Electronic Colloquium on Computational Complexity, Revision 2 of Report No. 70 (2022)

On Solving Sparse Polynomial Factorization Related Problems

Pranav Bisht * Ilya Volkovich†

Abstract

In a recent result of Bhargava, Saraf and Volkovich [FOCS'18; JACM'20], the first factor sparsity bound for constant individual degree polynomials was shown. In particular, it was shown that any factor of a polynomial with at most s terms and individual degree bounded by d can itself have at most $s^{O(d^2 \log n)}$ terms. It is conjectured, though, that the "true" sparsity bound should be polynomial (i.e. $s^{\mathrm{poly}(d)}$). In this paper we provide supporting evidence for this conjecture by presenting polynomial-time algorithms for several problems that would be implied by a polynomial-size sparsity bound. In particular, we give efficient (deterministic) algorithms for identity testing of $\Sigma^{[2]}\Pi\Sigma\Pi^{[\mathsf{ind\text{-}deg}\ d]}$ circuits and testing if a sparse polynomial is an exact power. Hence, our algorithms rely on different techniques.

Keywords: Sparse Polynomials, Identity Testing, Derandomization, Factor-Sparsity, Multivariate Polynomial Factorization.

1 Introduction

Polynomial Factorization is one of the core problems in algebraic complexity: given a multivariate polynomial $f \in \mathbb{F}[x_1, x_2 \ldots, x_n]$ over a field $\mathbb{F}$, output all its irreducible factors. In addition to being a natural problem, its importance is highlighted by various applications such as: list decoding [Sud97, GS99], derandomization [KI04], cryptography [CR88] and others. In the seminal works of [Kal89, KT90], efficient randomized factorization algorithms were presented. Yet, coming up with an efficient *deterministic* factorization algorithm remains a long-standing open question.

Indeed, one aspect of the computational problem is the representation of the input polynomial. One natural way to represent a polynomial is by listing all its terms and coefficients. This is known as *dense* representation. Yet, even if the individual degree of every variable is bounded by a small constant d, the total number of terms can be exponentially large, reaching $(d+1)^n$. Nonetheless,

*Department of CSE, IIT Kanpur, India, pbisht@cse.iitk.ac.in

†Computer Science Department, Boston College, Chestnut Hill, MA. Email: ilya.volkovich@bc.edu

ISSN 1433-8092

in many applications [Zip79, GK85, BOT88, GJR10] the actual number of non-zero terms in a polynomial is much smaller - poly(n). Such polynomials are referred to as *sparse* polynomials, which will be the focus of our paper.

A key question that precedes the design of efficient factorization algorithms for sparse polynomial is whether a factor of a sparse polynomial is (itself) sparse. Indeed, this question was first studied by von zur Gathen and Kaltofen in [GK85] that gave a randomized factorization algorithm where the runtime depends on the number of terms in the output factors. In the same paper they provided an example inspired by geometric series (see below) of a family of polynomials that have factors with a super-polynomial (quasi-polynomial) number of terms. We denote by $\|f\|$ the *sparsity* of f. That is, the number of non-zero terms in f.

Example 1.1 ([GK85])**.** Let $n \geq 1$. Consider the polynomial $f(\overline{x}) = \prod_{i\in[n]} (x_i^n - 1)$ which can be written as a product of $g(\overline{x}) = \prod_{i\in[n]} (1 + x_i + \ldots + x_i^{n-1})$ and $h(\overline{x}) = \prod_{i\in[n]} (x_i - 1)$.
Observe that $\|f\| = \|h\| = 2^n$ while $\|g\| = n^n$, resulting in a quasi-polynomial blow-up[1].

Furthermore, for fields with finite characteristics the blow-up can be significantly larger:

Example 1.2 ([Vol15])**.** For a prime p, let $f \in \mathbb{F}_p[x_1, \ldots x_n]$, and let $0 < d < p$. Consider

$$f(\overline{x}) = (x_1 + x_2 + \ldots + x_n)^p = x_1^p + x_2^p + \ldots + x_n^p$$
$$g(\overline{x}) = (x_1 + x_2 + \ldots + x_n)^d$$

Notice that g is a factor of f, but $\|f\| = n$ and $\|g\| = \binom{n+d-1}{d} = n^{\Omega(d)}$.

Based on the above, we should first try to obtain a "sparsity-bound" on factors of sparse polynomials with constant (i.e. bounded) individual degree. More formally, for some fixed d, we require that $\deg_{x_i} \leq d$, for all variables x_i. The simplest case (when $d = 1$) corresponds to the so-called *multilinear* polynomials. In [SV10], it was shown that a factor of an s-sparse[2] multilinear polynomial is itself s-sparse. Subsequently, in [Vol17], this result was extended to the case of *multiquadratic* polynomials (i.e. when $d = 2$). In a recent work of [BSV20], a quasi-polynomial-size sparsity bound was given for *any* fixed d. Specifically, it was shown that a factor of an s-sparse polynomial with individual degree bounded by d is $s^{O(d^2 \log n)}$-sparse. In addition, [BSV20] designed a factorization algorithm whose runtime is efficient in terms of the sparsity bound. As a result they obtained a deterministic quasi-polynomial-time factorization algorithm for sparse polynomials with bounded individual degree. In the same paper it was also conjectured that the "true" sparsity bound should be polynomial rather than quasi-polynomial. More formally:

Conjecture 1.3. *There exists a universal constant $k \in \mathbb{N}$ such that for any $s, d \in \mathbb{N}$, any factor of an s-sparse polynomial with individual degree bounded by d has at most s^{d^k} terms.*

[1] Although g is not irreducible, this issue can be resolved using standard techniques. For example, by considering the product $f + yh = (g + y)h$ for a new variable y.

[2] A polynomial is s-sparse, it if contains at most s non-zero terms.

In this paper we provide supporting evidence for this conjecture by presenting deterministic polynomial-time algorithms for some problems that reduce to sparse polynomial factorization. It is to be noted that invoking the aforementioned factorization algorithm of [BSV20] with a polynomial-size sparsity bound would imply a (deterministic) *polynomial-time* algorithm for sparse polynomial factorization and hence polynomial-time algorithms for these problems. In the absence of a polynomial-size sparsity bound, we design our algorithms using new techniques.

1.1 Our Results

We will now describe our main results. In what follows, $\mathbb{F}$ is an arbitrary field (finite or otherwise).

1.1.1 Identity Testing for $\Sigma^{[2]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$ Circuits

The Polynomial Identity Testing (PIT) problem asks to decide whether a given input polynomial is identically zero. The input is usually given in the form of an algebraic circuit (see Appendix A for definition). The PIT algorithm is called *white-box* if one can look 'inside' the circuit. The algorithm is called *black-box* if the circuit is given via an oracle access, where one is only allowed to evaluate the polynomial on a chosen set of input points. PIT is one of the few natural problems which have a simple efficient randomized algorithm [DL78, Sch80, Zip79] but lack a deterministic one. Indeed, it has been a long standing open question to come up with an efficient deterministic algorithm for this problem. Our first result is an efficient (deterministic) identity testing algorithm for the class of $\Sigma^{[2]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$ circuits, where a $\Sigma^{[2]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$ circuit C of size s computes a polynomial of the form:

$$C = \prod_{i=1}^{r} g_i + \prod_{j=1}^{m} h_j$$

where each polynomial (g_i and h_j) is an s-sparse polynomial with individual degree at most d (for some fixed d). Note, though, that r and m, and hence the total degree of C, can be arbitrary (i.e. polynomially) large. In particular, the polynomial computed by C may not itself be sparse. This class generalizes the model considered in [Vol17], where $m = 1$ and the g_i-s are irreducible polynomials. For the formal definition of our circuit model and further discussion, see Section 4.1.

