

Identity Testing for constant-width, and commutative, read-once oblivious ABPs

Rohit Gurjar^{*2}, Arpita Korwar^{†1}, and Nitin Saxena^{‡1}

¹Department of Computer Science and Engineering, IIT Kanpur, India

²Aalen University, Germany

January 28, 2016

Abstract

We give improved hitting-sets for two special cases of Read-once Oblivious Arithmetic Branching Programs (ROABP). First is the case of an ROABP with known variable order. The best hitting-set known for this case had cost $(nw)^{O(\log n)}$ where n is the number of variables and w is the width of the ROABP. Even for a constant-width ROABP, nothing better than a quasi-polynomial bound was known. We improve the hitting-set complexity for the known-order case to $n^{O(\log w)}$. In particular, this gives the first polynomial time hitting-set for constant-width ROABP (known-order). However, our hitting-set works only over those fields whose characteristic is zero or large enough. To construct the hitting-set, we use the concept of the rank of partial derivative matrix. Unlike previous approaches whose basic building block is a monomial map, we use a polynomial map.

The second case we consider is that of commutative ROABP. The best known hitting-set for this case had cost $d^{O(\log w)}(nw)^{O(\log \log w)}$, where d is the individual degree. We improve this hitting-set complexity to $(ndw)^{O(\log \log w)}$. We get this by achieving rank concentration more efficiently.

1 Introduction

The polynomial identity testing (PIT) problem asks if a given multivariate polynomial is identically zero. The input to the problem is given via an arithmetic model computing a polynomial, for example, an arithmetic circuit or an arithmetic branching program. These are arithmetic analogues of boolean circuits and boolean branching programs, respectively. The degree of the given polynomial is assumed to be polynomially bounded. Usually, any such circuit or branching program can compute a polynomial with exponentially many monomials (exponential in the circuit size). Thus, one cannot compute the polynomial explicitly. However, given such an input, it is possible to efficiently evaluate the polynomial at a point in the field. This property enables a randomized polynomial identity test with one-sided error. It is known that evaluating a small-degree nonzero polynomial over a random point gives a nonzero value with a good probability [DL78, Sch80, Zip79]. Thus, the randomized test is to just evaluate the input polynomial, given as an arithmetic circuit or an arithmetic branching program at a random point.

Finding an efficient deterministic algorithm for PIT has been a major open question in complexity theory. The question is also related to arithmetic circuit lower bounds [Agr05,

^{*}rgurjar@cse.iitk.ac.in, supported by DFG grant TH 472/4 and TCS research fellowship

[†]arpk@cse.iitk.ac.in

[‡]nitin@cse.iitk.ac.in, supported by DST-SERB

HS80, KI03]. The PIT problem has been studied in two paradigms: (i) blackbox test, where one can only evaluate the polynomial at a chosen point, (ii) whitebox test, where one has access to the input circuit or arithmetic branching program. A blackbox test is essentially the same as finding a hitting-set – a set of points such that any nonzero polynomial evaluates to a nonzero value on at least one of the points in the set. This work concerns finding hitting-sets for a special model, called read-once oblivious arithmetic branching programs (ROABP).

An *arithmetic branching program (ABP)* is a directed layered graph, with edges going from a layer of vertices to the next layer. The first and the last layers have one vertex each, called the source and the sink. Each edge of the graph has a label, which is a simple polynomial, for example a univariate polynomial. For any path p , its weight is defined to be the product of labels on all the edges in p . The ABP is said to compute a polynomial which is the sum of weights of all the paths from the source to the sink. ABPs are a strong model for computing polynomials. It is known that for any arithmetic circuit with polynomially bounded degree, one can find an ABP of quasi-polynomial size computing the same polynomial (see for example [Koi12]). Apart from its size, another important parameter for an ABP is its width. The width of an ABP is the maximum number of vertices in any layer of the associated graph. Even when the width is restricted to a constant, the ABP model is quite powerful. Ben-Or and Cleve [BOC92] have shown that width-3 ABPs have the same expressive power as arithmetic formulas.

An ABP is called a *read-once oblivious ABP or ROABP* if every variable occurs in at most one layer of edges in the ABP. For an ROABP, one can assume without loss of generality that any variable occurs in exactly one layer of edges. The order of the variables in consecutive layers is said to be the *variable order* of the ROABP. The read-once property severely restricts the power of the ABP. There are polynomials known which can be computed by a simple depth-3 ($\Sigma\Pi\Sigma$) circuit but require an exponential size ROABP [KNS15]. Also note that there are polynomials which have a small ROABP in one variable order but require exponential size in another variable order. Nisan [Nis91] gave the exact characterization of the polynomials computed by width- w ROABPs in a certain variable order. In particular, they gave exponential lower bounds for this model. Their work is actually on non-commutative ABPs but the same results also apply to ROABP.

The question of whitebox identity testing of ROABPs has been settled by Raz and Shpilka [RS05], who gave a polynomial time algorithm for this. However, though ROABPs are a relatively well-understood model, we still do not have a polynomial time blackbox algorithm. The blackbox question is studied with two variations: one where we know the variable order of the ROABP and the other where we do not know it. For known-order ROABPs, Forbes and Shpilka [FS13] gave the first efficient blackbox test with $(ndw)^{O(\log n)}$ time complexity, where n is the number of variables, w is the width of the ROABP, and, d is the individual degree bound of each variable. For the unknown-order case, Forbes et al. [FSS14] gave an $n^{O(d \log w \log n)}$ -time blackbox test. Observe that their complexity is quasi-polynomial only when d is small. Subsequently, Agrawal et al. [AGKS15] removed the exponential dependence on the individual degree. They gave an $(ndw)^{O(\log n)}$ -time blackbox test for the unknown-order case. Note that these results remain quasi-polynomial even in the case of constant width. Studying ROABPs has also led to PIT results for other computational models, for example, sub-exponential size hitting-sets for depth-3 multilinear circuits [dOSV15] and sub-exponential time whitebox test for read- k oblivious ABPs [AFS⁺15]. It is possible that the results and techniques for ROABPs can help solve the PIT problem for more general models.

Another motivation to study ROABPs comes from their boolean analogues, called read-once ordered branching programs (ROBP). ROBPs have been studied extensively, with regard to the RL versus L question (randomized log-space versus log-space). The problem of finding hitting-sets for ROABP can be viewed as an analogue of finding pseudorandom generators (PRG) for ROBP. A pseudorandom generator for a boolean function f is an algorithm which can generate a probability distribution (with a small sample space) with the property that f cannot distinguish

it from the uniform random distribution (see [AB09] for details). Constructing an optimal PRG for ROBP, i.e., with $O(\log n)$ seed length or polynomial size sample space, would imply $\text{RL} = \text{L}$. This question has similar results as those for PIT of ROABPs, though no connection is known between the two questions. The best known PRG is of seed length $O(\log^2 n)$ ($n^{O(\log n)}$ size sample space), when variable order is known [Nis90, INW94, RR99]. On the other hand, in the unknown-order case, the best known seed length is of size $n^{1/2+o(1)}$ [IMZ12]. Finding an $O(\log n)$ -seed PRG even for constant-width known-order ROBPs has been a challenging open question.

Our first result addresses the analogous question in the arithmetic setting. We give the first polynomial time blackbox test for constant-width known-order ROABPs. However, it works only for zero or large characteristic fields. Our idea is inspired from the pseudorandom construction of Impagliazzo, Nisan and Wigderson [INW94] for ROBPs. While their result does not give better PRGs for the constant-width case, we are able to achieve this in the arithmetic setting.

