Electronic Colloquium on Computational Complexity, Report No. 84 (2010)

Deterministic Identity Testing of Read-Once Algebraic Branching Programs

Maurice Jansen* Youming Qiao* Jayalal Sarma M.N.*

Abstract

An algebraic branching program (ABP) is given by a directed acyclic graph with source and sink vertices s and t, respectively, and where edges are labeled by variables or field constants. An ABP computes the sum of weights of all directed paths from s to t, where the weight of a path is taken to be the product of the edge labels on the path. For a *read-once* ABP every variable appears at most once in the graph. More generally, we consider *preprocessed* RO-ABPs (PRO-ABP), which are obtained by allowing univariate polynomials on the edges (at most one non-constant polynomial $T_i(x_i)$ per variable x_i).

We study the problem of polynomial identity testing sums of k many PRO-ABPs (Σ_k-PRO-ABPs). For the main technical part of this paper we develop a recursive property of polynomials in terms of second order partial derivatives and zero substitutions, which we call *alignment*. Using this notion we obtain the following results, in case edges are labeled by univariate polynomials of degree at most d, and provided the underlying field has enough elements (more than $2k^2d^2n^5$ suffices):

1. Given free access to the PRO-ABPs in the sum, we get a deterministic algorithm that runs in time $O(dk^2n^7s^2) + (dn)^{O(k)}$, where s bounds the size of any largest PRO-ABP given on the input. This implies we have a deterministic polynomial time algorithm for testing whether the sum of a constant number of *poly*-degree bounded PRO-ABPs computes the zero polynomial or not.
2. Given black-box access to the PRO-ABPs computing the *individual* polynomials in the sum, we get a deterministic algorithm that runs in time $k^2(dn)^{O(\log n)} + (dn)^{O(k)}$.
3. Given only black-box access to the polynomial computed by the sum of the k PRO-ABPs, we obtain a $(dn)^{O(k+\log n)}$ time deterministic algorithm.

Items 1. and 3. above strengthen two main results of Shpilka and Volkovich [SV09] (Theorems 2 and 3, respectively), who considered polynomial identity testing of sums of k preprocessed *read-once formulas* (Σ_k-PRO-formulas).

1 Introduction

In this paper we study the polynomial identity testing problem (PIT): given an arithmetic circuit C with input variables $x_1, x_2 \ldots x_n$ over a field $\mathbb{F}$, test if $C(x_1, x_2, \ldots, x_n)$ computes the zero polynomial in the ring $\mathbb{F}[x_1, x_2, \ldots x_n]$. This is a well-studied algorithmic problem with a long history

*Institute for Theoretical Computer Science, Tsinghua University, Beijing, P.R. China. Email: `maurice.julien.jansen@gmail.com`, `jimmyqiao86@gmail.com`, `jayalal@tsinghua.edu.cn`. This work was supported in part by the National Natural Science Foundation of China Grant 60553001, and the National Basic Research Program of China Grant 2007CB807900,2007CB807901.

ISSN 1433-8092

and a variety of connections and applications. Efficient randomized algorithms were proposed independently by Schwartz [Sch80] and Zippel [Zip79]. Obtaining a deterministic algorithm for the problem seems surprisingly elusive.

Historically, the connection between derandomizing PIT and proving (algebraic) circuit lower bounds was first noticed in a 1980's paper by Heintz and Schnorr [HS80]. Then after a relatively quiet period, Kabanets and Impagliazzo [KI04] drew renewed attention to this, by showing that giving a deterministic polynomial time (even subexponential time) identity testing algorithm means either that NEXP $\not\subseteq$ P/*poly*, or that the permanent has no polynomial size arithmetic circuits. While advocating this research direction towards lower bounds, Agrawal [Agr05] showed that giving a black-box[1] derandomization of PIT implies the existence of an explicit multilinear polynomial that has no subexponential size arithmetic circuits. Recently, there has been a lot of progress in the area. We refer to a survey by Saxena [Sax09] for an overview.

1.1 Read-Once Formulas and Beyond

Shpilka and Volkovich [SV09, SV08] studied the arithmetic read-once formula model. An *arithmetic read-once formula* (RO-formula) is given by a tree whose nodes are taken from the set $\{+, \times\}$, and whose leaves are variables or field constants, subject to the restriction that each variable x_i appears at most once. More generally, *preprocessed* RO-formulas (PRO-formulas) are obtained by allowing univariate polynomials as edge labels instead of variables (at most one non-constant polynomial $T_i(x_i)$ per variable x_i). For "moderate" k, efficient deterministic PIT algorithms were given in [SV09, SV08] for sums of k many PRO-formulas (Σ_k-PRO-formulas).

Given the status of PIT, it is important to enlarge upon any known techniques that solve special cases of the problem (like those in [SV09, SV08]) for as much as possible, even if only to establish more clearly the cases of the problem where apparently radically new techniques are required. In this paper, we consider the generalization of RO-formulas to *global read-once algebraic branching programs*[2].

An algebraic branching program (ABP) is a layered directed acyclic graph with two special vertices s and t. Each edge is assigned a weight, which is an element of $X \cup \mathbb{F}$, where X is a set of variables. For a path in the graph its weight is taken to be the product of the weight on its edges. The output of the ABP is defined to be the sum of weights of all paths from s to t. The ABP is said to be *global read-once* if each variable appears on at most one edge. For simplicity we drop the adjective "global" for the rest of this paper, and merely talk about read-once ABPs (RO-ABPs). We call a polynomial $f \in \mathbb{F}[X]$ a *RO-ABP-polynomial*, if there exists a RO-ABP which computes f. Similarly as with RO-formulas, we also consider the generalization to preprocessed RO-ABPs (PRO-ABPs), by allowing edge labels that are univariate polynomials (at most one non-constant polynomial $T_i(x_i)$ per variable x_i).

We note that RO-ABPs are a natural generalization to consider. Applying a construction by Valiant [Val79], if f can be computed by a RO-formula of size s, then f can be computed by a RO-ABP of size $O(s)$. Non-black box identity testing a single RO-ABP is easily solved by phrasing it as a reachability problem (See Section 3.1). In this case it is more interesting to consider black-box PIT. PIT of RO-ABPs can be seen to be a special case of the more general problem of black-box identity testing "read-once determinantal expressions", i.e. expressions of the form

[1] In the black-box model one can only query an oracle holding the circuit C for the output of C on a given input.
[2] See Section 2 for a formal definition.

$\det M(x_1, x_2, \ldots, x_n)$, where each variable x_i appears at most once in the matrix M. It is well-known the bipartite perfect matching problem (BIPARTITE-PM) reduces to identity testing such expressions. By giving a black-box PIT algorithm for this, Agrawal, Hoang and Thierauf [AHT07] put BIPARTITE-PM for graphs with polynomially bounded number of perfect matchings inside NC^2. They conjectured this approach to work for the general problem. Perhaps further progress on black-box PIT of RO-ABPs can be a first step towards the more general case of black-box testing read-once determinantal expressions.

We remark that RO-ABPs are strictly more powerful than RO-formulas. Appendix C shows a RO-ABP computing $g = x_1x_2 + x_2x_3 + \cdots + x_{2n-1}x_{2n}$. Example 3.12 in [SV08] shows that g can not be computed by a RO-formula, if $n \geq 2$. We note that, like the RO-formula, the RO-ABP model is still not universal, e.g. for $n \geq 3$, $\sum_{1 \leq i < j \leq n} x_i x_j$, is not an RO-ABP-polynomial (See Appendix A).

For black-box identity testing a single RO-formula, the following construction is given in [SV09]: Let $A = \{a_1, a_2, \ldots, a_n\} \subseteq \mathbb{F}$ be a set of size n. For every $i \in [n]$, let $u_i(w)$ be the ith Lagrange interpolation polynomial on A. Then $u_i(w)$ is a polynomial of degree $n-1$ satisfying that $u_i(a_j) = 1$ if $j = i$ and 0 otherwise. For every $i \in [n]$ and $k \geq 1$, define $G_k^i(y_1, y_2, \ldots, y_k, z_1, z_2, \ldots, z_k) = \sum_{j \in [k]} u_i(y_j) z_j$. and let $G_k(y_1, y_2, \ldots, y_k, z_1, z_2, \ldots, z_k) : \mathbb{F}^{2k} \to \mathbb{F}^n$, be defined by $G_k = (G_k^1, G_k^2, \ldots, G_k^n)$. We refer to the polynomial mapping G_k as the kth-order SV-generator, or SV-generator for short. Given its track record[3], it is important to investigate how far this construction will take us towards our ultimate PIT goals. We demonstrate it takes us further than was known previously, by showing it provides a black-box test for PRO-ABP-polynomials. Namely, we have the following lemma. (For a proof see Section 3.2):

Lemma 1. *Let $d > 0$ be an integer, and assume that $|\mathbb{F}| > d$. If $f \in \mathbb{F}[X]$ is a nonzero polynomial with $|Var(f)| \leq 2^m$, for some $m \geq 0$, that is computable by a PRO-ABP that has univariate polynomials with degrees bounded by d, then $f(G_{m+1}) \not\equiv 0$.*

The above lemma implies that we have an *explicit hitting set* S of size $(nd)^{O(\log n)}$, such that any nonzero PRO-ABP-polynomial in n variables with individual degrees bounded by d evaluates to a nonzero value for at least one element of S.

1.2 Main Results

To make further progress, we consider sums of k many PRO-ABPs. In this case we manage to give an explicit hitting-set of size $(dn)^{O(k+\log n)}$, resulting in the following theorem:

Theorem 1. *Let $f = \sum_{i \in [k]} f_i$ be a sum of k PRO-ABP-polynomials in n variables with individual degrees at most d. Let $\mathbb{F}$ be a field with $|\mathbb{F}| > k^2 d^2 n^5 + kdn^4$. Given black-box access to f, it can be decided deterministically in time $(dn)^{O(k+\log n)}$ whether $f \equiv 0$.*

This strengthens a main result of [SV09], namely Theorem 3, which provides a deterministic $(dn)^{O(k+\log n)}$ time black-box PIT algorithm for Σ_k-PRO-formulas. By previous remarks, any Σ_k-PRO-formula computable polynomial is Σ_k-PRO-ABP computable, with negligible blow-up in size.

[3]Very recently, this generator has also been applied to identity testing multilinear depth 4 circuits with bounded top fan-in [KMSV09].

Moreover, we sketch[4] in Appendix B a dimension argument that shows there exists a RO-ABP-polynomial in n variables that requires $k = \Omega(n)$ in the Σ_k-PRO-formula model. Hence we have a strict separation for $k = o(n)$.

In the non-black-box setting we will prove the following result:

Theorem 2. *Let $\{A_i\}_{i\in[k]}$ be a set of k PRO-ABPs in n variables with individual degrees bounded by d. Let $\mathbb{F}$ be a field with $|\mathbb{F}| > dkn^2$. Given $\{A_i\}_{i\in[k]}$ on the input, it can be decided deterministically in time $O(dk^2n^7s^2) + (dn)^{O(k)}$ whether $\sum_{i\in[k]} f_i \equiv 0$, where f_i is the PRO-ABP-polynomial computed by A_i, for $i \in [k]$.*

Since the construction in [Val79] can be computed efficiently, this strengthens Theorem 2 in [SV09]. Finally, if black-box access is granted to the individual f_i's, which we call the *semi-black-box* setting, we obtain the following result:

Theorem 3. *Let $\{f_i\}_{i\in[k]}$ be a set of k PRO-ABP-polynomials in n variables with individual degrees bounded by d. Let $\mathbb{F}$ be a field with $|\mathbb{F}| > dkn^2$. Given black-box access to each individual f_i, it can decided deterministically in time $k^2(dn)^{O(\log n)} + (dn)^{O(k)}$ whether $\sum_{i\in[k]} f_i \equiv 0$.*

1.3 Techniques

The results for Σ_k-RO-ABP and Σ_k-PRO-ABP PIT are obtained through the *hardness of representation* approach of [SV09, SV08]. There the PIT algorithms are derived from a statement that $x_1x_2\dots x_n$ cannot be expressed as a sum of $k \le n/3$ RO-formula computable polynomials $\{f_i\}_{i\in[k]}$, if the polynomials f_i satisfy some special property. We do not need to define this special property for the discussion here, except that we should name it: $\bar{0}$-justification.

Unfortunately, the property of $\bar{0}$-justification, does not work for the Σ_k-RO-ABP model. With some thought it can be seen that the monomial $x_1x_2\dots x_n$ is expressible as the sum of three $\bar{0}$-justified RO-ABP-polynomials. Our main technical contribution is the development of a new "special property", called *alignment*. For this property we show a hardness of representation theorem can still be proved. Moreover, we show it can be enforced simultaneously for a collection of PRO-ABP-polynomials by means of an efficiently computable coordinate shift.

