigact-stoc/stoc2027/releases/download/papers-1/1410.pdf, 2026. 3, 4, 6

[BS25] Vishwas Bhargava and Devansh Shringi. Faster & deterministic FPT algorithm for worst-case tensor decomposition. In *52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025)*, volume 334 of *LIPIcs*, pages 28:1–28:20, 2025. 2

[BS26] Shubham Bhardwaj and Ramprasad Saptharishi. Hitting sets for polynomials with small partial derivative spaces. ECCC Report TR26-217, Electronic Colloquium on Computational Complexity, 2026. 3, 4, 5

[BSV21] Vishwas Bhargava, Shubhangi Saraf, and Ilya Volkovich. Reconstruction algorithms for low-rank tensors and depth-3 multilinear circuits. In *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing (STOC 2021)*, pages 809–822, 2021. 2

[BT24] Vishwas Bhargava and Anamay Tengse. Explicit commutative ROABPs from partial derivatives. In *44th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2024)*, volume 323 of *LIPIcs*, pages 10:1–10:15, 2024. 2

[DG24] Pranjal Dutta and Sumanta Ghosh. Advances in polynomial identity testing. *SIGACT News Complexity Theory Column 121*, 55(2):53–88, 2024. 1

[DL78] Richard A. DeMillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7(4):193–195, 1978. 1

[FS12] Michael A. Forbes and Amir Shpilka. On identity testing of tensors, low-rank recovery and compressed sensing. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC 2012)*, pages 163–172, 2012. 3

[FS13a] Michael A. Forbes and Amir Shpilka. Explicit Noether normalization for simultaneous conjugation via polynomial identity testing. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2013)*, volume 8096 of *Lecture Notes in Computer Science*, pages 527–542. Springer, 2013. 2

[FS13b] Michael A. Forbes and Amir Shpilka. Quasipolynomial-time identity testing of non-commutative and read-once oblivious algebraic branching programs. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2013)*, pages 243–252, 2013. 2

[FSS14] Michael A. Forbes, Ramprasad Saptharishi, and Amir Shpilka. Hitting sets for multilinear read-once algebraic branching programs, in any order. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC 2014)*, pages 867–875, 2014. Full version: *Pseudorandomness for multilinear read-once algebraic branching programs, in any order*, arXiv:1309.5668. 2, 3, 4, 5, 10

[GG20] Zeyu Guo and Rohit Gurjar. Improved explicit hitting-sets for ROABPs. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2020)*, volume 176 of *LIPIcs*, pages 4:1–4:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. 2, 3

[GKKS16] Ankit Gupta, Pritish Kamath, Neeraj Kayal, and Ramprasad Saptharishi. Arithmetic circuits: A chasm at depth 3. *SIAM Journal on Computing*, 45(3):1064–1079, 2016. 1

[GKS17] Rohit Gurjar, Arpita Korwar, and Nitin Saxena. Identity testing for constant-width, and any-order, read-once oblivious arithmetic branching programs. *Theory of Computing*, 13(2):1–21, 2017. 2, 3

[GKST17] Rohit Gurjar, Arpita Korwar, Nitin Saxena, and Thomas Thierauf. Deterministic identity testing for sum of read-once oblivious arithmetic branching programs. *Computational Complexity*, 26(4):835–880, 2017. 2

[GV20] Rohit Gurjar and Ben Lee Volk. Pseudorandom bits for oblivious branching programs. *ACM Transactions on Computation Theory*, 12(2):8:1–8:12, 2020. Preliminary version in ECCC TR17-122. arXiv:1708.02054. 2

[HS80] J. Heintz and C. P. Schnorr. Testing polynomials which are easy to compute (extended abstract). In *Proceedings of the Twelfth Annual ACM Symposium on Theory of Computing*, STOC '80, pages 262–272, New York, NY, USA, 1980. Association for Computing Machinery. 1