Theorem (Theorem 3.6). *Let \mathcal{C} be the class of n -variate, individual degree d polynomials in $\mathbb{F}[\mathbf{x}]$ computed by a width- w ROABP in the variable order (x_1, x_2, \dots, x_n) . Then there is a $dn^{O(\log w)}$ -time hitting-set for \mathcal{C} , when $\text{char}(\mathbb{F}) = 0$ or $\text{char}(\mathbb{F}) > ndw^{\log n}$.*

Our test actually works for any width. Its time complexity is better than the previous results on ROABP, when $w < n$ and is same in the other case. Our main technique uses the notion of rank of the partial derivative matrix defined by Nisan [Nis91]. We show that for a nonzero bivariate polynomial $f(x_1, x_2)$ computed by a width- w ROABP, the univariate polynomial $f(t^w, t^w + t^{w-1})$ is nonzero. Our argument is that any bivariate polynomial which becomes zero on $(t^w, t^w + t^{w-1})$ has rank more than w , while a polynomial computed by a width- w ROABP has rank w or less. Then, we use the map $(x_1, x_2) \mapsto (t^w, t^w + t^{w-1})$ recursively in $\log n$ rounds to achieve the above mentioned hitting-set. Our technique has a crucial difference from previous works on ROABPs [FSS14, FS13, AGKS15]. The basic building block in all the previous techniques is a monomial map, i.e., each variable is mapped to a univariate monomial. On the other hand we use a polynomial map. Our approach can potentially lead to a polynomial time hitting-set for ROABPs. The goal would be to obtain a univariate n -tuple $(p_1(t), \dots, p_n(t))$, such that any polynomial which becomes zero on $(p_1(t), \dots, p_n(t))$ must have rank or evaluation dimension higher than w . We conjecture that $(t^r, (t+1)^r, \dots, (t+n-1)^r)$ is one such tuple, where r is polynomially large (Conjecture 3.8).

It is also possible that our ideas for the arithmetic setting can help constructing an optimal PRG for constant-width ROBP.

Our second result is for a special case of ROABPs, called commutative ROABPs. An ROABP is commutative if its edge layers can be exchanged without affecting the polynomial computed. In particular, if all paths from the source to the sink are vertex disjoint, then the ROABP is commutative. Note that for a commutative ROABP, knowing the variable order is irrelevant. Commutative ROABPs have slightly better hitting-sets than the general case, but still no polynomial time hitting-set is known. The previously best known hitting-set for them has time complexity $d^{O(\log w)}(nw)^{O(\log \log w)}$ [FSS14]. We improve this to $(ndw)^{O(\log \log w)}$.

Theorem (Theorem 4.10). *There is an $(ndw)^{O(\log \log w)}$ -time hitting-set for n -variate commutative ROABPs with width w and individual degree d .*

To get this result we follow the approach of Forbes et al. [FSS14], which uses the notion of rank concentration. We achieve rank concentration more efficiently using the basis isolation technique of Agrawal et al. [AGKS15]. The same technique also yields a more efficient concentration in depth-3 set-multilinear circuits (see Section 2 for the definition). However, it is not clear if it gives better hitting-sets for them. The best known hitting-set for them has complexity $n^{O(\log n)}$ [ASS13].

2 Preliminaries

2.1 Definitions and Notations

\mathbb{N} denotes the set of all non-negative integers, i.e., $\{0, 1, 2, \dots\}$. $[n]$ denotes the set $\{1, 2, \dots, n\}$. $\llbracket d \rrbracket$ denotes the set $\{0, 1, \dots, d\}$. \mathbf{x} will denote a set of variables, usually the set $\{x_1, x_2, \dots, x_n\}$. For a set of n variables \mathbf{x} and for an exponent $\mathbf{a} = (a_1, a_2, \dots, a_n) \in \mathbb{N}^n$, $\mathbf{x}^{\mathbf{a}}$ will denote the monomial $\prod_{i=1}^n x_i^{a_i}$. The *support* of a monomial $\mathbf{x}^{\mathbf{a}}$, denoted by $\text{Supp}(\mathbf{a})$, is the set of variables appearing in that monomial, i.e., $\{x_i \mid i \in [n], a_i > 0\}$. The *support size* of a monomial is the cardinality of its support, denoted by $\text{supp}(\mathbf{a})$. A monomial is said to be ℓ -support if its support size is ℓ . For a polynomial $P(\mathbf{x})$, the coefficient of a monomial $\mathbf{x}^{\mathbf{a}}$ in $P(\mathbf{x})$ is denoted by $\text{coef}_P(\mathbf{x}^{\mathbf{a}})$. In particular, $\text{coef}_P(1)$ denotes the constant term of the polynomial P .

For a monomial $\mathbf{x}^{\mathbf{a}}$, $\sum_i a_i$ is said to be its *degree* and a_i is said to be its *degree in variable* x_i for each i . Similarly for a polynomial P , its degree (or degree in x_i) is the maximum degree (or maximum degree in x_i) of any monomial in P with a nonzero coefficient. We define the *individual degree* of P to be $\text{indv-deg}(P) = \max_i \{\text{deg}_{x_i}(P)\}$, where deg_{x_i} denotes degree in x_i .

To better understand polynomials computed by ROABPs, we often use polynomials over an algebra \mathbb{A} , i.e., polynomials whose coefficients come from \mathbb{A} . Matrix algebra is the vector space of matrices equipped with the matrix product. $\mathbb{F}^{m \times n}$ represents the set of all $m \times n$ matrices over the field \mathbb{F} . Note that the algebra of $w \times w$ matrices, has dimension w^2 .

We often view a vector/matrix with polynomial entries, as a polynomial with vector/matrix coefficients. For example,

$$D(x, y) = \begin{pmatrix} 1+x & y-xy \\ x+y & 1+xy \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} 1 + \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} x + \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} y + \begin{pmatrix} 0 & -1 \\ 0 & 1 \end{pmatrix} xy.$$

Here, the coef_D operator will return a matrix for any monomial, for example, $\text{coef}_D(y) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. For a polynomial $D(\mathbf{x}) \in \mathbb{A}[\mathbf{x}]$ over an algebra, its *coefficient space* is the space spanned by its coefficients.

For a matrix R , $R(i, j)$ denotes its entry in the i -th row and j -th column.

As mentioned earlier, a deterministic blackbox PIT is equivalent to constructing a hitting-set. A set of points $\mathcal{H} \in \mathbb{F}^n$ is called a *hitting-set* for a class \mathcal{C} of n -variate polynomials if for any nonzero polynomial P in \mathcal{C} , there exists a point in \mathcal{H} where P evaluates to a nonzero value. An $f(n)$ -time hitting-set would mean that the hitting-set can be generated in time $f(n)$ for input size n .

2.2 Arithmetic Branching Programs

An ABP is a directed graph with $q+1$ layers of vertices $\{V_0, V_1, \dots, V_q\}$ and a start node u and an end node t such that the edges are only going from u to V_0 , V_{i-1} to V_i for any $i \in [q]$ and V_q to t . The edges have univariate polynomials as their weights and as a convention, the edges going from u and those coming to t have weights from the field \mathbb{F} . The ABP is said to compute the polynomial $C(\mathbf{x}) = \sum_{p \in \text{paths}(u,t)} \prod_{e \in p} W(e)$, where $W(e)$ is the weight of the edge e .

The ABP has width w if $|V_i| \leq w$ for all $i \in [q]$. Without loss of generality we can assume $|V_i| = w$ for each $i \in [q]$.

It is well-known that the sum over all paths in a layered graph can be represented by an iterated matrix multiplication. To see this, let the set of nodes in V_i be $\{v_{i,j} \mid j \in [w]\}$. It is easy to see that the polynomial computed by the ABP is the same as $U^T (\prod_{i=1}^q D_i) T$, where $U, T \in \mathbb{F}^{w \times 1}$ and D_i is a $w \times w$ matrix for $1 \leq i \leq q$ such that

$$\begin{aligned} U(\ell) &= W(u, v_{0,\ell}) \text{ for } 1 \leq \ell \leq w \\ D_i(k, \ell) &= W(v_{i-1,k}, v_{i,\ell}) \text{ for } 1 \leq \ell, k \leq w \text{ and } 1 \leq i \leq q \\ T(k) &= W(v_{q,k}, t) \text{ for } 1 \leq k \leq w \end{aligned}$$

2.2.1 Read-once Oblivious ABP

An ABP is called a *read-once oblivious ABP (ROABP)* if the edge weights in different layers are univariate polynomials in distinct variables. Formally, the entries in D_i come from $\mathbb{F}[x_{\pi(i)}]$ for all $i \in [q]$, where π is a permutation on the set $[q]$. Here, q is the same as n , the number of variables. The order $(x_{\pi(1)}, x_{\pi(2)}, \dots, x_{\pi(n)})$ is said to be the variable order of the ROABP.