Observe that the identity testing problem for this circuit class reduces to polynomial factorization of sparse polynomials with bounded individual degree. Therefore, by invoking the factorization algorithm of [BSV20], we can get an algorithm whose runtime is efficient in terms of the sparsity bound. Plugging in the best bound of [BSV20] results in a quasi-polynomial-time algorithm. Our next result gives a *polynomial-time* algorithm for this model. In addition, our algorithm operates in the *black-box* setting, whereas the described factorization-based algorithm is a white-box algorithm.

Theorem 1. *There exists a deterministic algorithm that given a black-box access to a* $\Sigma^{[2]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$ *circuit* C *of size* s *determines if* $C \equiv 0$, *in time* $\text{poly}((sd)^{d^3}, n)$.

An important ingredient in our algorithm is a result that links the gcd of two polynomials, their subresultant and the resultant of their coprime parts - in the **multivariate** setting. See Section 3 for the formal definitions.

Theorem 2. *Let $A, B \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be two polynomials such that $A = f \cdot g$ and $B = h \cdot g$ and let x_i be a variable. Then*

$$S_{x_i}(d, A, B) = g \cdot \mathrm{Res}_{x_i}(f, h) \cdot \mathrm{lc}_{x_i}(g)^{m'+n'-1}$$

here $m = \deg_{x_i}(A)$, $n = \deg_{x_i}(B)$, $d = \deg_{x_i}(g)$, $m' = \deg_{x_i}(f) = m - d$ and $n' = \deg_{x_i}(h) = n - d$. In addition:

- $\mathrm{Res}_{x_i}(f, h)$ *is the resultant of f and h w.r.t the variable x_i.*
- $\mathrm{lc}_{x_i}(g)$ *is the leading coefficient of g when written as a polynomial in x_i*
- *And finally, $S_{x_i}(d, A, B)$ is the d-th subresultant of A and B.*

To put the result in context, consider two univariate polynomials $A, B \in \mathbb{F}[x]$. A classical result in the Theory of Resultants (see e.g. [GCL92, GG99, CLO15]) states that:

1. $\mathrm{Res}(A, B) \equiv 0$ if and only if $\gcd(A, B)$ is non-trivial.
2. For any polynomials $u, v \in \mathbb{F}[x]$ we have that $\mathrm{Res}(uA, vB) = \mathrm{Res}(A, B)$. Furthermore, if u and v are field elements (i.e. $u, v \in \mathbb{F}$) then $\mathrm{Res}(uA, vB) = \mathrm{Res}(A, B) \cdot u^{\deg B} \cdot v^{\deg A}$.
3. The *j-th Subresultant* $S(j, A, B) \equiv 0$ whenever $j < \deg(\gcd(A, B))$.
4. There exists a non-zero field element $\alpha \in \mathbb{F}$ such that $S(j, A, B) = \alpha \cdot \gcd(A, B)$, when $j = \deg(\gcd(A, B))$.

In the multivariate setting one can always regard multivariate polynomials as polynomials in a single variable with coefficients being rational functions in the remaining variables. Yet, in this case α is no longer a mere 'field element' as it can now be an arbitrary rational function in the remaining variables! From that perspective, our result can be seen as explicitly expressing α as a polynomial (and not even a rational function) in the remaining variables. We believe that this explicit relation could be of interest in its own right.

To illustrate the aforementioned problem and provide more intuition on our result, let us write the polynomials A and B in the statement of Theorem 2 as $A = (uf) \cdot (g/u)$ and $B = (uh) \cdot (g/u)$, where u is a rational function that **does not** depend on x_i. Observe that the introduction of u does not affect the degrees of x_i. We obtain the following invariant:

$$\begin{aligned} S_{x_i}(d, A, B) &= \frac{g}{u} \cdot \mathrm{Res}_{x_i}(uf, uh) \cdot \mathrm{lc}_{x_i}\left(\frac{g}{u}\right)^{m'+n'-1} = \frac{g}{u} \cdot \mathrm{Res}_{x_i}(f, h) \cdot u^{m'+n'} \cdot \frac{\mathrm{lc}_{x_i}(g)^{m'+n'-1}}{u^{m'+n'-1}} \\ &= g \cdot \mathrm{Res}_{x_i}(f, h) \cdot \mathrm{lc}_{x_i}(g)^{m'+n'-1}. \end{aligned}$$

1.1.2 Exact Powers

Our next result pertains to exact powers of polynomials. A polynomial $f \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ is an *exact power* if there exists (another) polynomial $g \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ and $e \in \mathbb{N}$ such that $f = g^e$. We note that despite the rich structure, the best known sparsity bound for exact roots (i.e. $\|g\|$ in terms of $\|f\|$) is the general sparsity bound of size $s^{O(d^2 \log n)}$ by [BSV20]. Hence, one can use the factorization algorithm of [BSV20] to test if a given sparse polynomial is an exact power, in quasi-polynomial time. Similarly, a polynomial-size sparsity bound, even for the case of exact roots, would imply a polynomial-time algorithm for exact-power testing problem. We provide a polynomial-time algorithm for exact-power testing that **does not** rely on this sparsity bound.

Theorem 3. *There is a deterministic algorithm that given a sparse polynomial $f \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ of individual degree d as an input, decides whether $f = g^e$ for some polynomial $g \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ and $e \in \mathbb{N}$, in time* $\mathrm{poly}(s^{d^2}, n)$.

We remark that the algorithm only performs exact-power testing and **does not** output a "witness" polynomial g. Indeed, a polynomial-time algorithm that actually outputs g would imply a polynomial-size sparsity bound on exact roots! In addition, the runtime of our algorithm is polynomial in the bit-complexity of the field elements since it does not rely on univariate polynomial factorization. For instance, for finite fields we get the runtime of $\mathrm{poly}(\log |\mathbb{F}|)$ vs $\mathrm{poly}(|\mathbb{F}|)$.

1.1.3 Improved Sparsity Bounds for Co-factors of Multilinear Polynomials

Given two polynomials $f, h \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ such that $f = gh$, g is called a *quotient polynomial* or a *co-factor* of h. We study the problem of multilinear co-factor sparsity: suppose f is s-sparse and h is multilinear. How sparse/dense can g be? We remark that any (even non-constructive) efficient upper bound on the sparsity of g allows us to compute g efficiently by interpolating the ratio f/h using a reconstruction algorithm for sparse polynomials (e.g. [KS01]) and verifying the result.

The motivation to study this problem is two-fold: first of all, by previous results (see e.g. [BSV20]) a multilinear factor of an s-sparse polynomial (of any degree) is itself s-sparse. This suggests more structure for multilinear co-factors we could potentially exploit. Second, a polynomial-size sparsity bound on multilinear co-factors g (even when the individual degree of g is $d = 2$) would imply a polynomial-size sparsity bound for (**all** factors of) polynomials with individual degree $d = 3$. We note that the multicubic ($d = 3$) case is the first instance where we do not have a polynomial-size factor-sparsity bound yet. Indeed, multilinear co-factors can be seen as the "bottle-neck" for this case. The formal argument is given in Section 1.4.

To state our next result we need the following technical definition. We say that a polynomial $h \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ has a *unique projection of length k* if there exist k variables $x_{i_1}, x_{i_2}, \ldots, x_{i_k}$ and k corresponding exponents $e_1, e_2, \ldots, e_k$ such that h has a unique monomial that contains the pattern $x_{i_1}^{e_1} x_{i_2}^{e_2} \cdot \ldots \cdot x_{i_k}^{e_k}$ (see Definition 6.5 for more details).