Regarding the latter, consider $f = f_1 + f_2 + \ldots + f_k$, where each f_i is a PRO-ABP-polynomial. Then $\forall v \in \mathbb{F}^n$, $f \equiv 0 \iff f(x_1+v_1, x_2+v_2, \ldots, x_n+v_n) \equiv 0$. With some technical work, we will establish a *sufficient* condition for alignment. With it we show that we can compute a coordinate shift v such that all $f_i(x+v)$ are aligned. Such a shift v is called a *simultaneous alignment*. In the case of having only black-box access to f, we will show we have a "small" set of candidates containing at least one simultaneous alignment. The PIT algorithms will follow from this.

The rest of this paper is organized as follows. Section 2 contains preliminaries and Section 3 presents the identity testing algorithms for a single PRO-ABP in the black-box and non-black-box setting. In Section 4 we develop the tools regarding alignment. Section 5 contains the hardness of representation theorems for RO-ABPs and PRO-ABPs. Then in Section 6 we show how to compute a simultaneous alignment. From these developments we put the PIT algorithms together in Section 7.

[4] For simplicity, the sketch is given for algebraically closed $\mathbb{F}$, but can be adapted to other infinite fields of interest like $\mathbb{Q}$ and $\mathbb{R}$.

2 Preliminaries

Let $X = \{x_1, x_2, \ldots, x_n\}$ be a set of variables and let $\mathbb{F}$ be a field. For $\mathbb{F}$ with more than d elements, we define $\mathcal{W}^n_{k,d} = \{y \in S^n \mid wt(y) \leq k\}$, where $wt(y)$ counts the number of nonzeros in y, and where we have fixed some arbitrary subset S of size $d+1$ of $\mathbb{F}$ that contains 0.

An algebraic branching program (ABP) is a 4-tuple $A = (G, w, s, t)$, where $G = (V, E)$ is an edge-labeled directed acyclic graph for which the vertex set V can be partitioned into levels $L_0, L_1, \ldots, L_d$, where $L_0 = s$ and $L_d = t$. Vertices s and t are called the source and sink of B, respectively. Edges may only go between consecutive levels L_i and L_{i+1}. The label function $w : E \to X \cup \mathbb{F}$ assigns variables or field constants to the edges of G. For a path p in G, we extend the weight function by $w(p) = \prod_{e \in p} w(e)$. Let $P_{i,j}$ denote the collection of all directed paths p from i to j in G. The program A computes the polynomial $\hat{A} := \sum_{p \in P_{s,t}} w(p)$. The size of A is defined to be $|V|$. An ABP is said to be *global read-once* if $|w^{-1}(x_i)| \leq 1$, for each $x_i \in X$. That is, every variable is read at most once by the program. For simplicity, we will drop the adjective "global" in the rest of this paper, and we will merely speak about read-once ABPs (RO-ABPs). A polynomial $f \in \mathbb{F}[X]$ is called a *RO-ABP-polynomial*, if there exists a RO-ABP that computes f. We use the following notation: for x_i present on arc (v, w) in a RO-ABP A: $begin(x_i) = v$ and $end(x_i) = w$. We let $source(A)$ and $sink(A)$ stand for the source and sink of A. For any nodes v, w in A, we denote the subprogram with source v and sink w by $A_{v,w}$. A *layer* of a RO-ABP A is any subgraph induced by two consecutive levels L_i and L_{i+1} in A. We will assume RO-ABPs are in the form given by the following straightforwardly proven lemma:

Lemma 2. *If $f \in \mathbb{F}[X]$ is a RO-ABP-polynomial, then f can be computed by a RO-ABP A, where every layer contains at most one variable-labeled edge.*

For any fixed integer parameter d, we generalize the RO-ABP model to *preprocessed* read-once ABPs (PRO-ABP) by allowing univariate poynomials as edges labels. Let $\mathcal{T}_d$ be the set of monic univariate polynomials T of degree at most d with $T(0) = 0$. Let $Z = \{z_i : i \in [n]\}$ be a set of indeterminates. A *preprocessed RO-ABP-polynomial (PRO-ABP-polynomial)* is any polynomial $f \in \mathbb{F}[X]$ that can be written as $f = g(T_1(x_1), T_2(x_2), \ldots, T_n(x_n))$, where $g \in \mathbb{F}[Z]$ is a RO-ABP-polynomial, and each $T_i \in \mathcal{T}_d$. In this case $(g, \{T_i(x_i)\}_{i \in [n]})$ is called the d-decomposition of f, and we say f is a d-decomposable PRO-ABP-polynomial. Note that both the classes of RO-ABP-polynomials and PRO-ABP-polynomials are closed[5] under coordinate shifting, i.e. for any f, $f(x_1 + v_1, \ldots, x_n + v_n)$ stays within the class, for all $v \in \mathbb{F}^n$. The proof of the following proposition is left as an exercise to the reader.

Proposition 1. *If for $g, h \in \mathbb{F}[Z]$ we have that $(g, \{T_i(x_i)\}_{i \in [n]})$ and $(h, \{U_i(x_i)\}_{i \in [n]})$ are d-decompositions of a PRO-ABP-polynomial $f \in \mathbb{F}[X]$, then 1) $g = h$, and 2) $\forall x_i \in Var(f), T_i = U_i$.*

Let f be a polynomial in the ring $\mathbb{F}[X]$. For $\alpha \in \mathbb{F}$, $f|_{x_i = \alpha}$ denotes the polynomial $f(x_1, x_2, \ldots x_{i-1}, \alpha, x_{i+1}, \ldots, x_n)$. Extending this to sets of variables, for a subset $I \subseteq [n]$ and an assignment $a \in \mathbb{F}^n$, $f|_{x_I = a_I}$ is the the polynomial resulting from setting the variable x_i to a_i in f for every $i \in I$. The following two notions are taken from [SV09]. We say that a polynomial f *depends on a variable* x_i if there exists an $a \in \mathbb{F}^n$ and $b \in \mathbb{F}$, such that

[5] Related to this, note that taking $\mathcal{T}_d$ to be the set of arbitrary univariate polnomials of degree at most d instead, would not change the class of PRO-ABP-polynomials, so wlog. we restrict ourselves to monic univariate polynomials having constant term zero.

$f(a_1, a_2, a_{i-1}, a_i, a_{i+1}, \ldots, a_n) \neq f(a_1, a_2, a_{i-1}, b, a_{i+1}, \ldots, a_n)$. The set of variables x_i that f depends on is denoted by $Var(f)$.

For $\alpha \in \mathbb{F}$ and $f \in \mathbb{F}[X]$, the *partial derivative with respect to x_i and direction α*, denoted by $\frac{\partial f}{\partial_\alpha x_i}$, is defined as $f|_{x_i=\alpha} - f|_{x_i=0}$. By convention, if we do not mention the direction, it means $\alpha = 1$. For a set of variables J, $\partial_J f$ denotes taking partial w.r.t. all variables in J (and direction $\alpha = 1$). Setting values to variables commutes with taking partial derivatives in the following way: $\forall i \neq j$, $\frac{\partial f}{\partial_\alpha x_i}|_{x_j=a} = \frac{\partial (f|_{x_j=a})}{\partial_\alpha x_i}$. We will freely use the properties listed for this notion in [SV09].

Proposition 2. *Suppose $f \in \mathbb{F}[X]$ has individual degrees bounded by r. For any $S \subseteq \mathbb{F}$ with $|S| > r$, we have that f depends on $x_i \Leftrightarrow \exists \alpha \in S, \frac{\partial f}{\partial_\alpha x_i} \not\equiv 0$.*

The above proposition follows from the "Combinatorial Nullstellensatz":

Lemma 3 (Lemma 2.1 in [Alo99])**.** *Let $f \in \mathbb{F}[X]$ be a nonzero polynomial such that the degree of f in x_i is bounded by r_i, and let $S_i \subseteq \mathbb{F}$ be of size at least $r_i + 1$, for all $i \in [n]$. Then there exists $(s_1, s_2, \ldots, s_n) \in S_1 \times S_2 \times \ldots \times S_n$ with $f(s_1, s_2, \ldots, s_n) \neq 0$.*

Now we prove the following lemma.

Lemma 4. *Let $\alpha \in \mathbb{F}\backslash\{0\}$ be given. We have 1) If $f \in \mathbb{F}[X]$ is a RO-ABP-polynomial, then $\frac{\partial f}{\partial_\alpha x_i}$ is a RO-ABP-polynomial, and 2) Suppose $f \in \mathbb{F}[X]$ is a PRO-ABP-polynomial with d-decomposition $(g, \{T_i(x_i)\}_{i\in[n]})$. Then $\frac{\partial f}{\partial_\alpha x_i}$ is a PRO-ABP-polynomial with d-decomposition $(T_i(\alpha)\frac{\partial g}{\partial z_i}, \{T_i(x_i)\}_{i\in[n]})$.*

We start by proving the first item. Let $p = |Var(f)|$. In case $p = 0$ it is trivial. Assume $p > 0$. If $x_i \notin Var(f)$, then $\frac{\partial f}{\partial_\alpha x_i} \equiv 0$, in which case the property trivially holds. Now suppose $x_i \in Var(f)$. Hence x_i must appear somewhere in A. Say x_i is on the arc (v_1, w_1) from level L_j to L_{j+1}, where $L_j = \{v_1, v_2, \ldots, v_{m_1}\}$ and $L_{j+1} = \{w_1, w_2, \ldots, w_{m_2}\}$, for certain j, m_1, m_2. We can write

$$f = \sum_{a\in[m_1]} \sum_{b\in[m_2]} f_{s,v_a} w(v_a, w_b) f_{w_b,t}, \tag{1}$$

where for any nodes p and q in A, $f_{p,q}$ is the polynomial computed by subprogram $A_{p,q}$. Then

$$\begin{aligned}
\frac{\partial f}{\partial_\alpha x_i} &= f_{|x_i=\alpha} - f_{|x_i=0} \\
&= \sum_{a\in[m_1]} \sum_{b\in[m_2]} f_{s,v_a} w(v_a, w_b)_{|x_i=\alpha} f_{w_b,t} - \sum_{a\in[m_1]} \sum_{b\in[m_2]} f_{s,v_a} w(v_a, w_b)_{|x_i=0} f_{w_b,t} \\
&= \sum_{a\in[m_1]} \sum_{b\in[m_2]} f_{s,v_a} \left(w(v_a, w_b)_{|x_i=\alpha} - w(v_a, w_b)_{|x_i=0} \right) f_{w_b,t} \\
&= \alpha f_{s,v_1} f_{w_1,t}.
\end{aligned}$$

Hence we obtain a valid RO-ABP computing $\frac{\partial f}{\partial_\alpha x_i}$ from A by setting the label of the wire (v_1, w_1) to α, and removing all other wires between layers L_j and L_{j+1}.

The second item follows easily by writing $g = z_i \frac{\partial g}{\partial z_i} + g_{|z_i=0}$. $\square$

3 Identity Testing a Single PRO-ABP

In this section we describe identity testing algorithms for a single PRO-ABP, in the Non-Black-Box and Black-Box setting.

3.1 Non-Black-Box Testing a Single RO-ABP

Consider a RO-ABP A. Denote the source and sink of A by s and t, respectively. Suppose that x_i labels the edge (s_i, t_i). Wlog. assume that the order of variable layers in A is $x_1, x_2, \ldots, x_n$. We have the following easy proposition:

Proposition 3. *Suppose* $1 \leq i_1 < i_2 < \cdots < i_k \leq n$. *For a RO-ABP A, $x_{i_1} x_{i_2} \ldots x_{i_k}$ appears in $\hat{A}$ if and only if the constant terms in $\hat{A}(s, s_{i_1})$, $\hat{A}(t_{i_m}, s_{i_{m+1}})$, for all $m \in [k-1]$, and $\hat{A}(t_k, t)$ are not zero.*

We build a directed graph $G_A = (V, E)$ for RO-ABP A with vertex set $V = \{s, t, x_1, x_2, \ldots, x_n\}$. Edges are given as follows:

1. (s, x_i), if the constant term in $\hat{A}(s, s_i)$ is nonzero.
2. (x_i, t), if the constant term in $\hat{A}(t_i, t)$ is nonzero.
3. (x_i, x_j), $i < j$, if the constant term in $\hat{A}(t_i, s_j)$ is nonzero.

We have the following corollary of Proposition 3:

Corollary 1. $\hat{A}(x_1, \ldots, x_n) \equiv 0$ *if and only if t is not reachable form s in G_A.*

The algorithm for testing A is to construct G_A and to test connectivity. This can be done in time $O(n^2 s^2)$, where s bounds the size of A.

3.2 Black-Box Testing a Single PRO-ABP

In this subsection, we give a proof of Lemma 1 which demonstrates that the generator described in [SV09] provides a black-box test for PRO-ABP polynomials. We restate the lemma first.