Viewing $D_i(x_{\pi(i)}) \in \mathbb{F}^{w \times w}[x_{\pi(i)}]$ as a polynomial over the matrix algebra, we can write the polynomial computed by an ROABP as

$$C(\mathbf{x}) = U^\top D_1(x_{\pi(1)}) D_2(x_{\pi(2)}) \cdots D_n(x_{\pi(n)}) T.$$

An equivalent representation of a width- w ROABP can be

$$C(\mathbf{x}) = D_1(x_{\pi(1)}) D_2(x_{\pi(2)}) \cdots D_n(x_{\pi(n)}),$$

where $D_1 \in \mathbb{F}^{1 \times w}[x_{\pi(1)}]$, $D_i \in \mathbb{F}^{w \times w}[x_{\pi(i)}]$ for $2 \leq i \leq n-1$ and $D_n \in \mathbb{F}^{w \times 1}[x_{\pi(n)}]$.

2.2.2 Commutative ROABP

An ROABP $U^\top (\prod_{i=1}^q D_i) T$ is a commutative ROABP, if all D_i s are polynomials over a commutative subalgebra of the matrix algebra. For example, if the coefficients in the polynomials D_i s are all diagonal matrices. Note that the order of the variables becomes insignificant for a commutative ROABP. A polynomial computed by a commutative ROABP can be computed by an ROABP in any variable order.

2.2.3 Set-multilinear Circuits

A depth-3 set-multilinear circuit is a circuit of the form

$$C(\mathbf{x}) = \sum_{i=1}^k l_{i,1}(\mathbf{x}_1) l_{i,2}(\mathbf{x}_2) \cdots l_{i,q}(\mathbf{x}_q),$$

where $l_{i,j}$ s are linear polynomials and $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_q$ form of partition of \mathbf{x} . It is known that these circuits are subsumed by ROABPs [FSS14]. However, they are incomparable to commutative ROABPs. Consider the corresponding polynomial over a k -dimensional algebra

$$D(\mathbf{x}) = D_1(\mathbf{x}_1) D_2(\mathbf{x}_2) \cdots D_q(\mathbf{x}_q),$$

where $D_j = (l_{1,j}, l_{2,j}, \dots, l_{k,j})$ and the algebra product is coordinate-wise product. It is easy to see that $C = (1, 1, \dots, 1) \cdot D$. Note that the polynomials D_i s are over a commutative algebra. Hence, some of our techniques for commutative ROABPs also work for set-multilinear circuits.

3 Hitting-set for Known-order ROABP

3.1 Bivariate ROABP

To construct a hitting-set for ROABPs, we start with the bivariate case. Recall that a bivariate ROABP is of the form $U^\top D_1(x_1) D_2(x_2) T$, where $U, T \in \mathbb{F}^{w \times 1}$, $D_1 \in \mathbb{F}^{w \times w}[x_1]$ and $D_2 \in \mathbb{F}^{w \times w}[x_2]$. It is easy to see that a bivariate polynomial $f(x_1, x_2)$ computed by a width- w ROABP can be written as $f(x_1, x_2) = \sum_{r=1}^w g_r(x_1) h_r(x_2)$. To give a hitting-set for this, we will use the notion of a partial derivative matrix defined by Nisan [Nis91] in the context of lower bounds. Let $f \in \mathbb{F}[x_1, x_2]$ have its individual degree bounded by d . The *partial derivative matrix* M_f for f is a $(d+1) \times (d+1)$ matrix with

$$M_f(i, j) = \text{coef}_f(x_1^i x_2^j) \in \mathbb{F},$$

for all $i, j \in [d]$. It is known that the rank of M_f is equal to the smallest possible width of an ROABP computing f [Nis91].

Lemma 3.1 (rank \leq width). *For any polynomial $f(x_1, x_2) = \sum_{r=1}^w g_r(x_1)h_r(x_2)$, $\text{rank}(M_f) \leq w$.*

Proof. Let us define $f_r = g_r h_r$, for all $r \in [w]$. Clearly, $M_f = \sum_{r=1}^w M_{f_r}$, as $f = \sum_{r=1}^w f_r$. We will show that $\text{rank}(M_{f_r}) \leq 1$, for all $r \in [w]$. As $f_r = g_r(x_1)h_r(x_2)$, its coefficients can be written as a product of coefficients from g_r and h_r , i.e.,

$$\text{coef}_{f_r}(x_1^i x_2^j) = \text{coef}_{g_r}(x_1^i) \text{coef}_{h_r}(x_2^j).$$

Now, it is easy to see that

$$M_{f_r} = u_r v_r^\top,$$

where $u_r, v_r \in \mathbb{F}^{d+1}$ with $u_r = (\text{coef}_{g_r}(x_1^i))_{i=0}^d$ and $v_r = (\text{coef}_{h_r}(x_2^i))_{i=0}^d$.

Thus, $\text{rank}(M_{f_r}) \leq 1$ and $\text{rank}(M_f) \leq w$. \square

One can also show that if $\text{rank}(M_f) = w$ then there exists a width- w ROABP computing f . We skip this proof as we will not need it. Now, using the above lemma we give a hitting-set for bivariate ROABPs.

Lemma 3.2. *Let $\text{char}(\mathbb{F}) = 0$, or $\text{char}(\mathbb{F}) > d$. Let $f(x_1, x_2) = \sum_{r=1}^w g_r(x_1)h_r(x_2)$ be a nonzero bivariate polynomial over \mathbb{F} with individual degree d . Then $f(t^w, t^w + t^{w-1}) \neq 0$.*

Proof. Let $f'(t)$ be the polynomial after the substitution, i.e., $f' = f(t^w, t^w + t^{w-1})$. Any monomial $x_1^i x_2^j$ will be mapped to the polynomial $t^{wi}(t^w + t^{w-1})^j$, under the mentioned substitution. The highest power of t coming from this polynomial is $t^{w(i+j)}$. We will cluster together all the monomials for which this highest power is the same, i.e., $i + j$ is the same. The coefficients corresponding to any such cluster of monomials will form a *diagonal* in M_f . The set $\{M_f(i, j) \mid i + j = k\}$ is defined to be the k -th *diagonal* of M_f , for all $0 \leq k \leq 2d$. Let ℓ be the highest number such that ℓ -th diagonal has at least one nonzero element, i.e.,

$$\ell = \max\{i + j \mid M_f(i, j) \neq 0\}.$$

As $\text{rank}(M_f) \leq w$ (from Lemma 3.1), we claim that the ℓ -th diagonal has at most w nonzero elements. To see this, let $\{(i_1, j_1), (i_2, j_2), \dots, (i_{w'}, j_{w'})\}$ be the set of indices where the ℓ -th diagonal of M_f has nonzero elements, i.e., the set $\{(i, j) \mid M_f(i, j) \neq 0, i + j = \ell\}$. As $M_f(i, j) = 0$ for any $i + j > \ell$, it is easy to see that the rows $\{M_f(i_1), M_f(i_2), \dots, M_f(i_{w'})\}$ are linearly independent. Thus, $w' \leq \text{rank}(M_f) \leq w$.