Theorem 4. *Let $f \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a polynomial of sparsity s and individual degree at most d such that $f = gh$. Suppose, in addition, that h is a multilinear polynomial with a unique projection of length k. Then the sparsity of g is bounded by $s^{O(dk)}$.*

We remark that Example 1.2 with $d = p - 1$ (resulting in a lower bound of $n^{\Omega(p)}$) showcases the tightness of our result as here f is n-sparse and $h = x_1 + \ldots + x_n$ has a unique projection of length 1 (e.g. x_1) which results in an upper bound of $n^{O(p)}$ for g.

We can also extend Theorem 4 to the case of a co-factor of a power of a multilinear polynomial. See Theorem 6.25 for the formal statement. Subsequently, we show that every multilinear s-sparse polynomial has a unique projection of length $O(\log s)$ (see Lemma 6.9). By plugging in this result into Theorem 4, we obtain a new sparsity bound of size $s^{O(d \log s)}$ for all multilinear co-factors.

Corollary 1.4. *Let $f \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a polynomial of sparsity s and individual degree at most d such that $f = gh$. Suppose, in addition, that h is a multilinear polynomial. Then the sparsity of g is bounded by $s^{O(d \log s)}$.*

The obtained bound is slightly better than the general sparsity bound of size $s^{O(d^2 \log n)}$ by [BSV20] when $s = \text{poly}(n)$. Although our overall improvement may seem incremental (e.g. it does not allow us to "get rid" of the $\log n$ in the exponent) our main contribution here is conceptual: identifying a combinatorial property - the length of the shortest unique projection - that governs the bound on the sparsity of multilinear co-factors.

1.2 Related works

For the sparse polynomial factorization problem, [BSV20] have shown that factors of an s-sparse polynomial of individual degree d, have their sparsity bounded by $s^{O(d^2 \log n)}$. Currently, this is the best known bound for factor-sparsity when $d \geq 3$. For restricted classes of symmetric polynomials, Bisht and Saxena [BS21] recently improved this bound to $s^{O(d^2 \log d)}$.

In [GK85], another problem was posed alongside the sparse factorization problem, in the hope that it might be easier. This problem is referred to as *testing sparse factorization*. Given $m + 1$ sparse polynomials $f, g_1, \ldots, g_m$, it asks to test whether $f = g_1 \cdot \ldots \cdot g_m$. The work of [SSS13] gives a polynomial-time algorithm for this problem, in the special case where every g_i is a sum of univariate polynomials. [Vol17] gives a polynomial-time algorithm when f (and therefore every g_i) has constant individual degree and each g_i is an irreducible polynomial. Our PIT result is connected to this problem. In Theorem 1, we give a polynomial-time algorithm to test whether $\prod_{i=1}^{r} f_i = \prod_{j=1}^{m} g_j$, where each f_i, g_j is a sparse polynomial with constant individual degree. Note that now LHS is also a product of polynomials. Moreover, there is no restriction placed on g_j-s except that they have bounded individual degree.

The depth-4 $\Sigma\Pi\Sigma\Pi$ circuit class (see Appendix A for definition) is extremely important in the context of the PIT problem, as it is known that a polynomial-time black-box PIT for this class implies a quasi-polynomial-time black-box PIT for general VP circuits [AV08, AGS19]. For a long

time, no PIT algorithm better than the trivial $d^{O(n)}$ time algorithm was known for this class, until the recent breakthrough result of Limaye et al. [LST22], which gives a sub-exponential time algorithm. Various restricted versions of depth-4 circuits are studied to get close to polynomial-time PIT algorithms. For example, Peleg and Shpilka [PS21] give a polynomial-time PIT algorithm for $\Sigma^{[3]}\Pi\Sigma\Pi^{[2]}$ circuits, where the top fan-in is 3 and the bottom fan-in is 2. Recently, Dutta et al. [DDS21] gave a quasi-polynomial-time PIT for $\Sigma^{[k]}\Pi\Sigma\Pi^{[d]}$ circuits, where the top fan-in k and bottom fan-in d are allowed to be any fixed constants. In this model, the restriction on bottom fan-in implies that the bottom $\Sigma\Pi$ computes polynomials of total degree at most d. We give polynomial-time PIT algorithm for $\Sigma^{[2]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$ model, where the top fan-in is 2 and the bottom $\Sigma\Pi$ computes polynomials with *individual* degree at most d. We note that the individual degree restriction is much weaker than the total degree restriction. Indeed, even for the case of individual degree bounded by 1 (i.e. multilinear polynomials) the total degree can still be $\Omega(n)$! [SV18] gave a polynomial-time PIT algorithm for the class of multilinear $\Sigma^{[k]}\Pi\Sigma\Pi$ circuits, with constant top fan-in k, where every gate in the circuit computes a multilinear polynomial. Yet, even a white-box polynomial-time PIT for *general* $\Sigma^{[2]}\Pi\Sigma\Pi$ circuits is still open.

Another related problem is that of *divisibility testing*, which gives two multivariate polynomials f and h and asks to decide whether h divides f. [For15] gives a quasi-polynomial-time algorithm when f is sparse and h is a quadratic polynomial (and hence also sparse). We note that the quadratic restriction on h is much stronger than a constant individual degree restriction, although there is no constant degree restriction for f here. [Vol17] gives a polynomial-time algorithm when both f, h are sparse and have constant individual degree. In the proof of Corollary 1.4, we solve a 'search' version of the divisibility testing problem, i.e. we actually compute f/h in quasi-polynomial time, when f is sparse with constant individual degree and h is a multilinear factor of f.

1.3 Our Techniques & Proof Ideas

1.3.1 Identity Testing for $\Sigma^{[2]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$ Circuits

Let $C = \prod_{i=1}^{r} g_i + \prod_{j=1}^{m} h_j$ where g_i-s and h_j-s are s-sparse polynomials in $\mathbb{F}[x_1, x_2, \ldots, x_n]$ of individual degree at most d. Clearly, if $C \equiv 0$ then it will evaluate to zero on any input. Now suppose $C \not\equiv 0$. Our goal is to find a point $\mathbf{a} \in \mathbb{F}^n$ such that $C(\mathbf{a}) \neq 0$. Our approach relies on the uniqueness of factorization property of the ring of multivariate polynomials. Specifically, we have that

$$\prod_{i=1}^{r} g_i \neq -\prod_{j=1}^{m} h_j$$

Consequently, wlog there exists an irreducible polynomial (factor) u and $\ell > 0$ such that u^ℓ divides the LHS but does not divide the RHS. Our goal is to preserve this "situation" while reducing the number of variables. Clearly, a random projection will be sufficient. However, we wish to obtain

a deterministic algorithm. To this end, we are looking for a projection that does not introduce new dependencies between factors. That is, for every i, j: if $v \mid g_i$ and $u \mid h_j$ satisfying $\gcd(u, v) = 1$ we need to ensure that $\gcd(u', v') = 1$, when u' and v' are the projections of u and v, respectively. The main tool for that is the *Resultant*. Indeed, one of the fundamental properties of the resultant is that

$$\mathrm{Res}(A, B) \not\equiv 0 \text{ if and only if } \gcd(A, B) = 1.$$

In the multivariate setting, this condition roughly translates into:

$$[\forall x_k : \mathrm{Res}_{x_k}(u, v) \not\equiv 0] \implies \gcd(u', v') = 1.$$

In other words, we need to hit all the resultants of the form $\mathrm{Res}_{x_k}(u, v)$ when $v \mid g_i$ and $u \mid h_j$. By definition, $\mathrm{Res}_{x_k}(u, v)$ is a determinant of $2d \times 2d$ matrix where each entry is a coefficient of u or v. Hence, $\mathrm{Res}_{x_k}(u, v)$ is $t^{O(d)}$-sparse polynomial with individual degree at most $O(d^2)$, where t is an upper bound on the sparsities of u and v. Consequently, we can use a hitting set generator for sparse polynomials (e.g. [KS01]) to hit the resultant. As u and v are factors of s-sparse polynomials of individual degree d, the best upper by [BSV20] will be $t = s^{O(d^2 \log s)}$. This will result in a quasi-polynomial-time algorithm.