[KI04] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1–2):1–46, 2004. 1

[Koi12] Pascal Koiran. Arithmetic circuits: The chasm at depth four gets wider. *Theoretical Computer Science*, 448:56–65, 2012. 1

[KS01] Adam R. Klivans and Daniel A. Spielman. Randomness efficient identity testing of multivariate polynomials. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC 2001)*, pages 216–223, 2001. 4

[KS06] Adam R. Klivans and Amir Shpilka. Learning restricted models of arithmetic circuits. *Theory of Computing*, 2(10):185–206, 2006. 2

[KST26] Nimrod Kaplan, Amir Shpilka, and Yann Tal. PIT and reconstruction for sums of powers. arXiv:2611.07106, 2026. 1

[Mul12] Ketan Mulmuley. Geometric complexity theory V: Equivalence between blackbox derandomization of polynomial identity testing and derandomization of Noether's normalization lemma. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2012)*, pages 629–638, 2012. 2

[Nis91] Noam Nisan. Lower bounds for non-commutative computation. In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing (STOC 1991)*, pages 410–418, 1991. 2, 5, 10

[Nis92] Noam Nisan. Pseudorandom generators for space-bounded computation. *Combinatorica*, 12(4):449–461, 1992. 2

[Ore22] Øystein Ore. Über höhere kongruenzen. *Norsk Mat. Forenings Skrifter*, 1(7):15, 1922. 1

[PSV24] Shir Peleg, Amir Shpilka, and Ben Lee Volk. Tensor reconstruction beyond constant rank. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, volume 287 of *LIPIcs*, pages 87:1–87:20, 2024. 2

[RS05] Ran Raz and Amir Shpilka. Deterministic polynomial identity testing in non-commutative models. *Computational Complexity*, 14(1):1–19, 2005. 2

[RT22] C. Ramya and Anamay Tengse. On finer separations between subclasses of read-once oblivious ABPs. In *39th International Symposium on Theoretical Aspects of Computer Science (STACS 2022)*, volume 219 of *LIPIcs*, pages 53:1–53:23, 2022. 2

[Sax08] Nitin Saxena. Diagonal circuit identity testing and lower bounds. In *Automata, Languages and Programming (ICALP 2008)*, volume 5125 of *Lecture Notes in Computer Science*, pages 60–71. Springer, 2008. 2

[Sax09] Nitin Saxena. Progress on polynomial identity testing. *Bulletin of the EATCS*, 99:49–79, 2009. 1

[Sax14] Nitin Saxena. Progress on polynomial identity testing – II. In *Perspectives in Computational Complexity: The Somenath Biswas Anniversary Volume*, volume 26 of *Progress in Computer Science and Applied Logic*, pages 131–146. Birkhäuser, 2014. arXiv:1401.0976. 1

[Sch80] Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4):701–717, 1980. 1

[SS90] Jeanette P. Schmidt and Alan Siegel. The spatial complexity of oblivious k-probe hash functions. *SIAM Journal on Computing*, 19(5):775–786, 1990. 4, 5

[SS12] Nitin Saxena and C. Seshadhri. Blackbox identity testing for bounded top-fanin depth-3 circuits: the field doesn't matter. *SIAM Journal on Computing*, 41(5):1285–1298, 2012. 3

[SS26] Shubhangi Saraf and Devansh Shringi. Fast algorithms for low rank tensor decomposition. ECCC Report TR26-295, Electronic Colloquium on Computational Complexity, 2026. 2

[ST24] Chandan Saha and Bhargav Thankey. Hitting sets for orbits of circuit classes and polynomial families. *ACM Transactions on Computation Theory*, 2024. Preliminary version in APPROX/RANDOM 2021, LIPIcs 207, pages 50:1–50:26. 10

[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. 1

[Tav15] Sébastien Tavenas. Improved bounds for reduction to depth 4 and depth 3. *Information and Computation*, 240:2–11, 2015. 1

[Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *Symbolic and Algebraic Computation (EUROSAM 1979)*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226. Springer, 1979. 1

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il