Now, we claim that there exists an r with $w(\ell - 1) < r \leq w\ell$ such that $\text{coef}_{f'}(t^r) \neq 0$. To see this, first observe that the highest power of t which any monomial $x_1^i x_2^j$ with $i + j < \ell$ can contribute is $t^{w(\ell-1)}$. Thus, for any $w(\ell - 1) < r \leq w\ell$, the term t^r can come only from the monomials $x_1^i x_2^j$ with $i + j \geq \ell$. We can ignore the monomials $x_1^i x_2^j$ with $i + j > \ell$ as $\text{coef}_{f'}(x_1^i x_2^j) = M_f(i, j) = 0$, when $i + j > \ell$. Now, for any $i + j = \ell$, the monomial $x_1^i x_2^j$ goes to

$$t^{w(\ell-j)}(t^w + t^{w-1})^j = \sum_{p=0}^j \binom{j}{p} t^{w\ell-p}.$$

Hence, for any $0 \leq p < w$,

$$\text{coef}_{f'}(t^{w\ell-p}) = \sum_{a=1}^{w'} M_f(i_a, j_a) \binom{j_a}{p}.$$

Writing this in the matrix form we get

$$[\text{coef}_{f'}(t^{w\ell}) \cdots \text{coef}_{f'}(t^{w\ell-w+1})] = [M_f(i_1, j_1) \cdots M_f(i_{w'}, j_{w'})]C,$$

where C is a $w' \times w$ matrix with $C(a, b) = \binom{j_a}{b-1}$, for all $a \in [w']$ and $b \in [w]$. If all the rows of C are linearly independent then clearly, $\text{coef}_{f'}(t^r) \neq 0$ for some $w(\ell - 1) < r \leq w\ell$. We show the linear independence in Claim 3.3. To show this linear independence we need to assume that the numbers $\{j_a\}_a$ are all distinct. Hence, we need the field characteristic to be zero or strictly greater than d , as j_a can be as high as d for some $a \in [w']$.

Claim 3.3. *Let C be a $w \times w$ matrix with $C(a, b) = \binom{j_a}{b-1}$, for all $a \in [w]$ and $b \in [w]$, where $\{j_a\}_a$ are all distinct numbers. Then C has full rank.*

Proof. We will show that for any nonzero vector $\alpha := (\alpha_1, \alpha_2, \dots, \alpha_w) \in \mathbb{F}^{w \times 1}$, $C\alpha \neq 0$. Consider the polynomial $h(y) = \sum_{b=1}^w \alpha_b \frac{y(y-1)\dots(y-b+2)}{(b-1)!}$. As $h(y)$ is a nonzero polynomial with degree bounded by $w - 1$, it can have at most $w - 1$ roots. Thus, there exists an $a \in [w]$ such that $h(j_a) = \sum_{b=1}^w \alpha_b \binom{j_a}{b-1} \neq 0$. \square

\square

As mentioned above, the hitting-set proof works only when the field characteristic is zero or greater than d . We given an example over a small characteristic field, which demonstrates that the problem is not with the proof technique, but with the hitting-set itself. Let the field characteristic be 2. Consider the polynomial $f(x_1, x_2) = x_2^2 + x_1^2 + x_1$. Clearly, f has a width-2 ROABP. For a width-2 ROABP, the map in Lemma 3.2 would be $(x_1, x_2) \mapsto (t^2, t^2 + t)$. However, $f(t^2, t^2 + t) = 0$ (over \mathbb{F}_2). Hence, the hitting-set does not work.

Now, we move on to getting a hitting-set for an n -variate ROABP.

3.2 n -variate ROABP

Observe that the map given in Lemma 3.2 works irrespective of the degree of the polynomial, as long as the field characteristic is large enough. We plan to obtain a hitting-set for general n -variate ROABP by applying this map recursively. For this, we use the standard divide and conquer technique. First, we make pairs of consecutive variables in the ROABP. For each pair (x_{2i-1}, x_{2i}) , we apply the map from Lemma 3.2, using a new variable t_i . Thus, we go to $n/2$ variables from n variables. In Lemma 3.4, we show that after this substitution the polynomial remains nonzero. Moreover, the new polynomial can be computed by a width- w ROABP. Thus, we can again use the same map on pairs of new variables. By repeating the halving procedure $\log n$ times we get a univariate polynomial. In each round the degree of the polynomial gets multiplied by w . Hence, after $\log n$ rounds, the degree of the univariate polynomial is bounded by $w^{\log n}$ times the original degree. Without loss of generality, let us assume that n is a power of 2.

Lemma 3.4 (Halving the number of variables). *Let $\text{char}(\mathbb{F}) = 0$, or $\text{char}(\mathbb{F}) > d$. Let $f(\mathbf{x}) = D_1(x_1)D_2(x_2) \cdots D_n(x_n)$ be a nonzero polynomial computed by a width- w and individual degree- d ROABP, where $D_1 \in \mathbb{F}^{1 \times w}[x_1]$, $D_n \in \mathbb{F}^{w \times 1}[x_n]$ and $D_i \in \mathbb{F}^{w \times w}[x_i]$ for all $2 \leq i \leq n - 1$. Let the map $\phi: \mathbf{x} \rightarrow \mathbb{F}[\mathbf{t}]$ be such that for any index $1 \leq i \leq n/2$,*

$$\begin{aligned} \phi(x_{2i-1}) &= t_i^w, \\ \phi(x_{2i}) &= t_i^w + t_i^{w-1}. \end{aligned}$$

Then $f(\phi(\mathbf{x})) \neq 0$. Moreover, the polynomial $f(\phi(\mathbf{x})) \in \mathbb{F}[t_1, t_2, \dots, t_{n/2}]$ is computed by a width- w ROABP in the variable order $(t_1, t_2, \dots, t_{n/2})$.

Proof. Let us apply the map in $n/2$ rounds, i.e., define a sequence of polynomials ($f = f_0, f_1, \dots, f_{n/2} = f(\phi(\mathbf{x}))$) such that the polynomial f_i is obtained by making the replacement $(x_{2i-1}, x_{2i}) \mapsto (\phi(x_{2i-1}), \phi(x_{2i}))$ in f_{i-1} for each $1 \leq i \leq n/2$. We will show that for each $1 \leq i \leq n/2$, if $f_{i-1} \neq 0$ then $f_i \neq 0$. Clearly this proves the first part of the lemma.

Note that f_{i-1} is a polynomial over variables $\{t_1, \dots, t_{i-1}, x_{2i-1}, \dots, x_n\}$. As $f_{i-1} \neq 0$, there exists a constant tuple $\alpha \in \mathbb{F}^{n-i-1}$ such that after replacing the variables $(t_1, \dots, t_{i-1}, x_{2i+1}, \dots, x_n)$ with α , f_{i-1} remains nonzero. After this replacement we get a polynomial f'_{i-1} in the variables (x_{2i-1}, x_{2i}) . As f is computed by the ROABP $D_1 D_2 \cdots D_n$, the polynomial f'_{i-1} can be written as $U^T D_{2i-1}(x_{2i-1}) D_{2i}(x_{2i}) T$ for some $U, T \in \mathbb{F}^{w \times 1}$. In other words, f'_{i-1} has a bivariate ROABP of width w . Thus, $f'_{i-1}(\phi(x_{2i-1}), \phi(x_{2i}))$ is nonzero from Lemma 3.2. But, $f'_{i-1}(\phi(x_{2i-1}), \phi(x_{2i}))$ is nothing but the polynomial obtained after replacing the variables $(t_1, \dots, t_{i-1}, x_{2i+1}, \dots, x_n)$ in f_i with α . Thus, f_i is nonzero. This finishes the proof.

Now, we argue that $f(\phi(\mathbf{x}))$ has a width w ROABP. Let $D'_i := D_{2i-1}(t_i^w) D_{2i}(t_i^w + t_i^{w-1})$ for all $1 \leq i \leq n/2$. Clearly, $D'_1 D'_2 \cdots D'_{n/2}$ is an ROABP computing $f(\phi(\mathbf{x}))$ in variable order $(t_1, t_2, \dots, t_{n/2})$, as $D'_1 \in \mathbb{F}^{1 \times w}[t_1]$, $D'_{n/2} \in \mathbb{F}^{w \times 1}[t_{n/2}]$ and $D'_i \in \mathbb{F}^{w \times w}[t_i]$ for all $2 \leq i \leq n/2 - 1$. \square

By applying the map ϕ in Lemma 3.4, we reduced an n -variate ROABP to an $(n/2)$ -variate ROABP, while preserving the non-zerosness. The resulting ROABP has same width w , but the individual degree goes up to become $2dw$, where d is the original individual degree. As our map ϕ is degree independent, we can apply the same map again on the variables $\{t_i\}_{i=1}^{n/2}$. It is easy to see that when the map ϕ is repeatedly applied in this way $\log n$ times, we get a nonzero univariate polynomial of degree $ndw^{\log n}$. Next lemma puts it formally. For ease of notation, we use the variable numbering from 0 to $n-1$. Let $p_0(t) = t^w$ and $p_1(t) = t^w + t^{w-1}$.