Another idea would be to use the multiplicative properties of the resultant and hit $\mathrm{Res}_{x_k}(h_j, g_i)$ instead. Indeed, $\mathrm{Res}_{x_k}(h_j, g_i) \not\equiv 0 \implies \mathrm{Res}_{x_k}(u, v) \not\equiv 0$ and since g_i and h_j are s-sparse, $\mathrm{Res}_{x_k}(h_j, g_i)$ is $s^{O(d)}$-sparse and this would get a polynomial-time algorithm. The main issue is that we could have $\mathrm{Res}_{x_k}(u, v) \not\equiv 0$ while $\mathrm{Res}_{x_k}(h_j, g_i) \equiv 0$. For example, if $h_j = uf$ and $g_i = vf$ for the same polynomial f. Going back, one may ask whether we could show a better sparsity bound on $\mathrm{Res}_{x_k}(u, v)$. While we do not quite do that, we instead show that $\mathrm{Res}_{x_k}(u, v)$ is a *factor* of *some* $s^{O(d)}$-sparse polynomial of individual degree at most $O(d^2)$. As the ring of polynomials forms an integral domain, this allows us to use a polynomial-size hitting set generator for sparse polynomials.

To achieve the above goal, suppose for simplicity that $g_i = u^{a_1} \cdot v^{b_1}$ and $h_j = u^{a_2} \cdot v^{b_2}$, for some non-negative integers a_1, b_1, a_2, b_2. If all these numbers are strictly positive, we run into the same issue we have encountered earlier. That is, $\mathrm{Res}_{x_k}(u, v) \not\equiv 0$ while $\mathrm{Res}_{x_k}(h_j, g_i) \equiv 0$. To address that, we apply Theorem 2 (our key technical contribution) which allows us to "extract" the gcd. For example, if $g_i = uv^2$ and $h_j = u^2 v$, we can write $g_i = v \cdot uv$ and $h_j = u \cdot uv$ and obtain that $\mathrm{Res}_{x_k}(u, v)$ is a factor of $S_{x_k}(\deg_{x_k}(uv), g_i, h_j)$, which is an $s^{O(d)}$-sparse polynomial (see Observation 3.19). However, a sole gcd extraction may be insufficient. Consider the case when $g_i = uv^2$ and $h_j = uv$. Repeating the same argument will just yield a trivial statement that $\mathrm{Res}_{x_k}(v, 1) = 1$ is a factor of a sparse polynomial. To overcome this difficulty, we apply the previous argument on powers of g_i and h_j. That is, on $g_i^z = u^{za_1} \cdot v^{zb_1}$ and $h_j^t = u^{ta_2} \cdot v^{tb_2}$. The idea now would be to isolate the powers of u from the powers of v. Within the same example, consider $g_i^2 = u^2 v^4 = v \cdot u^2 v^3$ and $h_j^3 = u^3 v^3 = u \cdot u^2 v^3$. Now, by Theorem 2, $\mathrm{Res}_{x_k}(u, v)$ is a factor of $S_{x_k}(\deg_{x_k}(u^2 v^3), g_i^2, h_j^3)$. More generally, we show how to find appropriate "small" z and t using linear algebra.

Unfortunately, though, this could be made possible only when h_j and g_i satisfy certain "non-degeneracy" condition w.r.t u and v. More formally, when the matrix $E \triangleq \begin{bmatrix} a_1 & a_2 \\ b_1 & b_2 \end{bmatrix}$ has full rank (see Lemma 4.5). Our final crucial observation is that we can actually ignore "degenerate" pairs u, v. To this end, we prove a technical lemma (Lemma 2.3) which could be of independent interest.

1.3.2 Exact Power Testing

Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be an s-sparse polynomial of constant individual degree d. We show how to test whether $f = g^e$, for some other polynomial $g \in \mathbb{F}[x_1, \ldots, x_n]$ and some $e \in \mathbb{N}$. We utilize the notion of *reverse-monic* polynomials for this result. We call a polynomial h reverse-monic, if there exists some $i \in [n]$, such that $h|_{x_i=0} = 1$ (see Definition 6.1). If our input polynomial f is reverse-monic, we show that g is $s^{O(d)}$-sparse. Moreover, we also get an algorithm to compute this exact root g. We prove this in Lemmas 5.4 & 5.9 using a formal expansion that can be thought of as a generalization of the Binomial Expansion:

$$(1+x)^{\frac{1}{e}} = \sum_{i=0}^{\infty} \binom{\frac{1}{e}}{i} x^i.$$

In general though, our input polynomial f may not be reverse-monic. We first convert f into a reverse-monic polynomial $\hat{f}$ with respect to some variable x_i, using a known standard transformation (see Definition 5.10). This step only incurs a slight sparsity blow-up of s^d. One important property of this transformation is that it preserves the "exact power" structure. That is, if $f = g^e$, then $\hat{f} = h^e$, for some polynomial h. We then compute this e-th root of the reverse-monic $\hat{f}$, as mentioned previously.

However, we are still not quite done. It can happen that a polynomial f which was not an exact power, may become an exact power after the reverse-monic transformation. We need an additional condition to get the converse implication. We show that if both $\hat{f}$ and $f_{|_{x_i=0}}$ are exact powers, then we can correctly conclude that f is also an exact power (Claim 5.12). This gives us a recursive algorithm, as $f_{|_{x_i=0}}$ is a polynomial in $(n-1)$ variables. This procedure is described formally in Algorithm 3.

1.3.3 Co-Factor Sparsity Bound

For the co-factor bounds, our results build on the division elimination techniques of [Str73]. Let us outline our approach. To this end, let $f, h \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be s-sparse polynomials such that $h(0, \ldots, 0) = 1$ and suppose that $f = gh$ for some polynomial $g \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ with individual degree at most d. Consider the following formal expansion: $\frac{1}{(1-x)} = \sum_{j=0}^{\infty} x^j$. Then we have:

$$g = \frac{f}{h} = \frac{f}{(1-(1-h))} = \sum_{j=0}^{\infty} f(1-h)^j$$

when the equality is an equality of formal sums of monomials. The key observation is that $(1-h)$ does not contain any constants, hence total degree of every monomial in $(1-h)^j$ (and hence in every summand $f(1-h)^j$) is at least j. Consequently, we can "discard" the tail $\sum_{j=dn+1}^{\infty} f(1-h)^j$ since every monomial in g has a total degree of at most dn. Indeed, g will be formed by a subset of monomials of $\sum_{j=0}^{dn} f(1-h)^j$. This allows us to obtain an upper bound on the sparsity of g:

$$\|g\| \leq \sum_{j=0}^{dn} s^{j+1} \leq s^{dn+2}.$$

Clearly, the outlined approach has two major flaws:

1. It requires that $h(0,\ldots,0) = 1$ (or more generally, $h(0,\ldots,0) \neq 0$). And even then:
2. The obtained bound is exponential in n.