Lemma 5. *Let $d > 0$ be an integer, and assume that $|\mathbb{F}| > d$. If $f \in \mathbb{F}[X]$ is a nonzero polynomial with $|Var(f)| \leq 2^m$, for some $m \geq 0$, that is computable by a PRO-ABP that has univariate polynomials with degrees bounded by d, then $f(G_{m+1}) \not\equiv 0$.*

Proof. Let $p = |Var(f)|$. The proof proceeds by induction on p. The bases $p = 0$ and $p = 1$ trivially hold. Suppose $p > 1$. Hence $m \geq 1$. Consider arbitrary PRO-ABP A computing f. Wlog. assume that any nonconstant edge label in A is given by a monic univariate polynomial $T_i(x_i)$ with $T_i(0) = 0$. Let s and t be the source and sink of A, respectively. Wlog. assume that only the p variables in $Var(f)$ are present in A, and assume A satisfies the condition yielded by Lemma 2. Observe that for some variable x_i there are at most $p/2$ variables in layers before the layer containing x_i, and at most $p/2$ variables in layers after. (If p is odd it splits $((p-1)/2), (p-1)/2)$ if p is even it splits $(p/2 - 1, p/2)$).

Say we have univariate polynomial $T_i(x_i)$ on the arc (v_1, w_1) from layer L_j to L_{j+1}, where $L_j = \{v_1, v_2, \ldots, v_{m_1}\}$ and $L_j = \{w_1, w_2, \ldots, w_{m_2}\}$, for certain j, m_1, m_2. We can write

$$f = \sum_{a=1}^{m_1} \sum_{b=1}^{m_2} f_{s,v_a} f_{w_b,t} w(v_a, v_b), \tag{2}$$

where for any nodes p and q in A, $f_{p,q}$ is the polynomial computed by subprogram of $A_{p,q}$. Consider $f' = f(G_m^1, \ldots, G_m^{i-1}, x_i, G_m^{i+1}, \ldots, G_m^n)$.

Claim 1. *f' depends on x_i.*

Proof. Since f depends on x_i, by Proposition 2, there exists nonzero $\alpha \in \mathbb{F}$ such that $f'' := \frac{\partial f}{\partial_\alpha x_i} \not\equiv 0$. By Proposition 2, it suffices to show that $\frac{\partial f'}{\partial_\alpha x_i} \not\equiv 0$. Note that $\frac{\partial f'}{\partial_\alpha x_i} = f''(G_m)$. We have that $\frac{\partial f}{\partial_\alpha x_i} = T_i(\alpha) f_{s,v_1} f_{w_1,t}$. Note that $|Var(f_{s,v_1})|$ and $|Var(f_{w_1,t})|$ are both at most $p/2$. Since $f'' \not\equiv 0$, both f_{s,v_1} and $f_{w_1,t}$ are not identically zero and $T_i(\alpha) \neq 0$. As $p/2 < p$, the induction hypothesis applies. Since $p/2 \leq 2^{m-1}$, it yields that $f_{s,v_1}(G_m) \not\equiv 0$ and $f_{w_1,t}(G_m) \not\equiv 0$. Therefore $f''(G_m) \not\equiv 0$. This proves the claim. □

Recall the set $A = \{a_1, \ldots, a_n\}$ used for the construction of the SV-generator. By Observation 5.2 in [SV09], $f(G_{m+1})_{|y_{m+1}=a_i} = f'_{|x_i = G_m^i + z_{m+1}}$. Since z_{m+1} does not appear in G_m^j for any j, we get by Claim 1 that $f(G_{m+1})_{|y_{m+1}=a_i} \not\equiv 0$. Hence $f(G_{m+1}) \not\equiv 0$. □

4 X-Aligned RO-ABP and PRO-ABP polynomials

A first requirement of our new "special property" is that we can bring out linear factors somehow. The following lemma shows that partial derivatives can be used for this.

Lemma 6. *Let $f \in \mathbb{F}[X]$ be a RO-ABP-polynomial with $|Var(f)| \geq 3$. Then for any $x_i \in Var(f)$, there exist distinct $x_j, x_k \in X \backslash \{x_i\}$ such that $\frac{\partial^2 f}{\partial x_j \partial x_k} = g \cdot (\beta x_i - \alpha)$, where g is a RO-ABP-polynomial that does not depend on x_i, and $\alpha, \beta \in \mathbb{F}$.*

Proof. Let A be a RO-ABP computing f. Wlog. assume all variables in X appear in A. By Lemma 2 assume wlog. that A has at most one variable per layer. Let $x_{r_1}, x_{r_2}, \ldots, x_{r_n}$ be the variables in X as they appear layer-by-layer, when going from the source to the sink of A. Consider an arbitrary $x_i \in Var(f)$. First, we handle the case that $i = r_m$, for some $1 < m < n$.

Let $j = r_{m-1}$ and $k = r_{m+1}$. So x_j and x_k are the variables right before and right after x_i in A, respectively. Assume that x_j and x_k label the edges (u, v) and (m, n) respectively. Then $\frac{\partial^2 f}{\partial x_j \partial x_k} = f_{s,u} f_{v,m} f_{n,t}$, where $f_{s,u} f_{v,m}$, and $f_{n,t}$ are computed by the subprograms $A_{s,u}, A_{v,m}$, and $A_{n,t}$, respectively. Observe that $f_{v,m}$ is of form $\beta x_i - \alpha$, for $\alpha, \beta \in \mathbb{F}$. Take $g = f_{s,u} f_{v,m}$, which is easily seen to be RO-ABP-computable by putting $A_{s,u}$ and $A_{v,m}$ in series.

The special case where $i = r_1$ ($i = r_n$), i.e. x_i is the first (last) variable in A, is handled similarly as above, by choosing $x_k \in X \backslash \{x_i, x_j\}$ arbitrarily and appealing to Lemma 4. □

Recall that one of our goals is to show that small sums of RO-ABP-polynomials satisfying the "special property" cannot represent $P_n := x_1 x_2 \ldots x_n$. In the above, if $\beta \neq 0$, setting $x_i = \alpha/\beta$, kills $\frac{\partial^2 f}{\partial x_j \partial x_k}$. As will become clear in the Hardness of Representation Theorem 4, it is extra nice to do so with $\alpha/\beta \neq 0$, since P_n stays self-similar under such a substitution. Note it also does when taking $\partial x_j \partial x_k$. We encapsulate this as follows:

Definition 1. *Let $S \subseteq X$. Every RO-ABP-polynomial $f \in \mathbb{F}[X]$ with $|Var(f)| \leq 2$ is X-pre-aligned on S. A RO-ABP-polynomial $f \in \mathbb{F}[X]$ with $|Var(f)| > 2$ is X-pre-aligned on S, if the following condition is satisfied: for every $x_i \in S$, there exist distinct $x_j, x_k \in X\backslash\{x_i\}$ such that $\frac{\partial^2 f}{\partial x_j \partial x_k} = g \cdot (\beta x_i - \alpha)$, where g is a RO-ABP-polynomial that does not depend on x_i, and $\alpha, \beta \in F$ satisfy that $\alpha = 0 \Rightarrow \beta = 0$.*

If f is X-pre-aligned on $Var(f)$, we simply say that f is X-pre-aligned. Another requirement, stemming from the "Vanishing Theorem" to be proved later, is that the "special property" holds recursively w.r.t. zero substitution. For technical reasons, we need to keep a separation between concepts (Think of the following as "special property++"). We make an inductive definition.

Definition 2. *Every RO-ABP-polynomial $f \in \mathbb{F}[X]$ with $|Var(f)| \leq 2$ is X-aligned. A RO-ABP-polynomial $f \in \mathbb{F}[X]$ with $|Var(f)| > 2$ is X-aligned, if the following conditions are satisfied: 1) f is X-pre-aligned, and 2) for every $x_i \in Var(f)$, $f_{|x_i=0}$ is $X\backslash\{x_i\}$-aligned.*

For a PRO-ABP-polynomial f with d-decomposition $(g, \{T_i(x_i)\}_{i\in[n]})$, where $g \in \mathbb{F}[Z]$, we say it is X-pre-aligned on S if g is Z-pre-aligned on S', where $S' = \{z_i \in Z : x_i \in S\}$. Note that due to Proposition 1, this is well-defined. Similarly, we say that f is X-aligned provided g is Z-aligned.

Next we prove our notions are well-behaved, which show yet more constraints that needed to be satisfied for the proof to go through. Mostly, it will be sufficient to establish these results for the unpreprocessed case only. Then finally, we also need to show that X-alignment can be enforced by coordinate shifting (simultaneously for several PRO-ABP-polynomials). This is left for Section 6.

Proposition 4. *If RO-ABP-polynomial $f \in \mathbb{F}[X]$ is X-pre-aligned, then $\forall \mu \in \mathbb{F}$, $\mu \cdot f$ is X-pre-aligned. The same statement holds with aligned instead of pre-aligned.*

One main requirement is that alignment is preserved when taking partial derivatives which is given by the following lemma.

Lemma 7. *For any RO-ABP-polynomial $f \in \mathbb{F}[X]$ and any $x_r \in X$, the following hold: 1) If f is X-pre-aligned, then $\frac{\partial f}{\partial x_r}$ is $(X\backslash\{x_r\})$-pre-aligned. 2) If f is X-aligned, then $\frac{\partial f}{\partial x_r}$ is $(X\backslash\{x_r\})$-aligned.*

Proof. We first show that Item 1 holds. Let $f' = \frac{\partial f}{\partial x_r}$ and $X' = X\backslash\{x_r\}$. By Lemma 4, we know that f' is a RO-ABP-polynomial. Assume that $|Var(f')| \geq 3$, since otherwise the statement holds trivially. Consider arbitrary $x_i \in Var(f')$. Then $x_i \in Var(f)$, so there exist distinct x_j and x_k in $X\backslash\{x_i\}$, such that $\frac{\partial^2 f}{\partial x_j \partial x_k} = g \cdot (\beta x_i - \alpha)$, where g is a RO-ABP-polynomial that does not depend on x_i, and $\alpha = 0 \Rightarrow \beta = 0$. Consider the following two cases:

Case I: "$r \notin \{j,k\}$". Hence $x_j, x_k \in X'\backslash\{x_i\}$. We have that $\frac{\partial^2 f'}{\partial x_j \partial x_k} = \frac{\partial^3 f}{\partial x_j \partial x_k \partial x_r} = \frac{\partial g}{\partial x_r} \cdot (\beta x_i - \alpha)$. By Lemma 4, $\frac{\partial g}{\partial x_r}$ is a RO-ABP-polynomial, and it clearly does not depend on x_i, so we conclude that f' is X'-pre-aligned on $\{x_i\}$.

Case II: "$r \in \{j,k\}$". Wlog. assume $r = j$. Then $x_k \in X'\backslash\{x_i\}$. Since $|Var(f')| \geq 3$, there must be at least one more variable x_l in $Var(f')$ distinct from each of x_k and x_i. Then $x_l \in X'\backslash\{x_i\}$. We have that $\frac{\partial f'}{\partial x_k} = g \cdot (\beta x_i - \alpha)$. Hence $\frac{\partial^2 f'}{\partial x_k \partial x_l} = \frac{\partial g}{\partial x_l} \cdot (\beta x_i - \alpha)$. We again conclude f' is X'-pre-aligned on $\{x_i\}$.

Item 2 is proved by induction on $|X|$. The base case is when $|X| \leq 3$. Then $|Var(f')| \leq 2$, and hence f' is X'-aligned. Now suppose $|X| > 3$. Assume $|Var(f')| > 2$, since otherwise it is trivial. By Item 1, we know f' is X'-pre-aligned. Consider an arbitrary $x_i \in Var(f')$. Then $x_i \in Var(f)$. We have that $f'_{|x_i=0} = \left(\frac{\partial f}{\partial x_r}\right)_{x_i=0} = \frac{\partial f_{|x_i=0}}{\partial x_r}$. Since $f_{|x_i=0}$ is $(X\backslash\{x_i\})$-aligned, we can apply the induction hypothesis to conclude that $\frac{\partial f_{|x_i=0}}{\partial x_r}$ is $(X\backslash\{x_i\})\backslash\{x_r\} = (X'\backslash\{x_i\})$-aligned. □

In addition to the above, we crucially need the following "Nearly Unique Nonalignment Lemma".

Lemma 8. *Let $f \in \mathbb{F}[X]$ be an X-pre-aligned RO-ABP-polynomial for which $\frac{\partial^2 f}{\partial x_p \partial x_q} \not\equiv 0$, for any distinct $x_p, x_q \in X$. Then there are at most two $\gamma \in \mathbb{F}$ such that $f_{|x_n=\gamma}$ is not $(X\backslash\{x_n\})$-pre-aligned.*

Before giving the proof, we need a lemma.

Lemma 9. *Let $f \in \mathbb{F}[X]$ be a RO-ABP-polynomial with $|Var(f)| \geq 3$ that is X-pre-aligned on S, for some $S \subseteq Var(f)$. Assume that for any distinct $x_p, x_q \in X$, $\frac{\partial^2 f}{\partial x_p \partial x_q} \not\equiv 0$. In any RO-ABP A computing f, for any $x_i \in S$,*

1. *if there exists a non-constant layer with variable x_a right before the x_i-layer, and there exists a non-constant layer with variable x_b right after the x_i-layer, then*

$$\frac{\partial^2 f}{\partial x_a \partial x_b} = g \cdot (\beta x_i - \alpha),$$

where g is a RO-ABP-polynomial that does not depend on x_i, and $\alpha, \beta \in F$ satisfy that $\alpha = 0 \Rightarrow \beta = 0$. Furthermore, $-\alpha$ equals the sum of weights of all paths from end(x_a) to begin(x_b) that do not go over x_i.