Lemma 3.5. *Let $\text{char}(\mathbb{F}) = 0$, or $\text{char}(\mathbb{F}) \geq ndw^{\log n}$. Let $f \in \mathbb{F}[\mathbf{x}]$ be a nonzero polynomial, with individual degree d , computed by a width- w ROABP in variable order $(x_0, x_1, \dots, x_{n-1})$. Let the map $\phi: \{x_0, x_1, \dots, x_{n-1}\} \rightarrow \mathbb{F}[t]$ be such that for any index $0 \leq i \leq n-1$,*

$$\phi(x_i) = p_{i_1}(p_{i_2} \cdots (p_{i_{\log n}}(t))),$$

where $i_{\log n} i_{\log n-1} \cdots i_1$ is the binary representation of i .

Then $f(\phi(\mathbf{x}))$ is a nonzero univariate polynomial with degree $ndw^{\log n}$.

Note that the map ϕ crucially uses the knowledge of the variable order. In the last round when we are going from two variables to one, the individual degree is $ndw^{\log n-1}$ and Lemma 3.2 requires $\text{char}(\mathbb{F})$ to be higher than the individual degree. Thus, having $\text{char}(\mathbb{F}) \geq ndw^{\log n}$ suffices. For a univariate polynomial, the standard hitting-set is to plug-in distinct field values as many as one more than the degree. Thus, we get the following theorem.

Theorem 3.6. *For an n -variate, individual degree d and width- w ROABP, there is a blackbox PIT with time complexity $O(ndw^{\log n})$, when the variable order is known and the field characteristic is zero or at least $ndw^{\log n}$.*

From this, we immediately get the following result for constant-width ROABPs. Note that when w is constant, the lower bound on the characteristic also becomes $\text{poly}(n)$.

Corollary 3.7. *There is a polynomial time blackbox PIT for constant width ROABPs, with known variable order and field characteristic being zero (or polynomially large).*

As mentioned earlier, our approach can potentially lead to a polynomial time hitting-set for ROABPs. We make the following conjecture for which we hope to get a proof on the lines of Lemma 3.2.

Conjecture 3.8. *Let $\text{char}(\mathbb{F}) = 0$. Let $f(\mathbf{x}) \in \mathbb{F}[\mathbf{x}]$ be an n -variate, degree- d polynomial computed by a width- w ROABP. Then $f(t^r, (t+1)^r, \dots, (t+n-1)^r) \neq 0$ for some r bounded by $\text{poly}(n, w, d)$.*

4 Commutative ROABP

In this section, we give better hitting-sets for commutative ROABPs. Recall that an ROABP is commutative if the matrices involved in the matrix product come from a commutative algebra. To elaborate, a commutative ROABP is of the form $U^\top D_1 D_2 \cdots D_n T$, where $U, T \in \mathbb{F}^{w \times 1}$ and $D_i \in \mathbb{F}^{w \times w}[x_i]$ is a polynomial over a commutative subalgebra of $\mathbb{F}^{w \times w}$ for each i . In simple words, $D_i D_j = D_j D_i$ for any $i, j \in [n]$. As the order of variables does not matter for a commutative ROABP, we take the standard variable order (x_1, x_2, \dots, x_n) . Here we work with the polynomial $D = D_1 D_2 \cdots D_n$ over the matrix algebra. With an abuse of notation, we say $D_1 D_2 \cdots D_n$ is an ROABP computing a polynomial over matrices.

Forbes et al. [FSS14] gave a $d^{O(\log w)}(ndw)^{O(\log \log w)}$ -time hitting-set for width- w , n -variate commutative ROABPs with individual degree bound d . Note that when d is small, this time complexity is much better than that for general ROABP, i.e., $(ndw)^{O(\log n)}$ [AGKS15]. However when d is $O(n)$, the complexity is comparable to the general case. We improve the time complexity for the commutative case to $(ndw)^{O(\log \log w)}$. This is significantly better than the general case for all values of d .

Forbes et al. [FSS14] constructed the hitting-set using the notion of rank-concentration defined by Agrawal et al. [ASS13].

Definition 4.1 ([ASS13]). *A polynomial $D(\mathbf{x})$ over an algebra is said to be ℓ -concentrated if its coefficients of $(< \ell)$ -support monomials span all its coefficients.*

Note that for a polynomial in $\mathbb{F}[\mathbf{x}]$, ℓ -concentration simply means that it has a monomial of $(< \ell)$ -support with a nonzero coefficient. For a polynomial which has low-support concentration, it is easy to construct hitting-sets. However, not every polynomial has a low-support concentration, for example $C(\mathbf{x}) = x_1 x_2 \cdots x_n$. Agrawal et al. [ASS13] observed that concentration can be achieved by a shift of variables, e.g., $C(\mathbf{x} + \mathbf{1}) = (x_1 + 1)(x_2 + 1) \cdots (x_n + 1)$ has 1-concentration. For a polynomial $C(\mathbf{x})$, shift by a tuple $\mathbf{f} = (f_1, f_2, \dots, f_n)$ would mean $C(\mathbf{x} + \mathbf{f}) = C(x_1 + f_1, x_2 + f_2, \dots, x_n + f_n)$. The first step of Forbes et al. [FSS14] is to show that for a given commutative width- w ROABP, $O(\log w)$ -concentration can be achieved by a shift with cost $nd^{O(\log w)}$. Their second step is to show that if a given commutative ROABP is $O(\log w)$ -concentrated then there is a hitting-set for it of size $(ndw)^{O(\log \log w)}$. We improve the first step by giving a shift with cost $(ndw)^{O(\log \log w)}$, which gives us the desired hitting-set.

First, we elaborate the first step of Forbes, Saptharishi and Shpilka [FSS14]. To achieve concentration they use the idea of Agrawal, Saha and Saxena [ASS13], i.e., achieving concentration in small sub-circuits implies concentration in the whole circuit. For the sake of completeness, we rewrite the lemma using the terminology of this paper.

Lemma 4.2 ([ASS13, FSS14]). *Let $D(\mathbf{x}) = D_1(x_1)D_2(x_2) \cdots D_n(x_n)$ be a product of univariate polynomials over a commutative algebra \mathbb{A}_k . Suppose there exists an ℓ such that for any $S \in [n]$ with $|S| = \ell$, the polynomial $\prod_{i \in S} D_i$ has ℓ -concentration. Then $D(\mathbf{x})$ has ℓ -concentration.*

Proof. For any set $S \subseteq [n]$, let us define a sub-circuit D_S of D as $\prod_{i \in S} D_i(x_i)$. We will show ℓ -concentration in all the sub-circuits D_S of D , using induction on the size of S .

Base Case: D_S is trivially ℓ -concentrated if $|S| < \ell$. In the case of $|S| = \ell$, D_S is ℓ -concentrated from the hypothesis in the lemma.

Induction Hypothesis: D_S has ℓ -concentration for any set S with $|S| < j$.