One way to address the former is by a random shift to the variable. However, this may significantly increase **both** the sparsity and the individual degree! We take a different approach. Our main observation is that the argument still works if we treat the polynomials as polynomials in "fewer" variables.

Formally, let $I \subseteq [n]$ of size $|I| = k$. We can regard the polynomials as polynomials in the variables x_I with coefficient in the remaining variables. In particular, suppose that $h|_{x_I=0_I} = 1$. In this case we say that h is *I-reverse monic*. Observe that every monomial in $(1-h)^j$ contains at least one variable from x_I. That is, the total x_I-degree of $(1-h)^j$ is at least j and hence (as before) we can discard the tail. Yet now, g depends "only" on k variables and thus its "total" degree is kd (and not nd). This way we obtain a better upper bound on the sparsity of g, if k is "small":

$$\|g\| \leq \sum_{j=0}^{kd} s^{j+1} \leq s^{kd+2}.$$

Of course, our approach still relies on the assumption that h is I-reverse monic for a "small" subset I. Although we are unable to lift this assumption, we can weaken it. As was noted earlier, if $h(0,\ldots,0) = \alpha \neq 0$ (i.e. when $I = [n]$) we can just divide by α as it is a field element. However, this is no longer possible for an arbitrary I (especially, if I is a small set). Yet, we observe that if $h|_{x_I=0_I} = \alpha$ and α is a non-zero *single monomial* (in the remaining variables) we can transform h into an I-reverse monic polynomial $\hat{h}$ with the exact same sparsity. The idea is to apply the transformation $x_i = \alpha \cdot x_i$ for all $i \in I$. Note that since α is a single monomial, this transformation is reversible. Indeed, there is an 1-1 correspondence between the monomials of h and $\hat{h}$. Given this connection, we refer to such h as I-reverse *pseudo*-monic.

Our final ingredient is (yet) another observation that for multilinear polynomials we can weaken the assumption that h is I-reverse *pseudo*-monic further by considering *unique projections*. That is,

monomials that have a "unique pattern". Formally, we want h to have exactly one monomial that contains the submonomial: $x_{i_1}^{e_1} x_{i_2}^{e_2} \cdots x_{i_k}^{e_k}$. We show that by "flipping" the variables in h we can transform it into another multilinear polynomial $\tilde{h}$ which is $\{i_1, i_2, \ldots, i_k\}$-reverse pseudo-monic. As a result, $\|g\| \leq s^{kd+2}$.

This is our main conceptual contribution: the upper bound on the sparsity of a multilinear co-factor g is governed by a *combinatorial property* of the set of monomials of h: the length of the shortest unique projection. As an application, we show that every s-sparse polynomial has a unique projection of length at most $\log s + 1$, thus we obtain a new, slightly stronger, sparsity bound on co-factors of multilinear polynomials.

1.4 Multilinear co-Factor Motivation

Theorem 4 and Corollary 1.4 in this paper apply to the factorization scenario of $f = gh$ where f is s-sparse and h is multilinear. First of all, note that by previous results (see [BSV20] and references within) h itself is s-sparse. So we are looking to bound the sparsity of g. As it turns out, this pattern is the "bottleneck" case for multicubic polynomials. In other words, showing a polynomial-size sparsity bound on g in this scenario would imply a polynomial-size sparsity bound on factors of general multicubic polynomials! In fact, it is sufficient to consider the case when the degree of g in every variable is exactly 2! We remark that getting polynomial-size sparsity bound is open for $d \geq 3$. The following lemma summarizes this formally.

Lemma 1.5. *Suppose there exists an absolute constant $a \geq 1$ such that for any multicubic polynomial f: if $g \mid f$ **and** f/g is multilinear then $\|g\| \leq \|f\|^a$. Then for any multicubic polynomial f if $g \mid f$ then $\|g\| \leq \|f\|^a$.*

Proof. We prove the following claim: Let $f \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a multicubic polynomial such that $f = uv$ and $v \not\equiv 0$. Then $\|u\| \leq \|f\|^a$. Note that the claim also covers the case when $f = u \equiv 0$. The proof is by induction on n (the number of variables in f). The base case is when $n = 0$ (i.e. $u, f, v \in \mathbb{F}$) where the claim follows trivially. Suppose $n \geq 1$. We have the following cases to consider:

- There exists a variable x_i s.t. $\deg_{x_i}(u) \geq 1$ but $\deg_{x_i}(v) = 0$. Let $1 \leq d \leq 3$ be the degree of x_i in u. In this case we can write:

  $$(u_d x_i^d + \ldots + u_0)v = uv = f = f_d {x_i}^d + \ldots + f_0.$$

  Here, u_j, f_j and v do not depend on x_i. Formally: $f_j = u_j v$ for $j \in \{0, \ldots, d\}$. By the induction hypothesis, we have that $\|u_j\| \leq \|f_j\|^a$ for $j \in \{0, \ldots, d\}$ and hence:

  $$\|u\| = \sum_{j=0}^{d} \|u_j\| \leq \sum_{j=0}^{d} \|f_j\|^a \leq \left(\sum_{j=0}^{d} \|f_j\| \right)^a = \|f\|^a.$$

- There exists a variable x_i s.t. $\deg_{x_i}(v) \geq 1$, but $\deg_{x_i}(u) = 0$. Pick $\alpha \in \mathbb{F}$ such that $v|_{x_i=\alpha} \not\equiv 0$. We have that: $$u \cdot v|_{x_i=\alpha} = u|_{x_i=\alpha} \cdot v|_{x_i=\alpha} = f|_{x_i=\alpha}.$$ By the induction hypothesis: $\|u\| \leq \|f|_{x_i=\alpha}\|^a \leq \|f\|^a$.

- There exists a variable x_i s.t. $\deg_{x_i}(u) = 1$. Wlog $\deg_{x_i}(v) \geq 1$. We can write $$(u_1 x_i + u_0)(v_d x_i^d + \ldots + v_e x_i^e) = uv = f = (f_{d+1} x_i^{d+1} + \ldots + f_e x_i^e).$$ Here, $d > e$ and $v_d, v_e \not\equiv 0$. In particular, we have that $u_1 v_d = f_{d+1}$ and $u_0 v_e = f_e$. By the induction hypothesis: $\|u\| = \|u_1\| + \|u_0\| \leq \|f_{d+1}\|^a + \|f_e\|^a \leq \|f\|^a$.

- WLOG we are left with the case that for each $i \in [n]$ we have that: $\deg_{x_i}(u) = 2$ and $\deg_{x_i}(v) = 1$. Based on our assumption, in this case $\|u\| \leq \|f\|^a$ and we are done. □

2 Preliminaries

Notations: We use the shorthand $[n]$ for the set $\{1, 2, \ldots, n\}$. We denote a vector $v = (v_1, \ldots, v_n)$ in short by $\mathbf{v}$ (as a column vector). We denote the n-fold Cartesian product of a set H by H^n. The set of non-negative real numbers is denoted by $\mathbb{R}_{\geq 0}$ and $\mathbb{R}^n_{\geq 0}$ denotes the space of n-dimensional non-negative real vectors. We will use $\log x$ for $\log_2 x$. We use the $\stackrel{\Delta}{=}$ symbol for definition.