Proof. Consider $x_i \in S$. Since f is X-pre-aligned on S, we know there exist distinct $x_j, x_k \in X\backslash\{x_i\}$ with $\frac{\partial^2 f}{\partial x_j \partial x_k} = h \cdot (\beta' x_i - \alpha')$, where h is a RO-ABP-polynomial that does not depend on x_i, and $\alpha', \beta' \in F$ satisfy that $\alpha' = 0 \Rightarrow \beta' = 0$. Since $\frac{\partial^2 f}{\partial x_j \partial x_k} \not\equiv 0$, it must be that $\alpha' \neq 0$.

Case I: In A, the x_i-layer lies in between the x_j-layer and x_k layer.

Wlog assume the x_i layer lies before the x_k-layer and after the x_j-layer (according to the order of the DAG underlying A). Write $\frac{\partial^2 f}{\partial x_j \partial x_k} = p_1 p_2 \cdot (q_1 q_2 x_i + q_3)$, where

- p_1 is the sum of weights over all paths in A from *source*(A) to *begin*(x_j), and p_2 is the sum of weights over all paths in A from *end*(x_k) to *sink*(A).
- q_3 is the sum of weights over all paths from *end*(x_j) to *begin*(x_k) that bypass the x_i-edge, q_1 is the sum of weights over all paths from *end*(x_j) to *begin*(x_i), and q_2 is the sum of weights over all paths from *end*(x_i) to *begin*(x_k).

Now we have that $p_1 p_2 \cdot (q_1 q_2 x_i + q_3) = h \cdot (\beta' x_i - \alpha')$. Since both $p_1 p_2$ and h do not depend on x_i, it must be that $(\beta' x_i - \alpha') \mid (q_1 q_2 x_i + q_3)$. Note that β' cannot equal 0, since then one of q_1, q_2 would be zero. The latter implies that $\frac{\partial^2 f}{\partial x_i \partial x_j} \equiv 0$ or $\frac{\partial^2 f}{\partial x_i \partial x_k} \equiv 0$, which is a contradiction. Since $\beta' \neq 0$, we can conclude that $q_3 = \mu q_1 q_2$ for some $\mu \in \mathbb{F}$, $\mu \neq 0$. Now we need the following claim:

Claim 2. *Given an RO-ABP A computing $f(x_1, \ldots, x_n)$, if for any distinct $x_p, x_q \in X$, $\frac{\partial^2 f}{\partial x_p \partial x_q} \not\equiv 0$, then $\prod_{i\in[n]} x_i$ appears in f. Furthermore, for two variables x_i and x_j, if x_i is before x_j in A, if we let S be the set of variables in between x_i and x_j, then $\prod_{x_m \in S} x_m$ is a term in the polynomial $\hat{A}(end(x_i), begin(x_j))$.*

Proof. Suppose the variable layers in A are arranged according to the permutation $\phi : [n] \to [n]$, that is, $x_{\phi(i)}$ labels the ith variable layer. Then we that

1. $\hat{A}(s, begin(x_{\phi(1)})) \not\equiv 0$ (Since otherwise $\frac{\partial^2 f}{\partial x_{\phi(1)} \partial x_{\phi(2)}} \equiv 0$),
2. Similarly $\hat{A}(end(x_{\phi(n)}), t) \not\equiv 0$, and
3. For $i \in [n-1]$, $\hat{A}(begin(x_{\phi(i)}), end(x_{\phi(i+1)})) \not\equiv 0$ (Since otherwise $\frac{\partial^2 f}{\partial x_{\phi(i)} \partial x_{\phi(i+1)}} \equiv 0$).

The coefficient of $\prod_{i\in[n]} x_i$ is just

$$\hat{A}(s, begin(x_{\phi(1)})) \cdot \hat{A}(end(x_{\phi(n)}), t) \prod_{i\in[n-1]} \hat{A}(begin(x_{\phi(i)}), end(x_{\phi(i+1)})),$$

and hence $\prod_{i\in[n]} x_i$ appears in f. A similar argument yields the statement for $\hat{A}(end(x_i), begin(x_j))$ and finishes the proof of the claim. $\square$

As in the proof of Lemma 6, write $\frac{\partial^2 f}{\partial x_a \partial x_b} = g \cdot (\beta x_i - \alpha)$, where g is a RO-ABP-polynomial that does not depend on x_i, and $-\alpha$ equals the sum of weights over all paths from $end(x_a)$ to $begin(x_b)$ not going over x_i. We have three cases:

1. Neither x_j nor x_k is the most adjacent variable to x_i in A. By above claim, x_a appears in a monomial of q_1, and x_b appears in a monomial q_2. Hence, there is a monomial in $q_1 q_2$ with $x_a x_b$. As $q_3 = \mu q_1 q_2$, for $\mu \neq 0$, the same can be said for q_3. But this implies $\alpha \neq 0$, as the coefficient of $x_a x_b$ is $-\alpha \cdot \hat{A}(end(x_j), begin(x_a))\hat{A}(end(x_b), begin(x_k))$.
2. x_j is not the most adjacent variable to x_i in A, but $x_k = x_b$. Then similarly $q_1 q_2$ has a monomial with x_a in it, and therefore the same holds for q_3. Therefore $\alpha \neq 0$, as the coefficient of x_a in q_3 is $-\alpha \cdot \hat{A}(end(x_j), begin(x_a))$.
3. $x_j = x_a$, but x_k is not the most adjacent variable to x_i in A. This is argued similarly as the second item.

This concludes the argument for this case.

Case II: In A, the x_i-layer lies before the x_j-layer and x_k-layer.

Wlog. assume that the x_j layer lies before the x_k layer. Similarly as in Case I, we write $\frac{\partial^2 f}{\partial x_j \partial x_k} = p_1 p_2 \cdot (q_1 q_2 x_i + q_3)$, but where now we have that

- $p_1 = \hat{A}_{end(x_j), begin(x_k)}$, and $p_2 = \hat{A}_{end(x_k), sink(A)}$,
- $q_1 = \hat{A}_{source(A), begin(x_i)}$,

- $q_2 = \hat{A}_{end(x_i),begin(x_j)}$,
- $q_3 = A[x_i \hat{=} 0]_{source(A),begin(x_j)}$.

Then $p_1p_2 \cdot (q_1q_2x_i + q_3) = h \cdot (\beta' x_i - \alpha')$. Since both p_1p_2 and h do not depend on x_i, it must be that $(\beta' x_i - \alpha') \mid (q_1q_2x_i + q_3)$. Similarly as before, we get $q_3 = \mu q_1 q_2$ for some $\mu \in \mathbb{F}$, $\mu \neq 0$.

The rest of the proof is similar to Case I. One argues that 1) when $x_j \neq x_b$, q_1q_2 contains a monomial with $x_a x_b$. To make $x_a x_b$ appear in a monomial q_3 we need $\alpha \neq 0$, and 2) when $x_j = x_b$, q_1q_2 contains a monomial with x_a, and to make x_a appear in a monomial of q_3, we need $\alpha \neq 0$.

Case III: In A, the x_i-layer lies after the x_j-layer and x_k-layer.
This case is symmetrical to Case II. □

We also need the following proposition:

Proposition 5. *Let $f \in \mathbb{F}[X]$ be a RO-ABP-polynomial with $|Var(f)| \geq 3$, and let $S \subseteq Var(f)$. Then f is X-pre-aligned on S if and only if $f' := (x_{n+1} + 1)f$ is $X \cup \{x_{n+1}\}$-pre-aligned on S.*

Proof. Let $X' = X \cup \{x_{n+1}\}$. It is easy to see that assuming f is X-pre-aligned on S, we have that f is X'-pre-aligned on S.

Conversely, assume f' is X'-pre-aligned on S. Let $x_i \in S$. Then there exist $x_j, x_k \in X' \backslash \{x_i\}$, such that $\frac{\partial^2 f'}{\partial x_j \partial x_k} = g(\beta x_i + \alpha)$, where g is a RO-ABP-polynomial that does not depend on x_i, and $\alpha = 0$ implies $\beta = 0$. If $x_{n+1} \notin \{x_j, x_k\}$, then $\frac{\partial^2 f'}{\partial x_j \partial x_k} = \frac{\partial^2 f}{\partial x_j \partial x_k}(x_{n+1}+1)$. Setting $x_{n+1} = 0$, we have that $\frac{\partial^2 f}{\partial x_j \partial x_k} = (g_{|x_{n+1}=0})(\beta x_i + \alpha)$. So we get the required X-pre-alignment of f on $\{x_i\}$. Otherwise, say wlog. $x_j = x_{n+1}$. We have that $\frac{\partial f}{\partial x_k} = \frac{\partial^2 f'}{\partial x_{n+1} \partial x_k} = g(\beta x_i + \alpha)$. One easily obtains the required X-pre-alignment of f on $\{x_i\}$, by taking one more ∂x_l, for some variable $x_l \in X \backslash \{x_i, x_k\}$, and then using Lemma 4. □

Proof of Lemma 8: The proof proceeds by induction on $|X|$. For the base case we take $|X| \leq 3$, in which case the statement clearly holds. Now suppose $|X| > 3$. Let $f' = f_{|x_n=\gamma}$, for some γ. Let $X' = X \backslash \{x_n\}$. Suppose f' is not X'-pre-aligned. Hence $|Var(f')| \geq 3$. We want to show this can happen for at most one γ.

Consider an arbitrary RO-ABP A computing f. Let $f_e = f(x_{n+1}+1)(x_{n+2}+1)(x_{n+3}+1)(x_{n+4}+1)$. Let $X_e := X \cup \{x_{n+1}, x_{n+2}, x_{n+3}, x_{n+4}\}$. By Proposition 5, f_e is X_e-pre-aligned on $Var(f)$. Let $f'_e := (f_e)_{|x_n=\gamma}$ and $X'_e := X' \cup \{x_{n+1}, x_{n+2}, x_{n+3}, x_{n+4}\}$. Note that $f'_e = f'(x_{n+1}+1)(x_{n+2}+1)(x_{n+3}+1)(x_{n+4}+1)$. So also by Proposition 5, f'_e is not X'_e-pre-aligned on $Var(f')$ if and only if f' is not X'-pre-aligned on $Var(f')$. We will show the former happens for at most one γ. So let us assume that f'_e is not X'_e-pre-aligned on $Var(f')$. We can easily obtain a RO-ABP A_e from A, which computes f_e. In this, we make sure x_{n+1} and x_{n+2} are the first and second variable in A_e, and x_{n+3} and x_{n+4} are the fore-last and last variable in A_e. For each $x_i \in Var(f')$, let x_{j_i} be the variable right after x_i in A^e, and let x_{k_i} be the variable before x_i in A_e. Note that we have made sure these always exist in A_e. Since f_e is X_e-pre-aligned on $Var(f)$, by Lemma 9, $\frac{\partial^2 f_e}{\partial x_{j_i} \partial x_{k_i}} = g \cdot (\beta_i x_i - \alpha_i)$, where g is a RO-ABP-polynomial that does not depend on x_i, and $\alpha_i = 0 \Rightarrow \beta_i = 0$. Furthermore, we have that α_i is the sum of weights of all paths from $end(x_{k_i})$ to $begin(x_n)$, which do not go over x_i in A_e. Consider the following two cases:

Case I: $n \notin \{j_i, k_i\}$, for any $x_i \in Var(f')$.

Then for any i, $\frac{\partial^2 f'_e}{\partial x_{j_i} \partial x_{k_i}} = (g_i)_{|x_n=\gamma} \cdot (\beta_i x_i - \alpha_i)$, which contradicts the assumption that f'_e is not X'_e-pre-aligned on $Var(f')$.

Case II: $n \in \{j_i, k_i\}$, for some $x_i \in Var(f')$.

By symmetry we can assume wlog. that $j_i = n$ (the case $k_i = n$ is handled similarly). Since $\frac{\partial^2 f}{\partial x_{j_i} \partial x_{k_i}} \not\equiv 0$, and $\alpha_i = 0$ implies $\beta_i = 0$, We have that $\alpha_i \neq 0$.