Induction Step: We will prove ℓ -concentration in D_S for a set S with $|S| = j$. Let $S = \{x_{i_1}, x_{i_2}, \dots, x_{i_j}\}$. Consider a monomial $\mathbf{x}^{\mathbf{a}} = x_{i_1}^{a_1} x_{i_2}^{a_2} \cdots x_{i_j}^{a_j}$ with support from the set S . Without loss of generality let us assume $a_1 \neq 0$. Now, let the set $S' = S \setminus \{x_{i_1}\}$ and let the monomial $\mathbf{x}^{\mathbf{a}'} = \mathbf{x}^{\mathbf{a}} / x_{i_1}^{a_1}$. As $|S'| = j - 1$, by the inductive hypothesis $D_{S'}$ is ℓ -concentrated. Thus,

$$\text{coef}_{D_{S'}}(\mathbf{x}^{\mathbf{a}'}) \in \text{span}\{\text{coef}_{D_{S'}}(\mathbf{x}^{\mathbf{b}}) \mid \text{Supp}(\mathbf{b}) \subseteq S', \text{supp}(\mathbf{b}) < \ell\}. \quad (1)$$

It is easy to see that for any monomial $\mathbf{x}^{\mathbf{b}}$ with its support in S' ,

$$\text{coef}_{D_S}(\mathbf{x}^{\mathbf{b}} x_{i_1}^{a_1}) = \text{coef}_{D_{S'}}(\mathbf{x}^{\mathbf{b}}) \text{coef}_{D_{i_1}}(x_{i_1}^{a_1}).$$

Thus, by multiplying $\text{coef}_{D_{i_1}}(x_{i_1}^{a_1})$ in (1), we get

$$\text{coef}_{D_S}(\mathbf{x}^{\mathbf{a}}) \in \text{span}\{\text{coef}_{D_S}(\mathbf{x}^{\mathbf{b}} x_{i_1}^{a_1}) \mid \text{Supp}(\mathbf{b}) \subseteq S', \text{supp}(\mathbf{b}) < \ell\}.$$

Hence,

$$\text{coef}_{D_S}(\mathbf{x}^{\mathbf{a}}) \in \text{span}\{\text{coef}_{D_S}(\mathbf{x}^{\mathbf{b}}) \mid \text{Supp}(\mathbf{b}) \subseteq S, \text{supp}(\mathbf{b}) \leq \ell\}. \quad (2)$$

Now, we claim that for any monomial $\mathbf{x}^{\mathbf{b}}$ with $\text{Supp}(\mathbf{b}) \subseteq S$ and $\text{supp}(\mathbf{b}) = \ell$,

$$\text{coef}_{D_S}(\mathbf{x}^{\mathbf{b}}) \in \text{span}\{\text{coef}_{D_S}(\mathbf{x}^{\mathbf{c}}) \mid \text{Supp}(\mathbf{c}) \subseteq S, \text{supp}(\mathbf{c}) < \ell\}. \quad (3)$$

To see this, let T be the support of the monomial $\mathbf{x}^{\mathbf{b}}$. As $|T| = \ell$, D_T has ℓ -concentration. Thus,

$$\text{coef}_{D_T}(\mathbf{x}^{\mathbf{b}}) \in \text{span}\{\text{coef}_{D_T}(\mathbf{x}^{\mathbf{c}}) \mid \text{Supp}(\mathbf{c}) \subseteq T, \text{supp}(\mathbf{c}) < \ell\}. \quad (4)$$

For any monomial $\mathbf{x}^{\mathbf{c}}$ with support in T , one can write

$$\text{coef}_{D_S}(\mathbf{x}^{\mathbf{c}}) = \text{coef}_{D_T}(\mathbf{x}^{\mathbf{c}}) \prod_{i \in S \setminus T} \text{coef}_{D_i}(1).$$

Note that the commutativity of the underlying algebra is crucial for this. Thus, multiplying (4) by $\left(\prod_{i \in S \setminus T} \text{coef}_{D_i}(1)\right)$, we get (3).

By combining (3) with (2), we get

$$\text{coef}_{D_S}(\mathbf{x}^{\mathbf{a}}) \in \text{span}\{\text{coef}_{D_S}(\mathbf{x}^{\mathbf{c}}) \mid \text{Supp}(\mathbf{c}) \subseteq S, \text{supp}(\mathbf{c}) < \ell\},$$

for any monomial $\mathbf{x}^{\mathbf{a}}$ with $\text{Supp}(\mathbf{a}) \subseteq S$. This proves ℓ -concentration in D_S .

Taking $S = [n]$, we get ℓ -concentration in D . \square

Now, the goal is just to achieve ℓ -concentration in an ℓ -variate ROABP (computing a polynomial over the matrix algebra). We would remark here that for an ℓ -variate polynomial over a k -dimensional algebra, one can hope to achieve ℓ -concentration only when $\ell \geq \log(k+1)$. To see this, consider the polynomial $D(\mathbf{x}) = \prod_{i=1}^{\ell} (1 + v_i x_i)$ over a k -dimensional algebra such that $k > 2^{\ell} - 1$. Suppose the vector v_i s are such that all the 2^{ℓ} coefficients of the polynomial D are linearly independent. There are only $2^{\ell} - 1$ coefficients of D with $(< \ell)$ -support. Hence, they cannot span the whole coefficient space of D , whatever the shift we use.

Agrawal et al. [ASS13] and Forbes et al. [FSS14] achieve ℓ -concentration in arbitrary ℓ -variate polynomials over a k -dimension algebra for $\ell = \log(k+1)$ by a shift with cost $d^{O(\ell)}$, where d is the individual degree. Forbes et al. [FSS14] use it to give a single shift on n variables such that it works for any choice of ℓ variables. This has cost $nd^{O(\ell)}$.

We give a new shift with cost $(ndw)^{O(\log \ell)} = (ndw)^{O(\log \log w)}$, for a width- w , ℓ -variate ROABP (w^2 is the dimension of the underlying algebra). The cost has n as a parameter because the shift works for any size ℓ subset of n variables. Like [ASS13, FSS14], we use a shift by univariate polynomials in a new variable t . In this case, the concentration is considered over the field $\mathbb{F}(t)$. Note that while the shift of [ASS13, FSS14] works for an arbitrary ℓ -variate polynomial, our shift works only for ℓ -variate ROABPs. The univariate map we use is the basis isolating weight assignment for ROABPs from Agrawal et al. [AGKS15]. We simply use the fact that for any polynomial over a k -dimensional algebra, shift by a basis isolating map achieves $\log(k+1)$ -concentration [GKST15].

Let us first recall the definition of a basis isolating weight assignment. Let M denote the set of all monomials over the variable set \mathbf{x} with individual degree $\leq d$. Any function $w: \mathbf{x} \rightarrow \mathbb{N}$ can be naturally extended to the set of all monomials as follows: $w(\prod_{i=1}^n x_i^{\gamma_i}) = \sum_{i=1}^n \gamma_i w(x_i)$, for any $(\gamma_i)_{i=1}^n \in \mathbb{N}^n$. Note that if the variable x_i is replaced with $t^{w(x_i)}$ for each i , then any monomial m just becomes $t^{w(m)}$. \mathbb{A}_k denotes a k -dimensional algebra.

Definition 4.3 ([AGKS15]). A weight function $w: \mathbf{x} \rightarrow \mathbb{N}$ is called a *basis isolating weight assignment* for a polynomial $D(\mathbf{x}) \in \mathbb{A}_k[\mathbf{x}]$, if there exists a set of monomials $S \subseteq M$ ($k' := |S| \leq k$) whose coefficients form a basis for the coefficient space of $D(\mathbf{x})$, such that

- for any $m, m' \in S$, $w(m) \neq w(m')$ and
- for any monomial $m \in M \setminus S$,

$$\text{coef}_D(m) \in \text{span}\{\text{coef}_D(m') \mid m' \in S, w(m') < w(m)\}.$$

Gurjar et al. [GKST15, Lemma 5.2] have shown that shifting by a basis isolating weight assignment achieves concentration.

Lemma 4.4 (Isolation to concentration). Let $A(\mathbf{x})$ be a polynomial over a k -dimensional algebra \mathbb{A}_k . Let w be a basis isolating weight assignment for $A(\mathbf{x})$. Then $A(\mathbf{x} + t^w)$ is ℓ -concentrated, where $\ell = \lceil \log(k+1) \rceil$ and t^w denotes the n -tuple $(t^{w(x_1)}, t^{w(x_2)}, \dots, t^{w(x_n)})$.

We now recall the construction of a basis isolating weight assignment for ROABP from [AGKS15]. Here, we present a slightly modified version of their Lemma 8, which easily follows from it.