Let $f \in \mathbb{F}[\mathbf{x}] = \mathbb{F}[x_1, x_2, \ldots, x_n]$ be an n-variate polynomial. The *individual degree* of a variable x_i in f, denoted by $\deg_{x_i}(f)$, is defined as the maximum degree of that variable in f, while the *individual degree* of f is the maximum among all the individual degrees, $\max_{i \in [n]} \deg_{x_i}(f)$. We will use $\mathbf{x}^\mathbf{e}$ to denote the monomial $x_1^{e_1} x_2^{e_2} \cdots x_n^{e_n}$. We define $\text{coeff}(\mathbf{x}^\mathbf{e})(f)$ as the coefficient of monomial $\mathbf{x}^\mathbf{e}$ in polynomial f. We define *support* of f as $\text{supp}(f) \stackrel{\Delta}{=} \{\mathbf{e} \mid \text{coeff}(\mathbf{x}^\mathbf{e})(f) \neq 0\}$. We define the *sparsity* of f as the number of non-zero terms in f. Let us denote *sparsity* of f as $\|f\|$, which is the same as $|\text{supp}(f)|$. We say that f *depends* on a variable x_i if there exist $\mathbf{a}, \mathbf{b} \in \mathbb{F}^n$ which differ only in the i-th coordinate such that $f(\mathbf{a}) \neq f(\mathbf{b})$. We define the set $\text{var}(f) \stackrel{\Delta}{=} \{i \mid f \text{ depends on } x_i\}$. For $i \in [n]$, we denote by $\text{lc}_{x_i}(f)$ the *leading coefficient* of f when written as a polynomial in x_i. Formally, let $f = \sum_{j=0}^{d} f_j \cdot x_i^j$ such that $\forall j$, f_j is a polynomial in rest of the variables and $f_d \not\equiv 0$. Then $\text{lc}_{x_i}(f) \stackrel{\Delta}{=} f_d$. Observe that if $f = g \cdot h$ then $\forall i \in [n] : \text{lc}_{x_i}(f) = \text{lc}_{x_i}(g) \cdot \text{lc}_{x_i}(h)$.

For a set $I \subseteq [n]$, we use x_I to denote the set of variables $\{x_i \mid i \in I\}$ and $x_{[n]\setminus I}$ to denote the set of remaining variables. We use the symbol $f|_{x_I=0_I}$ to denote the polynomial resulting from substituting 0 at all the x_I variables in f. Let $\mathbf{a} = (a_1, \ldots, a_n) \in \mathbb{F}^n$. For $i \in [n]$ we define a partial assignment $\mathbf{a}_{-i} \stackrel{\Delta}{=} (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$ and $f(x_i, \mathbf{a}_{-i}) \stackrel{\Delta}{=} f(a_1, \ldots, a_{i-1}, x_i, a_{i+1}, \ldots, a_n)$.

We use $f \equiv 0$ to denote that f is identically zero. For two polynomials $f, g \in \mathbb{F}[x_1, x_2, \ldots, x_n]$, we use the symbol $\gcd(f, g)$ to denote their greatest common divisor. Let $f \sim g$ denote equality of f and g up to multiplication by non-zero a field constant, i.e. $f = c \cdot g$, for some $c \neq 0 \in \mathbb{F}$. Using these notations, we can formulate the following known result.

Lemma 2.1 (See e.g. [BSV20] and references within). *Let $f, h \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ where h is a multilinear polynomial and $h \mid f$. Then $\|h\| \leq \|f\|$.*

2.1 Vectors and Interlacing

In this section we will state and prove some useful properties of vectors that may be of independent interest.

Definition 2.2 (Interlacing). *Let $a, b, c, d \in \mathbb{R}$. We say that two pair of points (a, b) and (c, d)* interlace *if $(a - c)(b - d) < 0$. Equivalently, $(a - c)$ and $(b - d)$ have opposite signs.*

Remark: In particular, interlacing implies that at least two of a, b, c, d are non-zero.
The following result relates the interlacing property with linear dependence.

Lemma 2.3. *Let $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n_{\geq 0}$ and $\mathbf{c}, \mathbf{d} \in \mathbb{R}^m_{\geq 0}$ such that $\left(\sum_{i=1}^n a_i, \sum_{i=1}^n b_i\right)$ interlaces with $\left(\sum_{j=1}^m c_j, \sum_{j=1}^m d_j\right)$. Then there exists $i \in [n], j \in [m]$ such that the matrix $\begin{bmatrix} a_i & c_j \\ b_i & d_j \end{bmatrix}$ has full rank.*

Proof. Without loss of generality, suppose $\sum_{i=1}^n a_i > \sum_{j=1}^m c_j$ and $\sum_{i=1}^n b_i < \sum_{j=1}^m d_j$. In particular, there exists $s \in [n], t \in [m]$ such that $a_s, d_t > 0$. Consider the $2 \times (n + m)$ matrix,

$$E \triangleq \begin{bmatrix} a_1 & \cdots & a_n & c_1 & \cdots & c_m \\ b_1 & \cdots & b_n & d_1 & \cdots & d_m \end{bmatrix}.$$

Suppose E is not full-rank. Since $\mathbf{a}$ and $\mathbf{d}$ are non-zero vectors, E is not the zero matrix and hence it must have rank 1. In that case, the first row is linearly dependent on the second row. Since all E's entries are non-negative, there exist $\alpha, \beta > 0$ such that $\alpha \cdot \mathbf{a} = \beta \cdot \mathbf{b}$ and $\alpha \cdot \mathbf{c} = \beta \cdot \mathbf{d}$. This implies:

$$\alpha \cdot \left(\sum_{i=1}^n a_i\right) = \beta \cdot \left(\sum_{i=1}^n b_i\right), \quad \alpha \cdot \left(\sum_{j=1}^m c_j\right) = \beta \cdot \left(\sum_{j=1}^m d_j\right).$$

Which in turn implies:

$$\beta \cdot \left(\sum_{i=1}^n b_i\right) = \alpha \cdot \left(\sum_{i=1}^n a_i\right) > \alpha \cdot \left(\sum_{j=1}^m c_j\right) = \beta \cdot \left(\sum_{j=1}^m d_j\right) \implies \sum_{i=1}^n b_i > \sum_{j=1}^m d_j$$

This contradicts the interlacing property. Hence, E must be full-rank. Let E' be a 2×2 rank-2 minor of E. If E' is of the form $E' = \begin{bmatrix} a_i & c_j \\ b_i & d_j \end{bmatrix}$ for some i, j we are done. Otherwise, suppose if E' is of the form $E' = \begin{bmatrix} a_i & a_j \\ b_i & b_j \end{bmatrix}$ or $E' = \begin{bmatrix} c_i & c_j \\ d_i & d_j \end{bmatrix}$ then by the exchange property, we can exchange one of the columns in E' with non-zero columns $\begin{bmatrix} c_t \\ d_t \end{bmatrix}$ or $\begin{bmatrix} a_s \\ b_s \end{bmatrix}$, respectively, to get a rank-2 minor of the required form. □

3 GCD, Resultants and Subresultants

In this section we discuss the classical tool of resultants and subresultants. In Theorem 3.9 we prove an interesting connection which is a somewhat more general version of Theorem 2. This result will be crucial in the design of our PIT algorithm and could be of interest in its own right.

The polynomial ring $\mathbb{F}[x_1, \ldots, x_n]$ is a unique factorization domain (UFD). Hence, the gcd of two polynomials is well defined up to a multiplication by field element. We can also define gcd with respect to a single variable x_i, where we treat rest of the variables as field elements. That is, $\gcd_{x_i}(f, g)$ is well defined up to multiplication by a rational function depending on the remaining variables. By convention, we choose the normalized gcd whenever we consider $\gcd_{x_i}(f, g)$ in this work. For example, let $f = x_1^2 x_2 + x_1 x_2^2$ and $g = x_1^2 x_2^2$, then $\gcd(f, g) = x_1 x_2$ while $\gcd_{x_2}(f, g) = x_2$. Technically, the former is gcd in $\mathbb{F}[x_1, x_2]$ and the latter is normalized gcd in $\mathbb{F}(x_1)[x_2]$.