We know that in A_e there still exists a variables layer, say with variables x_l, right after the x_{j_i}-layer. Let $b_i = begin(x_i), e_i = end(x_i), b_n = begin(x_n)$, and $e_n = end(x_n)$. Let $s = end(x_{k_i})$ and $t = begin(x_l)$. Then write:

$$\frac{\partial^2 f_e}{\partial x_l \partial x_{k_i}} = p_1 p_2 (c_{s,b_i} c_{e_i,b_n} c_{e_n,t} x_i x_n + c_{s,b_i} c_{e_i,t} x_i + c_{s,b_n} c_{e_n,t} x_n + c_{s,t}),$$

where in the above each constant $c_{v,w}$ is the sum of weights over all paths from v to w going over constant labeled edges only. Note that $c_{s,b_n} = \alpha_i \neq 0$. Furthermore, p_1 is the sum of weights of all paths from $source(A_e)$ to $begin(x_{k_i})$, and p_2 is the sum of weights over all paths from $end(x_l)$ to $sink(A_e)$. Then

$$\frac{\partial^2 f'_e}{\partial x_l \partial x_{k_i}} = p_1 p_2 ((c_{s,b_i} c_{e_i,b_n} c_{e_n,t} \gamma + c_{s,b_i} c_{e_i,t}) x_i + c_{s,b_n} c_{e_n,t} \gamma + c_{s,t}),$$

We have that f'_e can only not be X'_e-pre-aligned on $\{x_i\}$ if $c_{s,b_n} c_{e_n,t} \gamma + c_{s,t} = 0$. This can happen for more than one γ only if $c_{s,b_n} c_{e_n,t} = 0$. Since $c_{s,b_n} \neq 0$, this happens only if $c_{e_n,t} = 0$, but the latter implies that $\frac{\partial^2 f_e}{\partial x_l \partial x_n} \equiv 0$, which in turn implies that $\frac{\partial^2 f}{\partial x_l \partial x_n} \equiv 0$, which is a contradiction.

Finally, putting together from what we observed from the above two cases, note that, Case II can apply at most twice for a variable $x_i \in Var(f')$. Namely, possibly once for the variable right before x_n, and possibly once for the variable after x_n. We conclude the lemma holds. □

Finally, we require that we can drop linear factors, while maintaining the pre-alignment property, which is what the following lemma gives us:

Lemma 10. *Let $g, h \in \mathbb{F}[X]$ be RO-ABP-polynomials such that $h = g \cdot (\beta x_n - \alpha)$, for $\beta \in \mathbb{F} \backslash \{0\}$ and $\alpha \in \mathbb{F}$. If h is X-pre-aligned, then g is $(X \backslash \{x_n\})$-pre-aligned.*

Proof. If $|X| \leq 3$ it is trivial, so assume $|X| > 3$. Let $x_i \in X \backslash \{x_n\}$. Since h is X-pre-aligned, there exist $j, k \in [n] \backslash \{i\}$ such that $\frac{\partial^2 h}{\partial x_j \partial x_k} = h' \cdot (\beta' x_i - \alpha')$, where h' does not depend on x_i and $\alpha' = 0 \Rightarrow \beta' = 0$. We consider two cases. First, suppose $n \notin \{j, k\}$. Since $g = h_{|x_n=(1+\alpha)/\beta}$, we get $\frac{\partial^2 g}{\partial x_j \partial x_k} = h'_{|x_n=(1+\alpha)/\beta} \cdot (\beta' x_i - \alpha')$. Now suppose $n \in \{j, k\}$, and wlog. assume $j = n$. Then $\frac{\partial^2 h}{\partial x_j \partial x_k} = \beta \frac{\partial g}{\partial x_k} = h' \cdot (\beta' x_i - \alpha')$. Since $|X| > 3$, we can easily take partial w.r.t. another variable x_l so that $\frac{\partial^2 g}{\partial x_k \partial x_l}$ is of the required form. □

We are now ready to prove the hardness of representation theorems.

5 Hardness of Representation Theorems

The following theorem is an adaption of Theorem 6.1 in [SV09] to the notion of X-pre-alignment.

Theorem 4. *Let $n > 2$ be an integer and $X = \{x_i : i \in [n]\}$ be a set of indeterminates. Let $P_n = \prod_{i\in[n]} x_i$. If $\{f_i \in \mathbb{F}[X]\}_{i\in[k]}$ is a set of k many X-pre-aligned RO-ABP-polynomials for which $P_n = \sum_{i\in[k]} f_i$, then $n < 7k$.*

Proof. The proof proceeds by induction on k. For the base case $k = 1$, since $f_1 = P_n$, and f_1 is X-pre-aligned, and $n > 2$, for $x_i \in Var(P_n)$, whatever distinct $x_j, x_l \in X\backslash\{x_i\}$ we select, $\frac{\partial^2 f_1}{\partial x_j \partial x_l} = x_i \cdot \prod_{x_r \in X\backslash\{x_i,x_j,x_l\}} x_r$. This cannot be of the form $g \cdot (\beta x_i + \alpha)$ with g being an RO-ABP not depending on x_i, and $\alpha = 0 \Rightarrow \beta = 0$, as Definition 1 requires. Namely, since g does not depend on x_i, it must be that $\beta \neq 0$. Hence $\alpha \neq 0$, and thus $g \cdot (\beta x_i + \alpha)$ is not homogeneous. Since $x_i \cdot \prod_{x_r\in X\backslash\{x_i,x_j,x_l\}} x_r$ is homogeneous, this is a contradiction. Now assume $k > 1$. Suppose we can write $P_n = \sum_{i\in[k]} f_i$. For purpose of contradiction, assume that $n \geq 7k$. Hence $n \geq 14$.

Case I: "$\exists$ distinct $p, q, r \in [n]$ and $s \in [k]$, such that $\frac{\partial^3 f_s}{\partial x_p \partial x_q \partial x_r} \equiv 0$".

Wlog. assume that $p = n-2, q = n-1, r = n$ and $s = k$. Then $\sum_{i\in[k-1]} \frac{\partial^3 f_i}{\partial x_{n-2}\partial x_{n-1}\partial x_n} = P_{n-3}$.

By Lemma 7, all of the terms $\frac{\partial^3 f_i}{\partial x_{n-2}\partial x_{n-1}\partial x_n}$ are $(X\backslash\{x_{n-2}, x_{n-1}, x_n\})$-pre-aligned. By induction, it must be that $n - 3 < 7(k-1)$. Hence $n < 7k - 4$, which is a contradiction.

Case II: "$\forall$ distinct $p, q, r \in [n]$ and $s \in [k]$, we have that $\frac{\partial^3 f_s}{\partial x_p \partial x_q \partial x_r} \not\equiv 0$".

We know $\forall i$, $|Var(f_i)| \geq 3$. Since f_i is X-pre-aligned, there exist distinct $x_{j_i}, x_{l_i} \in X\backslash\{x_n\}$ such that $\frac{\partial^2 f_i}{\partial x_{j_i}\partial x_{l_i}} = g_i \cdot (\beta_i x_n - \alpha_i)$, where g_i is a RO-ABP-polynomial that does not depend on x_n, and $\alpha_i = 0 \Rightarrow \beta_i = 0$. Note that in this case, $g_i \not\equiv 0$, since otherwise a second order partial vanishes. Hence both j_i and l_i are certainly not equal to x_n. It must be that $\beta_i \neq 0$, since otherwise $\frac{\partial^3 f}{\partial x_{j_i}\partial x_{l_i}\partial x_n} \equiv 0$. Hence also $\alpha_i \neq 0$.

Claim 3. *g_i is $(X\backslash\{x_{j_i}, x_{l_i}, x_n\})$-pre-aligned.*

Proof. Assume that $|Var(g_i)| \geq 3$, since otherwise the claim is trivial. Let $h = g_i \cdot (\beta_i x_n - \alpha_i)$. By Lemma 7, h is $(X\backslash\{x_{j_i}, x_{l_i}\})$-pre-aligned. Since $\beta_i \neq 0$, applying Lemma 10 yields that g_i is $(X\backslash\{x_{j_i}, x_{l_i}, x_n\})$-pre-aligned. □

Now, let $A = \{\frac{\alpha_i}{\beta_i} : i \in [k]\}$. Define for $\gamma \in A$,

$$E_\gamma = \{i \in [k] : \gamma = \frac{\alpha_i}{\beta_i}\}$$

and

$$B_\gamma = \{i \in [k] : \gamma \neq \frac{\alpha_i}{\beta_i} \text{ and } (f_i)_{|x_n=\gamma} \text{ is not } (X\backslash\{x_n\})\text{-pre-aligned}\}.$$

Note that $\sum_{\gamma\in A} |E_\gamma| = k$. By Nearly Unique Nonalignment Lemma 8, $\sum_{\gamma\in A} |B_\gamma| \leq 2k$. Hence there exists $\gamma_0 \in A$ such that $|B_{\gamma_0}| \leq 2|E_{\gamma_0}|$. Let $I = E_{\gamma_0} \cup B_{\gamma_0}$, and let $J = \{j_i : i \in I\} \cup \{k_i : i \in I\}$. We have that $2 \leq |J| \leq 2|I| \leq 6|E_{\gamma_0}|$. Observe that $x_n \notin J$. Define for any i, $f'_i = \partial_J f_i$. We have the following three properties:

1. Each f'_i is an $(X\backslash J)$-pre-aligned RO-ABP-polynomial, due to Lemma 7.

2. For every $i \in I$, $f'_i = (\beta_i x_n - \alpha_i)h_i$, where h_i is a RO-ABP-polynomial. Namely, since $j_i, l_i \in J$, $f'_i = \partial_{J\backslash\{j_i,l_i\}}[g_i(\beta_i x_n - \alpha_i)] = (\beta_i x_n - \alpha_i) \cdot \partial_{J\backslash\{j_i,l_i\}} g_i$.

3. In the above, each h_i is an $(X\backslash(J \cup \{x_n\}))$-pre-aligned RO-ABP-polynomial. Namely, since g_i is $(X\backslash\{x_{j_i}, x_{l_i}, x_n\})$-pre-aligned. Hence, using Lemma 7, we get that h_i is an $(X\backslash(J \cup \{x_n\}))$-pre-aligned RO-ABP-polynomial.

For any i, define $f''_i = (f'_i)_{|x_n=\gamma_0}$. Then we have the following three properties:

1. $\forall i \in E_{\gamma_0}$, $f''_i \equiv 0$.

2. $\forall i \in B_{\gamma_0}$, $f''_i = (\beta_i\gamma_0 - \alpha_i)h_i$, so f''_i is an $(X\backslash(J \cup \{x_n\}))$-pre-aligned RO-ABP-polynomial, due to Proposition 4.

3. For every $i \in [k]\backslash I$, $(f_i)_{|x_n=\gamma_0}$ is $X\backslash\{x_n\}$-pre-aligned. Hence, since $n \notin J$, $f''_i = (f'_i)_{|x_n=\gamma_0} = \partial_J[f_{|x_n=\gamma_0}]$. So by Lemma 7, f''_i is an $(X\backslash(J \cup \{x_n\}))$-pre-aligned RO-ABP-polynomial.

Wlog. assume that $J = \{\tilde{n}+1, \tilde{n}+2, \dots, n-2, n-1\}$. Then $|J| = n - 1 - \tilde{n}$. Then $\sum_{i\in[k]} f''_i = (\partial_J P_n)_{|x_n=\gamma_0} = \gamma_0 \cdot P_{\tilde{n}}$. Let $\tilde{X} = \{x_1, \dots, x_{\tilde{n}}\}$. We have found a representation of $P_{\tilde{n}}$ as a sum of $\tilde{k}$ $\tilde{X}$-pre-aligned RO-ABP-polynomials, where $7\tilde{k} \le 7(k - |E_{\gamma_0}|) \le n - 7|E_{\gamma_0}| = n - 1 - 6|E_{\gamma_0}| + 1 - |E_{\gamma_0}| \le \tilde{n} + 1 - |E_{\gamma_0}| \le \tilde{n}$. This contradicts the induction hypothesis. □

Next we generalize to PRO-ABP-polynomials.

Theorem 5. *Let $n > 2$ be an integer and $X = \{x_i : i \in [n]\}$ be a set of indeterminates. Let $P_n = \prod_{i\in[n]} x_i$ and let $g \in \mathbb{F}[X]$ be a nonzero polynomial. If $\{f_i \in \mathbb{F}[X]\}_{i\in[k]}$ is a set of k many X-pre-aligned PRO-ABP-polynomials for which $g \cdot P_n = \sum_{i\in[k]} f_i$, then $n < 7k$.*

Proof. We use induction on k. Let d be such that each f_i is d-decomposable, and let $(h_i, \{T^i_j(x_j\}_{j\in[n]})$ be a d-decomposition of f_i, where $h_i \in \mathbb{F}[Z]$ and $Z = \{z_i : i \in [n]\}$. By definition, each h_i is Z-pre-aligned. Wlog. we can assume $|\mathbb{F}|$ is infinite, since if the statement holds over some infinite extension field of $\mathbb{F}$, then it holds for $\mathbb{F}$.

For $k = 0$ the statement is trivial. For $k = 1$, for purpose of contradiction suppose that $g \cdot P_n = f_1 = h_1(T^1_1(x_1), \dots, T^1_n(x_n))$, where h_1 is Z-pre-aligned on $Z = \{z_i : i \in [n]\}$. Hence for any $z_i \in Var(h_1)$, there exist $x_j, x_l \in Z\backslash\{z_i\}$, such that $\frac{\partial^2 h_1}{\partial z_j \partial z_l} = h'_1 \cdot (\beta z_i - \alpha)$, where h'_1 is a RO-ABP-polynomial not depending on z_i, and $\alpha = 0 \Rightarrow \beta = 0$.