Lemma 4.5. Let \mathbf{x} be a set of n variables. Let $D(\mathbf{x}) = D_1(x_{i_1})D_2(x_{i_2}) \cdots D_\ell(x_{i_\ell})$ be an ℓ -variate polynomial over a k -dimensional algebra \mathbb{A}_k . Then we can construct a basis isolating weight assignment for $D(\mathbf{x})$ with the cost being $(\text{poly}(k, n, d))^{\log \ell}$, where d is the individual degree.

The construction in [AGKS15, Lemma 8] actually gives a family \mathcal{B} of $(knd)^{O(\log \ell)}$ weight assignments such that for any ℓ -variate ROABP, at least one of them is basis isolating. However, we are interested in a single map which works for every ℓ -variate ROABP. To get a single shift for every ROABP, we follow the technique of [FSS14, GKST15] and take a Lagrange Interpolation of all the n -tuples in the family $\{t^w\}_{w \in \mathcal{B}}$.

Let $\mathcal{F} = \{\mathbf{f}_1(t), \mathbf{f}_2(t), \dots, \mathbf{f}_N(t)\}$ be this family of n -tuples, where $\mathbf{f}_i = \{f_{i,1}(t), f_{i,2}(t), \dots, f_{i,n}(t)\}$ for each i . Here, $N = (knd)^{O(\log \ell)}$. Let their degrees be bounded by D , i.e., $D = \max\{\deg(f_{i,j}) \mid i \in [N] \text{ and } j \in [n]\}$. From the construction in [AGKS15], $D = (knd)^{O(\log \ell)}$. Also, the family \mathcal{F} can be generated in time $(knd)^{O(\log \ell)}$.

Let $\mathbf{L}(y, t) \in \mathbb{F}[y, t]^n$ be the Lagrange interpolation of \mathcal{F} . That is, for all $j \in [n]$,

$$L_j = \sum_{i \in [N]} f_{i,j}(t) \prod_{\substack{i' \in [N] \\ i' \neq i}} \frac{y - \alpha_{i'}}{\alpha_i - \alpha_{i'}},$$

where $\{\alpha_i\}_{i \in [N]}$ are distinct field elements (we go to a large enough field extension where these many elements exist). Note that $L_j|_{y=\alpha_i} = f_{i,j}$. Thus, $\mathbf{L}|_{y=\alpha_i} = \mathbf{f}_i$. Also, $\deg_y(L_j) = N - 1$ and $\deg_t(L_j) \leq D$. The following lemma from [GKST15, Lemma 5.5] shows that a shift by the interpolation works for every polynomial simultaneously.

Lemma 4.6. Let $A(\mathbf{x})$ be a polynomial over \mathbb{A}_k such that there exists an $\mathbf{f} \in \mathcal{F}$ for which $A'(\mathbf{x}, t) = A(\mathbf{x} + \mathbf{f}) \in \mathbb{A}_k(t)[\mathbf{x}]$ is ℓ -concentrated. Then, $A''(\mathbf{x}, y, t) = A(\mathbf{x} + \mathbf{L}) \in \mathbb{A}_k(y, t)[\mathbf{x}]$ is ℓ -concentrated.

Proof. Let $\text{rank}_{\mathbb{F}}\{\text{coef}_A(\mathbf{x}^{\mathbf{a}}) \mid \mathbf{x}^{\mathbf{a}} \in M\} = k'$, for some $k' \leq k$, and $M_\ell = \{\mathbf{x}^{\mathbf{a}} \in M \mid \text{supp}(\mathbf{a}) < \ell\}$. We need to show that $\text{rank}_{\mathbb{F}(y,t)}\{\text{coef}_{A''}(\mathbf{x}^{\mathbf{a}}) \mid \mathbf{x}^{\mathbf{a}} \in M_\ell\} = k'$.

Since $A'(\mathbf{x})$ is ℓ -concentrated, we have that $\text{rank}_{\mathbb{F}(t)}\{\text{coef}_{A'}(\mathbf{x}^{\mathbf{a}}) \mid \mathbf{x}^{\mathbf{a}} \in M_\ell\} = k'$. Recall that $A'(\mathbf{x})$ is an evaluation of A'' at $y = \alpha_i$, i.e., $A'(\mathbf{x}, t) = A''(\mathbf{x}, \alpha_i, t)$ for some α_i . Thus, for all $\mathbf{x}^{\mathbf{a}} \in M$, we have $\text{coef}_{A'}(\mathbf{x}^{\mathbf{a}}) = \text{coef}_{A''}(\mathbf{x}^{\mathbf{a}})|_{y=\alpha_i}$.

Let $C \in \mathbb{F}[t]^{k \times |M_\ell|}$ be the matrix whose columns are $\text{coef}_{A'}(\mathbf{x}^\alpha)$, for $\mathbf{x}^\alpha \in M_\ell$. Let similarly $C' \in \mathbb{F}[y, t]^{k \times |M_\ell|}$ be the matrix whose columns are $\text{coef}_{A''}(\mathbf{x}^\alpha)$, for $\mathbf{x}^\alpha \in M_\ell$. Then we have $C = C'|_{y=\alpha_i}$.

As $\text{rank}_{\mathbb{F}(t)}(C) = k'$, there is a $k' \times k'$ submatrix in C , say indexed by (R, T) , such that $\det(C(R, T)) \neq 0$. Since $\det(C(R, T)) = \det(C'(R, T))|_{y=\alpha_i}$, it follows that $\det(C'(R, T)) \neq 0$. Hence, we have $\text{rank}_{\mathbb{F}(y, t)}(C') = k'$. Thus, the $(< \ell)$ -support coefficients of A'' span its coefficient space. \square

Hence, the Lagrange interpolation gives us a single shift which works for all ℓ -variate ROABPs.

Lemma 4.7. *Given n, d, w and $\ell = \log(w^2 + 1)$, in time $(ndw)^{O(\log \ell)}$ one can compute a polynomial tuple $\mathbf{f}(t) \in \mathbb{F}[t]^n$ of degree $(ndw)^{O(\log \ell)}$ such that for any ℓ -variate polynomial $A(\mathbf{x}) \in \mathbb{F}^{w \times w}[\mathbf{x}]$ of individual degree d that can be computed by an ROABP of width w , the polynomial $A(\mathbf{x} + \mathbf{f}(t))$ is ℓ -concentrated.*

Proof. Note that the dimension k of the underlying algebra is bounded by w^2 . After shifting the polynomial $A(\mathbf{x})$ by $\mathbf{L}(y, t)$ as defined above, its coefficients will be polynomials in y and t , with degree $d' = dn(ndw)^{O(\log \ell)}$. Consider the determinant polynomial $\det(C'(R, T))$ from the proof of Lemma 4.6. As $k' \leq k$, $\det(C'(R, T))$ has degree bounded by $d'' := kd'$. So, when we replace y by $t^{d''+1}$, it does not affect the non-zerosness of the determinant, and hence, the concentration is preserved. Thus, $\mathbf{f} = \mathbf{L}(t^{d''+1}, t)$ is an n -tuple of univariate polynomials in t that fulfils the claim of the lemma. \square

Combining Lemma 4.2 and Lemma 4.7 we get the following.

Lemma 4.8. *Given n, d, w , one can compute an n -tuple $\mathbf{f}(t)$ with cost $(ndw)^{O(\log \log w)}$ such that for any n -variate, individual degree- d polynomial $D(\mathbf{x}) \in \mathbb{F}^{w \times w}[\mathbf{x}]$ computed by a width- w commutative ROABP, $D(\mathbf{x} + \mathbf{f}(t))$ is $O(\log w)$ -concentrated.*

Note that if the polynomial $D(\mathbf{x}) \in \mathbb{F}^{w \times w}[\mathbf{x}]$ is ℓ -concentrated then the polynomial $C(\mathbf{x}) = U^\top D T$ is also ℓ -concentrated, where $U, T \in \mathbb{F}^{w \times 1}$. This is true because multiplication by U^\top and T are linear operations. Recall that for polynomial $C(\mathbf{x}) \in \mathbb{F}[\mathbf{x}]$, $O(\log w)$ -concentration means that there is a monomial with $O(\log w)$ -support which has a nonzero coefficient.