We now consider a somewhat more general scenario. Let $A(y), B(y) \in \mathcal{R}[y]$ be two non-zero polynomials of y-degree d and e, respectively in an arbitrary UFD $\mathcal{R}$. Suppose $A(y) = \sum_{i=0}^{d} a_i \cdot y^i$ and $B(y) = \sum_{j=0}^{e} b_j \cdot y^j$. Consider the $(d+e) \times (d+e)$ *Sylvester matrix* M whose first e rows are the e shifts of the row vector $(a_d, \ldots, a_0, 0, \ldots, 0)$ and next d rows are the d shifts of the row vector $(b_e, \ldots, b_0, 0, \ldots, 0)$.

$$M = \begin{bmatrix} a_d & a_{d-1} & \ldots & a_1 & a_0 & & & \\ & a_d & a_{d-1} & \ldots & a_1 & a_0 & & \\ & & \ldots & \ldots & \ldots & \ldots & & \\ & & & a_d & a_{d-1} & \ldots & a_1 & a_0 \\ b_e & b_{e-1} & \ldots & b_1 & b_0 & & & \\ & b_e & b_{e-1} & \ldots & b_1 & b_0 & & \\ & & \ldots & \ldots & \ldots & \ldots & & \\ & & & b_e & b_{e-1} & \ldots & b_1 & b_0 \end{bmatrix}.$$

Definition 3.1 (Resultant)**.** *The resultant* $\mathrm{Res}_y(A, B) \in \mathcal{R}$ *is defined to be the determinant of this Sylvester matrix. That is,* $\mathrm{Res}_y(A, B) = \det(M)$.

In our setting, $\mathcal{R}$ will be a polynomial ring, say $\mathbb{F}[x_1, \ldots, x_n]$ and $\mathrm{Res}_y(A, B)$ will be a polynomial free of y-variable. The following is the most fundamental property of the Resultant:

Lemma 3.2 (See e.g. [GCL92, GG99, CLO15])**.** *Let* $A, B \in \mathbb{F}[y, x_1, \ldots, x_n]$ *be two polynomials. Then* $\gcd_y(A, B) \neq 1$ *if and only if* $\mathrm{Res}_y(A, B) \equiv 0$*. That is, A and B have a non-trivial factor that depends on the variable y (*$\deg_y(\gcd(A, B)) > 0$*) iff the Resultant of A, B w.r.t. y is the identically zero polynomial.*

We state the following important connection between projection of resultant and resultant of projections.

Lemma 3.3. *Let $f, g \in \mathbb{F}[y, \mathbf{x}]$ be two polynomials and let $\mathbf{a} \in \mathbb{F}^n$. Then,*

$$\mathrm{Res}_y(f,g)(\mathbf{a}) \neq 0 \implies \mathrm{Res}_y(f(\mathbf{a}), g(\mathbf{a})) \neq 0.$$

Proof. Let $d \triangleq \deg_y(f)$, $e \triangleq \deg_y(g)$, $r \triangleq \deg_y(f(\mathbf{a}))$ and $t \triangleq \deg_y(g(\mathbf{a}))$. Then with some easy determinant calculations, one can show that:

$$\mathrm{Res}_y(f,g)(\mathbf{a}) = \begin{cases} \mathrm{Res}_y(f(\mathbf{a}), g(\mathbf{a})) & r = d, t = e \\ \left(\mathrm{lc}_y(f)(\mathbf{a})\right)^{e-t} \cdot \mathrm{Res}_y(f(\mathbf{a}), g(\mathbf{a})) & r = d, t < e \\ (-1)^{e(d-r)} \cdot \left(\mathrm{lc}_y(g)(\mathbf{a})\right)^{d-r} \cdot \mathrm{Res}_y(f(\mathbf{a}), g(\mathbf{a})) & r < d, t = e \\ 0 & r < d, t < e \end{cases}$$

Note that if $\mathrm{Res}_y(f,g)(\mathbf{a}) \neq 0$, then $\mathrm{Res}_y(f(\mathbf{a}), g(\mathbf{a}))$ divides it and hence the conclusion follows. □

Lemma 3.2 and Lemma 3.3 together give us the following useful Corollary:

Corollary 3.4. *Let $f(y, \mathbf{x})$ and $g(y, \mathbf{x})$ be two polynomials in $\mathbb{F}[y, \mathbf{x}]$. Let $\mathbf{a} \in \mathbb{F}^n$. Then,*

$$\mathrm{Res}_y(f,g)(\mathbf{a}) \neq 0 \implies \gcd_y(f(y, \mathbf{a}), g(y, \mathbf{a})) = 1.$$

We also require multiplicative property of the Resultant that essentially follows from the definition:

Lemma 3.5. *Let $A, B, u, v \in \mathbb{F}[y, x_1, \ldots, x_n]$ be polynomials. Then* $\mathrm{Res}_y(A,B) \mid \mathrm{Res}_y(uA, vB)$.

We now study few useful sub-matrices of the Sylvester matrix below.

Definition 3.6 (j-th principal resultant)**.** *Let M_j be the submatrix of M formed by deleting last j rows of A terms, last j rows of B terms and the last $2j$ columns. We call M_j to be the j-th principal resultant of A and B. Note that* $\mathrm{Res}_y(A,B) = M = M_0$.

We can now define the subresultant polynomial as follows.

Definition 3.7 (Subresultant)**.** *Let M_{ij} be the $(d + e - 2j) \times (d + e - 2j)$ submatrix of Sylvester matrix M formed by deleting:*

- *rows $e - j + 1$ to e (each having coefficients of $A(y)$),*
- *rows $d + e - j + 1$ to $d + e$ (each having coefficients of $B(y)$),*
- *columns $d + e - 2j$ to $d + e$, except for column $d + e - i - j$.*

Note that the j-th principal resultant M_j is exactly M_{jj}.
For $0 \leq j \leq e$, the j-th subresultant of $A(y), B(y) \in \mathcal{R}[y]$ is the polynomial in $\mathcal{R}[y]$ of degree j defined by

$$S_y(j, A, B) = \det(M_{0j}) + \det(M_{1j}) \cdot y + \ldots + \det(M_{jj}) \cdot y^j.$$

We state below known results in the theory of subresultants, which will be useful for us.

Lemma 3.8 (Lem 7.1 of [GCL92]). *Let $A(x), B(x) \in \mathcal{R}[x]$ be two polynomials over an arbitrary UFD $\mathcal{R}$. Let $\mathcal{K}$ be the field of fractions of $\mathcal{R}$. Suppose*

$$A(x) = Q(x) \cdot B(x) + R(x),$$

for some polynomials $Q, R \in \mathcal{K}[x]$ such that $\deg_x(A) = m$, $\deg_x(B) = n$, $\deg_x(Q) = m - n$, $\deg_x(R) = k$ and $m \geq n > k$. Let b and r denote the leading coefficients of $B(x)$ and $R(x)$ respectively. Then

$$S_x(j, A, B) = (-1)^{(m-j)(n-j)} \times \begin{cases} b^{m-k} \cdot S_x(j, B, R) & 0 \leq j < k \\ b^{m-k} \cdot r^{n-k-1} \cdot R(x) & j = k \\ 0 & k < j < n-1 \\ b^{m-n+1} \cdot R(x) & j = n-1. \end{cases}$$

That is, $S_x(j, A, B)$ equals to one of the above four expressions multiplied by the corresponding sign $(-1)^{(m-j)(n-j)}$.