Now let $\gamma \neq 0$ be such that $g_{|x_j=x_l=\gamma} \not\equiv 0$, which exists since $g \not\equiv 0$ and $|\mathbb{F}|$ is infinite. We have that

$$\frac{\partial^2 f_1}{\partial_\gamma x_j \partial_\gamma x_l} = T^1_j(\gamma)T^1_l(\gamma) \cdot h'_1(T^1_1(x_1), \dots, T^1_n(x_n)) \cdot (\beta T^1_i(x_1) - \alpha).$$

Therefore, $\frac{\partial^2 f_1}{\partial_\gamma x_j \partial_\gamma x_l}$ contains a term without x_i, if it is not identically zero. Also, $\frac{\partial^2 f_1}{\partial_\gamma x_j \partial_\gamma x_l} = \gamma^2 \cdot g_{|x_j=x_l=\gamma} \cdot \prod_{t\in[n]\backslash\{j,l\}} x_t$, which implies it is not identically zero, and every term contains the variable x_i. We have reached a contradiction and have proven the case $k = 1$.

For the induction step consider the case that $k \ge 2$. For purpose of contradiction that suppose $n \ge 7k$ and $g \cdot P_n = \sum_{i\in[k]} h_i(T^i_1(x_1), \dots, T^i_n(x_n))$. Let y and w be new variables. We make a distinction between two cases.

Case I: $\forall j \in [n], i \in [k]$, we have that $y \cdot g_{|x_j=y} T^i_j(w) \equiv w \cdot g_{|x_j=w} T^i_j(y)$.

This implies for $i_1 \neq i_2$ and any $j \in [n]$, that $T_j^{i_1}(y)/T_j^{i_2}(y) = T_j^{i_1}(w)/T_j^{i_2}(w)$. This means $T_j^{i_1}(y)/T_j^{i_2}(y) \in \mathbb{F}$, and since both polynomials are monic $T_j^{i_1}(y)/T_j^{i_2}(y) = 1$. Hence in this case there exists a single set $\{U_j \in \mathcal{T}_d\}_{j\in[n]}$ such that $\forall i \in [k]$, $h_i = f_i(U_1(x_1, \dots, U_n(x_n))$.

Observe that $y \cdot g_{|x_j=y}/U_j(y) = w \cdot g_{|x_j=w}/U_j(w)$ and that consequenty for some $g'_j \in \mathbb{F}[X]$ with $x_j \notin Var(g'_j)$, $y \cdot g_{|x_j=y}/U_j(y) = g'_j$. Hence $\forall i \in [n]$, $U_i(x_i)$ is a factor of $g \cdot P_n$, and

$$g \cdot P_n = c \cdot U_1(x_1)U_2(x_2)\dots, U_n(x_n) = cP_n(U_1(x_1), \dots, U_n(x_n)),$$

for some $c \in \mathbb{F}\backslash\{0\}$. Therefore,

$$\sum_{i\in[k]} h_i(U_1(x_1), \dots, U_n(x_n)) = \sum_{i\in[k]} f_i = cP_n(U_1(x_1), \dots, U_n(x_n)).$$

From which we conclude that $\sum_{i\in[k]} h_i(x_1, \dots, x_n) = cP_n(x_1), \dots, (x_n))$. By theorem 4 this implies that $n < 7k$, which is a contradiction.

Case II: $\exists j \in [n], i \in [k]$, such that $y \cdot g_{|x_j=y}T_j^i(w) \not\equiv w \cdot g_{|x_j=w}T_j^i(y)$.

Wlog. assume that $j = n$ and $i = k$. Since $|\mathbb{F}|$ is infinite, there exist $\alpha, \beta \in \mathbb{F}$, such that $\alpha \cdot g_{|x_n=\alpha}T_n^k(\beta) \not\equiv \beta \cdot g_{|x_n=\beta}T_n^k(\alpha)$. We have that $\frac{\partial(g\cdot P_n)}{\partial_\alpha x_n} = \alpha \cdot g_{|x_n=\alpha}P_{n-1}$ and $\frac{\partial(g\cdot P_n)}{\partial_\beta x_n} = \beta \cdot g_{|x_n=\beta}P_{n-1}$. Also

$$\frac{\partial(g \cdot P_n)}{\partial_\alpha x_n} = \sum_{i\in[k]} T_n^i(\alpha)\frac{\partial h_i}{\partial x_n}(T_1^i(x_1), \dots, T_n^i(x_n))$$

and

$$\frac{\partial(g \cdot P_n)}{\partial_\beta x_n} = \sum_{i\in[k]} T_n^i(\beta)\frac{\partial h_i}{\partial x_n}(T_1^i(x_1), \dots, T_n^i(x_n))$$

Hence

$$\begin{aligned} & P_{n-1}\left(\alpha \cdot g_{|x_n=\alpha}T_n^k(\beta) - \beta \cdot g_{|x_n=\beta}T_n^k(\alpha)\right) \\ = & \sum_{i\in[k-1]} (T_n^i(\alpha)T_n^k(\beta) - T_n^i(\beta)T_n^k(\alpha))\frac{\partial h_i}{\partial x_n}(T_1^i(x_1), \dots, T_n^i(x_n)) \end{aligned}$$

By Lemma 7, $\frac{\partial h_i}{\partial x_n}$ is an Z-pre-aligned RO-ABP-polynomial. We conclude that for some polynomial $g' \not\equiv 0$, we have a representation of $P_{n-1} \cdot g'$ as sum of $k-1$ X-pre-aligned PRO-ABP-polynomials. By induction hypothesis, it must be that $n < 7k - 6$. This is a contradiction. □

6 Computing a Simultaneous Alignment

A *simultaneous X-alignment* for a set of (P)RO-ABP-polynomials $\{f_i \in \mathbb{F}[X]\}_{i\in[k]}$ is any vector $v \in \mathbb{F}^n$ such that $f_i(x_1 + v_1, x_2 + v_2, \dots, x_n + v_n)$ is X-aligned for every $i \in [k]$. We present an algorithm for finding a simultaneous X-alignment for a set of PRO-ABP-polynomials. We assume that we have a polynomial identity testing algorithm $\mathrm{PIT}_{\text{PRO-ABP}}$ for testing a single PRO-ABP.

First, we establish a sufficient condition, so for a given RO-ABP-polynomial f we can make $f(x_1 + v_1, x_2 + v_2, \dots, x_n + v_n)$ X-aligned, by means of computing some shift $v \in \mathbb{F}^n$. For this, let us call a polynomial $f \in \mathbb{F}[X]$ *decent*, if for all $x_a, x_b \in Var(f)$ with $\frac{\partial^2 f}{\partial x_a \partial x_b} \not\equiv 0$, it holds that the monomial $x_a x_b$ appears in f with a nonzero constant coefficient.

Lemma 11. *A RO-ABP-polynomial $f \in \mathbb{F}[X]$ is X-aligned, if $|Var(f)| \leq 2$, or else for every $I \subseteq Var(f)$ with $|I| \leq |Var(f)| - 3$, $f_{|x_I=0}$ is decent.*

Proof (Induction on $|Var(f)|$). For $|Var(f)| \leq 2$ it is trivial. Now assume $|Var(f)| > 2$. Take $I = \emptyset$. Then we get that for any $x_a, x_b \in Var(f)$, if $\frac{\partial^2 f}{\partial x_a \partial x_b} \not\equiv 0$ then the monomial $x_a x_b$ appears in f with a nonzero constant coefficient. Let us first establish that f is X-pre-aligned. Consider an arbitrary $x_i \in Var(f)$. By Lemma 6, there exist distinct $x_j, x_k \in X \backslash \{x_i\}$ such that $p := \frac{\partial^2 f}{\partial x_j \partial x_k} = g \cdot (\beta x_i - \alpha)$, where g is a RO-ABP-polynomial that does not depend on x_i, and $\alpha, \beta \in F$.

If $\beta = 0$, then f is X-pre-aligned on $\{x_i\}$, so suppose $\beta \neq 0$. If p is identically zero, then we know $g \equiv 0$, so $\frac{\partial^2 f}{\partial x_j \partial x_k} = g \cdot (\beta x_i - \alpha')$, for any arbitrary $\alpha' \neq 0$. If p is not identically zero, then we know $x_j x_k$ is in f, which implies that $\alpha \neq 0$. We conclude that f is X-pre-aligned on $\{x_i\}$.

In the above, we find that f is X-pre-aligned on $\{x_i\}$ in any of the considered cases. Since x_i was arbitrarily taken from $Var(f)$, we conclude that f is X-pre-aligned.

Next, we show Condition 2 of Definition 2 holds. Consider $f' := f_{|x_i=0}$, for an arbitrary $x_i \in Var(f)$. We want to establish that the sufficient condition of Lemma 11 holds for $f' \in \mathbb{F}[X \backslash \{x_i\}]$, since then we can by apply the induction hypothesis and conclude that f' is $(X \backslash \{x_i\})$-aligned.

If $|Var(f')| \leq 2$ the sufficient condition of the Lemma 11 clearly holds for f'. Otherwise, consider $I' \subseteq Var(f')$ of size at most $|Var(f')| - 3$. Let $I = I' \cup \{x_i\}$. Then $|I| \leq |Var(f)| - 3$. Now consider $x_a, x_b \in Var(f'_{x_{I'}=0}) = Var(f_{x_I=0})$. Suppose $\frac{\partial^2 f'_{|x_{I'}=0}}{\partial x_a \partial x_b} \not\equiv 0$. Since the latter equals $\frac{\partial^2 f_{|x_I=0}}{\partial x_a \partial x_b} \not\equiv 0$, we know that $x_a x_b$ appears with a nonzero constant coefficient in $f_{|x_I=0}$. This implies $x_a x_b$ appears with a nonzero constant coefficient in $f_{|x_{I'}=0}$. Hence $f'_{x_{I'}=0}$ is decent.

We conclude the sufficient condition of the Lemma 11 holds for $f' \in \mathbb{F}[X \backslash \{x_i\}]$. Hence by the induction hypothesis we conclude that f' is $(X \backslash \{x_i\})$-aligned. $\square$

Lemma 12. *Any decent RO-ABP-polynomial $f \in \mathbb{F}[X]$ is X-aligned.*

Proof. We show that the condition of Lemma 11 is satisfied. If $|Var(f)| \leq 2$ this is clear. Otherwise, consider arbitrary $I \subseteq Var(f)$ with $|I| \leq |Var(f)| - 3$. Let $x_a, x_b \in Var(f_{|x_I=0})$, be such that $\frac{\partial^2 f_{|x_I=0}}{\partial x_a \partial x_b} \not\equiv 0$. We have that $x_a, x_b \in Var(f)$, and it must be that $\frac{\partial^2 f}{\partial x_a \partial x_b} \not\equiv 0$, since $\frac{\partial^2 f_{|x_I=0}}{\partial x_a \partial x_b} = \left(\frac{\partial^2 f}{\partial x_a \partial x_b}\right)_{|x_I=0}$. Hence $x_a x_b$ is in f. This implies that $x_a x_b$ is in $f_{|x_I=0}$. $\square$

The above lemma leads the way towards computing a simultaneous X-alignment as follows:

Corollary 2. *Let $\{f_i\}_{i \in [k]}$ be a set of RO-ABP-polynomials in $\mathbb{F}[X]$. If $v \in \mathbb{F}^n$ is a simultaneous nonzero of $\{\frac{\partial^2 f_i}{\partial x_a \partial x_b} \mid \frac{\partial^2 f_i}{\partial x_a \partial x_b} \not\equiv 0\}_{i \in [k], a, b \in [n]}$, then v is a simultaneous X-alignment for $\{f_i\}_{i \in [k]}$.*

Proof. Consider $\{f'_i = f_i(x_1 + v_1, x_2 + v_2, \ldots, x_n + v_n)\}_{i \in [k]}$. Due to Lemma 12, we only need to show that for every i, for every $x_a, x_b \in Var(f_i)$, if $\frac{\partial^2 f'_i}{\partial x_a \partial x_b} \not\equiv 0$ then the monomial $x_a x_b$ appears in f'_i with a nonzero constant coefficient. Observe that the monomial $x_a x_b$ appears in f'_i with a nonzero constant coefficient $\iff \frac{\partial^2 f'_i}{\partial x_a \partial x_b}(\bar{0}) \neq 0$. The latter holds, as $\frac{\partial^2 f'_i}{\partial x_a \partial x_b}(\bar{0}) = \frac{\partial^2 f_i}{\partial x_a \partial x_b}(v) \neq 0$. $\square$

The above corollary can be generalized to PRO-ABP-polynomials

Corollary 3. *Let $k > 0$ and $d > 0$ be integers and suppose $\mathbb{F}$ is a field with more than knd elements. Let $\{f_i\}_{i\in[k]}$ be a set of PRO-ABP-polynomials in $\mathbb{F}[X]$. Suppose f_i has d-decomposition $(g_i, \{T_j^i(x_j)\}_{j\in[n]})$, where $g_i \in \mathbb{F}[Z]$, for all $i \in [k]$. Suppose $\alpha \in \mathbb{F}\backslash\{0\}$ satisfies $T_j^i(\alpha) \neq 0$, for all $i \in [k]$ and $x_j \in Var(f_i)$. If $v \in \mathbb{F}^n$ is a simultaneous nonzero for $\{\frac{\partial^2 f_i}{\partial_\alpha x_a \partial_\alpha x_b} \mid \frac{\partial^2 f_i}{\partial_\alpha x_a \partial_\alpha x_b} \not\equiv 0\}_{i\in[k],a,b\in[n]}$, then v is a simultaneous X-alignment of $\{f_i\}_{i\in[k]}$.*