Lemma 4.8 gives a shift $\mathbf{f}(t)$ of univariate polynomials. To get a constant shift, we substitute $(ndw)^{O(\log \log w)}$ distinct values for t . As the degree in t is bounded by $(ndw)^{O(\log \log w)}$, at least for one value of t , the non-zerosness of the particular coefficient will be preserved.

Now, we move on to the second step of Forbes, Shpilka and Saptharishi [FSS14]. They give an $(ndw)^{O(\log \log w)}$ -time hitting-set for an already $O(\log w)$ -concentrated commutative ROABP. They do this by reducing the PIT question to an $O(\log w)$ -variate ROABP [FSS14, Lemma 7.6].

Lemma 4.9 ([FSS14]). *Let $C(\mathbf{x}) \in \mathbb{F}[\mathbf{x}]$ be an n -variate, individual degree- d polynomial computed by a width- w ROABP. Suppose $C(\mathbf{x})$ has an $(\leq \ell)$ -support monomial with a nonzero coefficient. Then, there is a $\text{poly}(n, w, d)$ -time computable m -variate map $\phi: \mathbf{x} \rightarrow \mathbb{F}[y_1, y_2, \dots, y_m]$ such that $C(\phi(\mathbf{x}))$ is a nonzero polynomial with degree $< d^2 n^4$, where $m = O(\ell^2)$. Moreover, $C(\phi(\mathbf{x}))$ is computed by a width- w , m -variate commutative ROABP.*

From the results of [FS13, AGKS15], we know that an m -variate, width- w commutative ROABP has an $(mdw)^{O(\log m)}$ -time hitting-set. Combining Lemma 4.8 and Lemma 4.9 with this fact and putting $m = O(\log^2 w)$, we get the following.

Theorem 4.10. *There is an $(ndw)^{O(\log \log w)}$ -time hitting-set for n -variate commutative ROABPs with width w and individual degree d .*

Concentration in Set-multilinear Circuits: Similar to Theorem 4.10, it would be interesting to achieve the same time complexity for set-multilinear circuits. Recall from Section 2.2.3 that a polynomial computed by a depth-3 set-multilinear circuit can be written as $(1, 1, \dots, 1) \cdot D$, where $D = D_1(\mathbf{x}_1)D_2(\mathbf{x}_2) \cdots D_q(\mathbf{x}_q)$ is a product of linear polynomials over a commutative algebra. It is easy to see that the same arguments as for commutative ROABP work here. Hence, we get the following result analogous to Lemma 4.8.

Corollary 4.11. *Given n, k , one can compute an n -tuple $\mathbf{f}(t)$ with cost $(nk)^{O(\log \log k)}$ such that for any n -variate polynomial $C(\mathbf{x})$ computed by a depth-3 set-multilinear circuit with top fan-in k , $C(\mathbf{x} + \mathbf{f}(t))$ is $O(\log k)$ -concentrated.*

However, it is not clear whether the second step of the hitting-set construction can be done for set-multilinear circuits, i.e., finding a better hitting-set by assuming that the polynomial is already concentrated (Lemma 4.9).

5 Discussion

For our first result (Theorem 3.6), there are three directions of improvement. Ideally, one would like to have all three at once.

1. Find a similar hitting-set for the unknown-order case. In fact, we conjecture that the same hitting-set (Lemma 3.5) works for the unknown-order case as well.
2. Get a hitting-set for all characteristic fields. It is easy to construct examples over small characteristic fields where our hitting-set does not work.
3. Reduce the time complexity to polynomial time. To achieve this, it seems one has to do away with the divide and conquer approach.

We conjecture a polynomial-time hitting-set for the unknown-order case in Conjecture 3.8.

As mentioned earlier, the ideas here can help in finding a better PRG for ROBPs. In particular, it is a big open question to find an $O(\log n)$ -seed-length PRG for constant-width ROBPs (analogous to Corollary 3.7).

References

- [AB09] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, New York, NY, USA, 1st edition, 2009.
- [AFS⁺15] Matthew Anderson, Michael Forbes, Ramprasad Saptharishi, Amir Shpilka, and Ben Lee Volk. Identity testing and lower bounds for read- k oblivious algebraic branching programs. Technical report, Electronic Colloquium on Computational Complexity (ECCC), 2015.
- [AGKS15] Manindra Agrawal, Rohit Gurjar, Arpita Korwar, and Nitin Saxena. Hitting-sets for ROABP and sum of set-multilinear circuits. *SIAM J. Comput.*, 44(3):669–697, 2015.
- [Agr05] Manindra Agrawal. Proving lower bounds via pseudo-random generators. In *FSTTCS*, volume 3821 of *Lecture Notes in Computer Science*, pages 92–105, 2005.
- [ASS13] Manindra Agrawal, Chandan Saha, and Nitin Saxena. Quasi-polynomial hitting-set for set-depth-formulas. In *STOC*, pages 321–330, 2013.

- [BOC92] Michael Ben-Or and Richard Cleve. Computing algebraic formulas using a constant number of registers. *SIAM J. Comput.*, 21(1):54–58, 1992.
- [DL78] Richard A. Demillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7(4):193 – 195, 1978.
- [dOSV15] Rafael Mendes de Oliveira, Amir Shpilka, and Ben Lee Volk. Subexponential size hitting sets for bounded depth multilinear formulas. In *30th Conference on Computational Complexity, CCC 2015, June 17-19, 2015, Portland, Oregon, USA*, pages 304–322, 2015.
- [FS13] Michael A. Forbes and Amir Shpilka. Quasipolynomial-time identity testing of non-commutative and read-once oblivious algebraic branching programs. In *FOCS*, pages 243–252, 2013.
- [FSS14] Michael A. Forbes, Ramprasad Satharishi, and Amir Shpilka. Hitting sets for multilinear read-once algebraic branching programs, in any order. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 867–875, 2014.
- [GKST15] Rohit Gurjar, Arpita Korwar, Nitin Saxena, and Thomas Thierauf. Deterministic identity testing for sum of read-once oblivious arithmetic branching programs. In *30th Conference on Computational Complexity, CCC 2015, June 17-19, 2015, Portland, Oregon, USA*, pages 323–346, 2015.
- [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. ACM.
- [IMZ12] Russell Impagliazzo, Raghu Meka, and David Zuckerman. Pseudorandomness from shrinkage. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 111–119, 2012.
- [INW94] Russell Impagliazzo, Noam Nisan, and Avi Wigderson. Pseudorandomness for network algorithms. In *Proceedings of the Twenty-sixth Annual ACM Symposium on Theory of Computing, STOC*, pages 356–364, New York, NY, USA, 1994. ACM.
- [KI03] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *STOC*, pages 355–364, 2003.
- [KNS15] Neeraj Kayal, Vineet Nair, and Chandan Saha. Separation between read-once oblivious algebraic branching programs (ROABPs) and multilinear depth three circuits. Technical report, Electronic Colloquium on Computational Complexity (ECCC), 2015.
- [Koi12] Pascal Koiran. Arithmetic circuits: The chasm at depth four gets wider. *Theoretical Computer Science*, 448:56–65, 2012.
- [Nis90] N. Nisan. Pseudorandom generators for space-bounded computations. In *Proceedings of the Twenty-second Annual ACM Symposium on Theory of Computing, STOC '90*, pages 204–212, New York, NY, USA, 1990. ACM.
- [Nis91] Noam Nisan. Lower bounds for non-commutative computation (extended abstract). In *Proceedings of the 23rd ACM Symposium on Theory of Computing, ACM Press*, pages 410–418, 1991.

- [RR99] Ran Raz and Omer Reingold. On recycling the randomness of states in space bounded computation. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing, May 1-4, 1999, Atlanta, Georgia, USA*, pages 159–168, 1999.
- [RS05] Ran Raz and Amir Shpilka. Deterministic polynomial identity testing in non-commutative models. *Computational Complexity*, 14(1):1–19, 2005.
- [Sch80] Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, October 1980.
- [Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *Proceedings of the International Symposium on Symbolic and Algebraic Computation, EUROSAM '79*, pages 216–226, London, UK, UK, 1979. Springer-Verlag.