Proof. For certain $U_1^i, \dots, U_n^i \in \mathcal{T}_d$, we can write $f_i(x_1 + v_1, \dots, x_n + v_n) = g_i(T_1^i(x_1 + v_1), \dots, T_n^i(x_n + v_n)) = g_i(U_1^i(x_1) + T_1^i(v_1), \dots, U_n^i(x_n) + T_n^i(v_n))$. So letting $g_i' = g_i(z_1 + T_1^i(v_1), \dots, z_n + T_n^i(v_n))$, gives us that $(g_i', \{U_1^i(x_j)\}_{j\in[n]})$ is a d-decomposition of $f_i(x_1+v_1, \dots, x_n+v_n)$. We want to select a single v such that every RO-ABP-polynomial g_i' is Z-aligned. Similarly as in the proof of Corollary 2, we can arrange this by ensuring that for each $i \in [k]$, $(T_1^i(v), \dots, T_n^i(v))$ is a common nonzero of $\{\frac{\partial^2 g_i}{\partial z_a \partial z_b} : \frac{\partial^2 g_i}{\partial z_a \partial z_b} \not\equiv 0\}_{a,b\in[n]}$. By Lemma 4, we have that $\frac{\partial^2 f_i}{\partial_\alpha x_a \partial_\alpha x_b} = T_a^i(\alpha) T_b^i(\alpha) \frac{\partial^2 g_i}{\partial z_a \partial z_b}(T_1^i(x_1), \dots, T_n^i(x_n))$. Note that $\frac{\partial^2 g_i}{\partial z_a \partial z_b} \not\equiv 0$ implies that $z_a, z_b \in Var(g_i)$. Since the $|\mathbb{F}| > d$, we have that for any $j \in [n]$, $x_j \in Var(f_i) \Leftrightarrow z_j \in Var(g_i)$. Hence, it suffices to find a single v that is a nonzero of $\{\frac{\partial^2 f_i}{\partial_\alpha x_a \partial_\alpha x_b} \mid \frac{\partial^2 f_i}{\partial_\alpha x_a \partial_\alpha x_b} \not\equiv 0\}_{i\in[k],a,b\in[n]}$. □

In order to apply the above corollary, we have to deal with the issue of finding the appropriate direction α. This is not too difficult as any set of size $knd + 1$ contains such an element. Namely, by Proposition 2, for any f_i and $x_j \in Var(f_i)$, there can be at most d values for α with $\frac{\partial f_i}{\partial_\alpha x_j} \equiv 0$. Using the procedure $\mathrm{PIT}_{\text{PRO-ABP}}$ on the f_i's we can single out one correct element in the interval $[knd + 1]$.

Now we proceeds similarly as in Lemma 4.3 of [SV09], but with first order partial derivatives replaced by second order ones. This yields the following theorem.

Theorem 6. *Let $d \geq 1$ be an integer, and suppose $\mathbb{F}$ is a field with $|\mathbb{F}| > dkn^2$. There exists an algorithm for finding a simultaneous X-alignment for a set of d-decomposable PRO-ABP polynomials $\{f_i \in \mathbb{F}[X]\}_{i\in[k]}$. The algorithm makes oracle calls to the procedure $\mathrm{PIT}_{\text{PRO-ABP}}$. The f_is are only accessed through this subroutine. The running-time of the algorithm is $O(dk^2n^5 \cdot t)$, where t is an upper bound on the time needed for any subroutine call to $\mathrm{PIT}_{\text{PRO-ABP}}$.*

Proof. We assume that we have a polynomial identity testing algorithm $\mathrm{PIT}_{\text{PRO-ABP}}$ for testing a single PRO-ABP, such that $\mathrm{PIT}_{\text{PRO-ABP}}$ outputs $True$ if $f \equiv 0$ and $False$ otherwise. We first state Algorithm 1 for computing $Var(f)$ of a PRO-ABP-polynomial f. Its correctness follows from Proposition 2.

Algorithm 1 Computing $Var(f)$.

Input: A d-decomposable PRO-ABP-polynomial $f \in \mathbb{F}[X]\}$.
Assumption: $|\mathbb{F}| > d$.
Output: $Var(f)$.
Oracle: PIT algorithm $\mathrm{PIT}_{\text{PRO-ABP}}$.

```
1: S = ∅
2: for all x ∈ X, α ∈ [d + 1] do
3:    If PIT_PRO-ABP(∂f/∂_α x) = False, add x to S
4: end for
5: return S
```

Algorithm 2 Alignment Finding.

Input: A set of d-decomposable PRO-ABP-polynomials $\{f_i \in \mathbb{F}[X]\}_{i\in[k]}$.
Assumption: $|\mathbb{F}| > dkn^2$.
Output: A simultaneous alignment v for $\{f_i\}_{i\in[k]}$.
Oracle: PIT algorithm $\mathrm{PIT}_{\text{PRO-ABP}}$.

1: Compute $Var(f_i)$, for each $i \in [k]$.
2: **for all** $t \in [knd+1]$ **do**
3:    If for all $i \in [k]$ and $x_j \in Var(f_i)$, $\mathrm{PIT}_{\text{PRO-ABP}}(\frac{\partial f_i}{\partial_t x_j}) = False$, set $\alpha = t$, exit for loop.
4: **end for**
5: $L = \emptyset$
6: **for all** f_i and (x_a, x_b), $a, b \in [n]$, $a \neq b$ **do**
7:    If $\mathrm{PIT}_{\text{PRO-ABP}}(\frac{\partial^2 f_i}{\partial_\alpha x_a \partial_\alpha x_b}) = False$, add it to L
8: **end for**
9: **for all** $j \in [n]$ **do**
10:    Find c such that for every $g \in L$, $\mathrm{PIT}_{\text{PRO-ABP}}(g\mid_{x_j=c}) = False$
11:    $v_j \leftarrow c$
12:    For every $g \in L$, $g \leftarrow g\mid_{x_j=c}$
13: **end for**
14: **return** v

Correctness of Algorithm 2: We first make two remarks, which pertain to applying Algorithm 2 in the setting where we only have black-box access to each f_i. Consider the first **for**-loop. Since we only have black-box access to f_i, the given pseudocode should be interpreted symbolically. Namely, by Lemma 4, $f' := \frac{\partial f_i}{\partial_t x_j}$ is a PRO-ABP. Note that black-box access to f_i is sufficient for being able to compute $f'(a)$ for any $a \in \mathbb{F}^n$. This is all the black-box algorithm $\mathrm{PIT}_{\text{PRO-ABP}}$ needs to decide whether $f' \equiv 0$. A similar remark pertains to line 7.

Also similarly, on line 12 the substitution is not actually carried out, but done symbolically. So it is just remembered that x_j is set to c. For example, suppose that up to some point in the execution the algorithm it has set $x_i = c_i$, for $i \in [m]$. Then on line 10, for evaluating $\mathrm{PIT}_{\text{RO-ABP}}(g\mid_{x_j=c})$, the black-box algorithm is granted access to a PRO-ABP in $n-m$ variables $g(c_1, c_2, \ldots, c_m, x_{m+1}, \ldots, x_n)$. The queries it makes can be answered with only black-box access to g.

First the algorithm finds an α such that for all $i \in [k]$ and $x_j \in Var(f_i)$, $\frac{\partial f_i}{\partial_t x_j} \not\equiv 0$. Note that oen can derive using Lemma 3 that for each f_i and any $x_j \in Var(f_i)$, there are at most d values t in $[knd+1]$ for which $\frac{\partial f_i}{\partial_t x_j} \equiv 0$. Hence for some $t \in [knd+1]$, all tests on line 3 will pass, and an α will be set. Suppose that $(h_i, \{T_1^i(x_j)\}_{j\in[n]})$ is a d-decomposition of f_i, for all $i \in [k]$. Since $\frac{\partial f_i}{\partial_t x_j} = T_j^i(t)\frac{\partial h_i}{\partial z_j}(T_1(x_1), \ldots, T_n(x_n))$, and since for any $x_j \in Var(f_i)$, $\frac{\partial h_i}{\partial z_j} \not\equiv 0$ (and therefore $\frac{\partial h_i}{\partial z_j}(T_1(x_1), \ldots, T_n(x_n)) \not\equiv 0$), the selected α is a common nonzero of $\{T_j^i : i \in [k], j \in Var(f_i)\}$. Now, by Corollary 3 it suffices to find a common nonzero of the set L.

First however, we need to explain how to find c such that $g\mid_{x_j=c} \not\equiv 0$. Let $V \subset \mathbb{F}$ with $|V| = dkn^2 + 1$ be given. We claim V always includes a good value. This is because we have at most kn^2 polynomials in L and each has individual degrees bounded by d. For specific polynomial in L, there are at most d one bad values due to Lemma 3. The algorithm can simply try all elements in V to get the required c. The correctness of the algorithm is now evident, from the observation that it simply maintains the invariant that all $g \in L$ are not identically zero.

The running time of the algorithm is as follows. Line 1 takes $O(dkn)$ time. For line 3 we need $O(k^2n^2d)$ calls to $\mathrm{PIT}_{\text{PRO-ABP}}$. For line 7 we need $O(kn^2)$ calls to $\mathrm{PIT}_{\text{PRO-ABP}}$. For line 10 we need $O(n \cdot (dkn^2+1) \cdot (kn^2)) = O(dk^2n^5)$ calls to $\mathrm{PIT}_{\text{PRO-ABP}}$. Thus the total running time of the algorithm is $O(dk^2n^5 \cdot t)$, where t is an upper bound on the time needed for any subroutine call to $\mathrm{PIT}_{\text{RO-ABP}}$. $\square$

For the black-box setting we need the following lemma.

Lemma 13. *Let $\mathbb{F}$ be a field with $|\mathbb{F}| > k^2d^2n^5 + kdn^4$, and let $V \subseteq \mathbb{F}$ with $|V| = k^2d^2n^5 + kdn^4 + 1$ be given. Let $\{f_i\}_{i\in[k]}$ be a set of PRO-ABP-polynomials in $\mathbb{F}[X]$. Let $G_m : \mathbb{F}^{2m} \to \mathbb{F}^n$ be the mth-order SV-generator with $m = \lceil \log n \rceil + 1$. Then $\mathcal{A}_k := G_m(V^{2m})$ contains a simultaneous X-alignment for $\{f_i\}_{i\in[k]}$.*

Proof. let $L = \{\frac{\partial^2 f_i}{\partial_\alpha x_a \partial_\alpha x_b} \mid \frac{\partial^2 f_i}{\partial_\alpha x_a \partial_\alpha x_b} \not\equiv 0\}_{i\in[k],a,b\in[n],\alpha\in[knd+1]}$. Let $P(x_1,\ldots,x_n) = \prod_{g\in L} g(x_1,\ldots,x_n)$. By Lemma 4, each $g \in L$ is a PRO-ABP-polynomial. Hence by Lemma 1, for $m = \lceil \log n \rceil + 1$, the SV-generator $(G_m^1, G_m^2, \ldots, G_m^n)$, satisfies that $g(G_m^1, G_m^2, \ldots, G_m^n) \not\equiv 0$, for all $g \in L$. So $P(G_m^1, G_m^2, \ldots, G_m^n) \not\equiv 0$.

Note that there are $2m$ variables in $P(G_m^1, \ldots, G_m^n)$, and the degree of every variable is bounded by $(knd+1)kn^2 \cdot dn^2 = k^2d^2n^5 + kdn^4$. Thus by Lemma 3, $\exists a \in V^{2m}, P(G_m^1(a), \ldots, G_m^n(a)) \neq 0$. Hence $\mathcal{A}_k = G_n(V^{2m})$ is ensured to contain a nonzero of P. Any nonzero of P is a simultaneous nonzero of all $g \in L$. By Corollary 3 and the remark after it regarding finding an appropriate α, $\mathcal{A}_k$ contains a simultaneous X-alignment for $\{f_i\}_{i\in[k]}$. $\square$

7 A Vanishing Theorem and the PIT Algorithms

Theorem 7. *Let $n > 2$ and $d > 0$ be integers. Let $\{f_i \in \mathbb{F}[X]\}_{i\in[k]}$ be a set of k many d-decomposable X-aligned PRO-ABPs. Let $f = \sum_{i\in[k]} f_i$. Then $f \equiv 0 \iff \forall w \in \mathcal{W}^n_{7k,d}, f(w) = 0$.*

Proof. (induction on n). We only argue "$\Leftarrow$". Assume that $\forall w \in \mathcal{W}^n_{7k,d}, f(w) = 0$. For $n < 7k$ it follows from Lemma 3 that $f \equiv 0$. Now assume that $n \geq 7k$. Consider a variable x_ℓ, for $\ell \in [n]$ and restriction of the polynomials f_i's and f to the subspace $x_\ell = 0$. Each of the $f_i' = f_i|_{x_\ell=0}$ are $(X \setminus \{x_\ell\})$-aligned. Let $f' = \sum_{i=1}^k f_i'$. Clearly, $\forall w \in \mathcal{W}^n_{7k,d}, f(w) = 0$ implies that $\forall w \in \mathcal{W}^{n-1}_{7k,d}, f'(w) = 0$. By induction, $f' = f|_{x_\ell=0} \equiv 0$, which implies that x_ℓ divides f. So we get that $P_n = \prod_{i=1}^k x_i$ divides f, i.e. for some polynomial g we have that $P_n \cdot g = f$. Thus $P_n \cdot g$ is the sum of k RO-ABPs which are also X-aligned. Since $n \geq 7k$, by Theorem 4, we get $g \equiv 0$. So $f \equiv 0$. $\square$

Now we explain how to get the PIT algorithms for Σ_k-PRO-ABP-polynomials given by $\{f_i \in \mathbb{F}[X]\}_{i\in[k]}$ with individual degrees bounded by d. We use that $\forall v \in \mathbb{F}^n$, $f \equiv 0 \iff f(x_1 + v_1, x_2 + v_2, \ldots, x_n + v_n) \equiv 0$. If we have a common alignment v for $\{f_i\}_{i\in[k]}$, we know that each $f_i(x_1+v_1, x_2+v_2, \ldots, x_n+v_n)$ is X-aligned. Then Theorem 7 is applicable, and it suffices to test on the set $\mathcal{W}^n_{7k,d}$. Based on three approaches to get a common alignment, we get the following:

(*Black-box Setting*) We have black-box access to $f = \sum_{i\in[k]} f_i$. Let $f_v(x_1,\ldots,x_n) = f(x_1 + v_1, \ldots, x_n + v_n)$. Then $f \equiv 0 \iff \forall v \in \mathcal{A}_k, \forall w \in \mathcal{W}^n_{7k,d}, f_v(w) = 0$, where $\mathcal{A}_k$ is given by Lemma 13. So we get running-time $(kdn)^{O(\log n + k)} \leq (dn)^{O(\log n + k)}$. This proves Theorem 1.

(*Non/Semi Black-box Settings*) As we showed in Section 3.1, for non black-box, $\mathrm{PIT}_{\text{RO-ABP}}$ takes time $O(n^2s^2)$ for a RO-ABP-polynomial of size s in n variables. For a PRO-ABP-polynomial f_j with decomposition $(g_j, \{T_i(x_i)\}_{i\in[n]})$, $f_j \equiv 0 \Leftrightarrow g_j \equiv 0$. Hence we get the same time bound for $\mathrm{PIT}_{\text{PRO-ABP}}$. By Lemma 1 and using Lemma 3, $\mathrm{PIT}_{\text{RO-ABP}}$ can be implemented in the black-box setting to run in time $(dn)^{O(\log n)}$, for RO-ABP-polynomials in n variables which are d-decomposable. Theorems 2 and 3 are now proved using these observations and applying Theorem 6.

References

[Agr05] M. Agrawal. Proving lower bounds via pseudo-random generators. In *Proc. 25th Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 92–105, 2005.

[AHT07] M. Agrawal, T.M. Hoang, and T. Thierauf. The polynomially bounded perfect matching problem is in NC^2. In *Proc. 24th Annual Symposium on Theoretical Aspects of Computer Science*, pages 489–499, 2007.

[Alo99] N. Alon. Combinatorial nullstellensatz. *Combinatorics, Probability and Computing*, 8(1–2):7–29, 1999.

[HS80] J. Heintz and C.P. Schnorr. Testing polynomials which are easy to compute (extended abstract). In *Proc. 12th Annual ACM Symposium on the Theory of Computing*, pages 262–272, 1980.

[KI04] V. Kabanets and R. Impagliazzo. Derandomizing polynomial identity testing means proving circuit lower bounds. *Computational Complexity*, 13(1–2):1–44, 2004.

[KMSV09] Z.S. Karnin, P. Mukhppadhyay, A. Shpilka, and Ilya Volkovich. Deterministic identity testing of depth 4 multilinear circuits with bounded top fan-in. Technical Report TR09–116, Electronic Colloquium on Computational Complexity (ECCC), November 2009.

[Sax09] N. Saxena. Progress of polynomial identity testing. Technical Report ECCC TR09-101, Electronic Colloquium in Computational Complexity, 2009.

[Sch80] J.T. Schwartz. Fast probabilistic algorithms for polynomial identities. *J. Assn. Comp. Mach.*, 27:701–717, 1980.

[SV08] A. Shpilka and I. Volkovich. Read-once polynomial identity testing. In *Proceedings of the 40th Annual STOC*, pages 507–516, 2008.

[SV09] A. Shpilka and I. Volkovich. Improved polynomial identity testing of read-once formulas. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques, volume 5687 of LNCS*, pages 700–713, 2009.

[Val79] L. Valiant. Completeness classes in algebra. In *Proc. 11th Annual ACM Symposium on the Theory of Computing*, pages 249–261, 1979.

[Zip79] R. Zippel. Probabilistic algorithms for sparse polynomials. In *Proceedings of the International Symposium on Symbolic and Algebraic Manipulation (EUROSAM '79)*, volume 72 of *Lect. Notes in Comp. Sci.*, pages 216–226. Springer Verlag, 1979.

A Example : RO-ABPs Are Not Universal

Proposition 6. *The degree-2 elementary symmetric polynomial* $e_n(x_1, x_2, \ldots, x_n) = \sum_{1\leq i<j\leq n} x_i x_j$, $n \geq 3$ *can not be computed by a RO-ABP.*

Proof. For the purpose of contradiction, suppose that some RO-ABP A computes e_n. For any x_i denote the edge it labels by $g_i = (s_i, t_i)$. We can define an ordering $<$ among g_i's, by taking $g_i < g_j$ if and only if the polynomial computed by the subprogram $A(t_i, s_j)$ has a nonzero constant term. Due to the fact that A is a DAG, we have for any i, j, if $x_i < x_j$, then not $x_j < x_i$.

The fact that for every (i, j) pair, $x_i x_j$ appears as a term in e_n implies that for any $i \neq j$, we have one of $x_i < x_j$ or $x_j < x_i$. Incidently, note this implies the ordering is transitive. Namely, if $x_i < x_j$ and $x_j < x_k$, then s_j must be reachable from t_i, and s_k must be reachable from t_j in A, but then s_i can not be reachable from t_k. Hence not $x_k < x_j$, which implies $x_j < x_k$.

In any case, observe there is a permutation $\phi : [n] \to [n]$ for which $x_{\phi(1)} < x_{\phi(2)} < \cdots < x_{\phi(n)}$. This implies that $\prod_{i\in[n]} x_i$ appears as a term in the polynomial computed by A, which is a contradiction. □

B Separation of Σ_k-PRO-Formula and Σ_k-PRO-ABP

We give the sketch of an argument that shows, there exists a RO-ABP-polynomial f in n variables that can not be computed by a sum of k many PRO-formulas, if $k = o(n)$. We assume that $\mathbb{F}$ is algebraically closed to keep the algebraic geometry simple. First, we observe that since any such f is multilinear, it suffices to argue that f can not be computed by a sum of k many RO-formulas. Namely, for any RO-formula-polynomial $g(z_1, \ldots, z_n)$ that depend on z_i, for any univariate polynomials $T_1, \ldots, T_n \in \mathcal{T}_d$, if T_i has degree e, then the individual degree of x_i in $g(T_1(x_1), \ldots, T_n(x_n))$ is e.

We think of multilinear polynomials as points in $\mathbb{F}^{2^n}$, as determined by the coefficients of its 2^n many monomials. Let $K = \mathbb{F}^{\binom{n}{2}}$. Let $\pi : \mathbb{F}^{2^n} \to K$ be the projection given by restricting to coefficient of monomials in the set $\{x_i x_j : 1 \leq i < j \leq n\}$. Let $V \subseteq F^{2^n}$ be the set of points corresponding to all RO-ABP-polynomials, and let $W_k \subseteq F^{2^n}$ correspond to the set of all Σ_k-RO-formulas.

Consider the following generic RO-ABP with $2n + 2$ nodes $\{v_1, \ldots, v_n, v_{n+1} = t\} \cup \{s = u_0, u_1, \ldots u_n\}$. We do not worry about levelling the ABP. For all $i \in [n]$, there is an edge from v_i to u_i with variable label x_i. For all $0 \leq i < j \leq n+1$, the edge from u_i to v_{i+1} carries the constant label $c_{i,j}$, yet to be determined. Let f be the output of this ABP. Observe that for every $i < j$, the monomial $x_i x_j$ of f has coefficient $c_{0,i} c_{i,j} c_{j,n+1}$. Since all $c_{i,j}$'s can be set independently, we thus have that $\pi(V) = K$.

For a single RO-formula in n variables it is not too difficult to see that it always can be simplified to have $O(n)$ many gates. Namely, there is no need to pre-compute constants, as we can use any element of $\mathbb{F}$ as a label. This means that we have an enumeration $F_1, F_2, \ldots, F_m$, of RO-formulas each having $O(n)$ many "generic constants", such that any RO-Formula can be obtained from some F_i by specifying values in $\mathbb{F}$ for these generic constants. Let R be a bound on the number of constants gates used in any F_i. We have that m is some large finite number depending on n, which counts the number of structurally different RO-formulas in n variables with at most R constant gates. Say F_i has $r \leq R = O(n)$ generic constants $c_1, \ldots, c_r$. The coefficients of $x^{\vec{a}} := x_1^{a_1} \ldots x_n^{a_n}$ of the polynomial computed by F_i is given by some polynomial $p^i_{\vec{a}}(c_1, \ldots, c_r)$. Let p^i be the polynomial map $\mathbb{F}^r \to \mathbb{F}^{2^n}$ given by the 2^n-tuple of polynomials $p^i_{\vec{a}}$, for all $a \in \{0,1\}^n$. We conclude that $\bigcup_{i \in [m]} \text{Image}(p^i) = W_1$. Eventhough m is a large number, this is a finite union, and hence its dimension[6] is bounded by the maximum dimension of any Image(p^i), and hence is at most $O(n)$. Applying the projection π cannot increase the dimension, so we conclude that $\pi(W_1)$ has dimension $O(n)$. However, $\pi(V) = K$, which is $\binom{n}{2}$-dimensional. Hence there exists $v \in V \backslash W_1$. The latter means there exists some RO-ABP-polynomial that is not a RO-formula (This fact has already been demonstrated by giving an explicit example, of course). To obtain the argument for sum's of k many RO-formulas one argues similarly, but now use the fact that the number of constants in any sum of k RO-formulas after simplifications can be bounded by $O(kn)$. Then we obtain a finite enumeration of Σ_k-RO-formulas, each having $O(kn)$ generic constants. Similarly as before, for $k = o(n)$ one fails to cover the entire space K for dimensional reasons.

C Figure 1

Figure 1 shows an RO-ABP computing $x_1x_2 + x_2x_3 + x_{n-1}x_n$, when n is even. The case when n is odd is dealt with similarly. Unlabeled edges are labeled with 1.

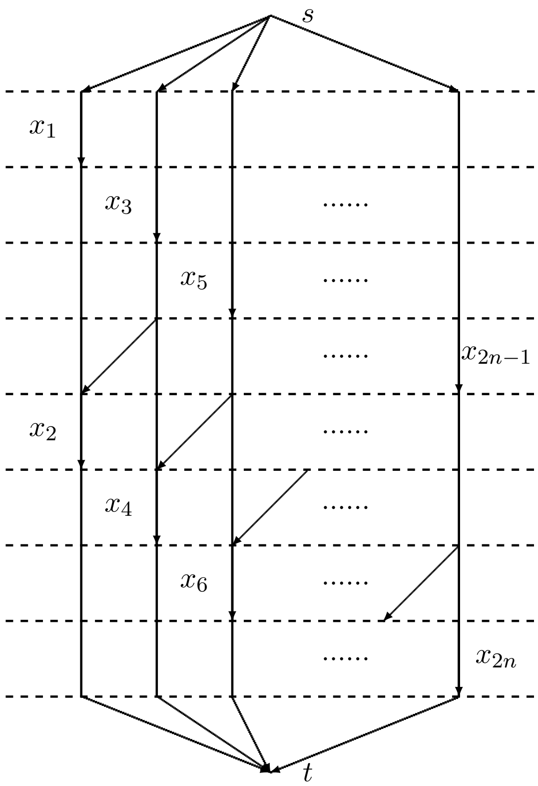

Figure 1: A RO-ABP computing $x_1x_2 + x_2x_3 + \ldots + x_{2n-1}x_{2n}$.

[6]To complete the argument formally, we would take the dimension of a a set $S \subseteq \mathbb{F}^q$, to mean the dimension of the algebraic set $\overline{S}$, where $\overline{S}$ denotes the closure of S in the Zariski topology.

ECCC ISSN 1433-8092
http://eccc.hpi-web.de