We now prove our main technical result which links the gcd of two polynomials, their subresultant and the resultant of their coprime parts. Theorem 2 is a special case of this result which, we believe, could be interesting in its own right.

Theorem 3.9. *Let $A(x), B(x) \in \mathcal{R}[x]$ be two polynomials over an arbitrary UFD $\mathcal{R}$. Suppose $A(x) = f(x) \cdot g(x)$ and $B(x) = h(x) \cdot g(x)$ with $\deg_x(A) = m$, $\deg_x(B) = n$, $\deg_x(g) = d$, $\deg_x(f) = m' = m - d$ and $\deg_x(h) = n' = n - d$. Then*

$$S_x(d, A, B) = g \cdot \mathrm{Res}_x(f, h) \cdot \mathrm{lc}_x(g)^{m'+n'-1}$$

Proof. Let $\mathcal{K}$ be the field of fractions of UFD $\mathcal{R}$. Consider Euclidean division of A by B in $\mathcal{K}[x]$ so that we get $A(x) = Q(x) \cdot B(x) + R(x)$, for some polynomials $Q, R \in \mathcal{K}[x]$ such that $\deg_x(R) < \deg_x(B)$. Note that since g divides both A and B, it must also divide R. Therefore, $R = g \cdot p$ for some polynomial $p(x) \in \mathcal{K}[x]$. Thus, we also get

$$f(x) = Q(x) \cdot h(x) + p(x) \tag{3.10}$$

Let $\deg_x(R) = k$ for some $k < n$ and let $\deg_x(p) = k' = k - d$. Now, we prove the theorem by induction on $\deg_x(p)$.

Base case: $\deg_x(p) = k' = 0$. In other words, $\deg_x(R) = k = d$. Thus using second case of Lemma 3.8, we get that:

$$S_x(d, A, B) = (-1)^{(m-d)(n-d)} \cdot b^{m-k} \cdot r^{n-k-1} \cdot R$$

$$
\begin{aligned}
&= (-1)^{m'.n'} \cdot \mathrm{lc}_x(h)^{m-k} \cdot \mathrm{lc}_x(g)^{m-k} \cdot \mathrm{lc}_x(p)^{n-k-1} \cdot \mathrm{lc}_x(g)^{n-k-1} \cdot pg \\
&= (-1)^{m'.n'} \cdot g \cdot \mathrm{lc}_x(h)^{m-k} \cdot \mathrm{lc}_x(p)^{n-k} \cdot \mathrm{lc}_x(g)^{m+n-2k-1} \\
S_x(d, A, B) &= (-1)^{m'.n'} \cdot g \cdot \mathrm{lc}_x(h)^{m-k} \cdot \mathrm{lc}_x(p)^{n-k} \cdot \mathrm{lc}_x(g)^{m'+n'-1} \qquad (3.11)
\end{aligned}
$$

The second last step above follows because $p = \mathrm{lc}_x(p)$ when $\deg_x(p) = 0$. Now, we shall compute $\mathrm{Res}_x(f, h)$. Note that $\mathrm{Res}_x(f, h) = S_x(0, f, h)$ by definition of subresultant. Considering (3.10) with $\deg_x(p) = 0$, we can use second case of Lemma 3.8 to get:

$$
\begin{aligned}
S_x(0, f, h) &= (-1)^{(\deg_x(f)-0).(\deg_x(h)-0)} \cdot \mathrm{lc}_x(h)^{\deg_x(f)-\deg_x(p)} \cdot \mathrm{lc}_x(p)^{\deg_x(h)-\deg_x(p)-1} \cdot p \\
&= (-1)^{m'.n'} \cdot \mathrm{lc}_x(h)^{m'} \cdot \mathrm{lc}_x(p)^{n'-1} \cdot p \quad [\text{as } \deg_x(p) = 0] \\
&= (-1)^{m'.n'} \cdot \mathrm{lc}_x(h)^{m'} \cdot \mathrm{lc}_x(p)^{n'} \quad [\text{as } p = \mathrm{lc}_x(p)] \\
\mathrm{Res}_x(f, h) &= (-1)^{m'.n'} \cdot \mathrm{lc}_x(h)^{m-k} \cdot \mathrm{lc}_x(p)^{n-k} \qquad (3.12)
\end{aligned}
$$

(3.11) and (3.12) together yield $S_x(d, A, B) = g \cdot \mathrm{Res}_x(f, h) \cdot \mathrm{lc}_x(g)^{m'+n'-1}$ for the base case.

Induction step: Now, we assume $\deg_x(p) = k' > 1$. In other words, $\deg_x(R) = k > d$. Therefore, by first case of Lemma 3.8:

$$
\begin{aligned}
S_x(d, A, B) &= (-1)^{(m-d)(n-d)} \cdot b^{m-k} \cdot S_x(d, B, R) \\
&= (-1)^{m'.n'} \cdot \mathrm{lc}_x(h)^{m-k} \cdot \mathrm{lc}_x(g)^{m-k} \cdot S_x(d, B, R) \qquad (3.13)
\end{aligned}
$$

Now consider Euclidean division of B by R in $\mathcal{K}[x]$ to get

$$
B(x) = Q'(x) \cdot R(x) + R'(x) \qquad (3.14)
$$

for some polynomial $R'(x) \in \mathcal{K}[x]$ with $\deg_x(R') < \deg_x(R)$. Since g divides both B and R, we deduce that g must also divide R'. Let $R' = g \cdot p'$ for some polynomial $p' \in \mathcal{K}[x]$. Thus from (3.14), we also get

$$
h(x) = Q'(x) \cdot p(x) + p'(x) \qquad (3.15)
$$

In (3.14) since $\deg_x(R') < \deg_x(R)$ or equivalently $\deg_x(p') < \deg_x(p)$, we can use induction hypothesis to deduce that,

$$
S_x(d, B, R) = g \cdot \mathrm{Res}_x(h, p) \cdot \mathrm{lc}_x(g)^{n'+k'-1} \qquad (3.16)
$$

Note that $\deg_x(p) = k' > 0$ in induction step, thus we can use first case of Lemma 3.8 on (3.10) to get

$$
\begin{aligned}
\mathrm{Res}_x(f, h) &= S_x(0, f, h) \\
&= (-1)^{(\deg_x(f)-0)(\deg_x(h)-0)} \cdot \mathrm{lc}_x(h)^{\deg_x(f)-\deg_x(p)} \cdot S_x(0, h, p) \\
&= (-1)^{m'.n'} \cdot \mathrm{lc}_x(h)^{m'-k'} \cdot \mathrm{Res}_x(h, p)
\end{aligned}
$$

$$\mathrm{Res}_x(h,p) = \frac{\mathrm{Res}_x(f,h)}{(-1)^{m'.n'} \cdot \mathrm{lc}_x(h)^{m'-k'}}. \tag{3.17}$$

Substituting (3.17) in (3.16), we get:

$$S_x(d,B,R) = g \cdot \frac{\mathrm{Res}_x(f,h)}{(-1)^{m'.n'} \cdot \mathrm{lc}_x(h)^{m'-k'}} \cdot \mathrm{lc}_x(g)^{n'+k'-1} \tag{3.18}$$

Substituting (3.18) back into (3.13), we get

$$\begin{aligned} S_x(d,A,B) &= (-1)^{m'.n'} \cdot \mathrm{lc}_x(h)^{m-k} \cdot \mathrm{lc}_x(g)^{m-k} \cdot g \cdot \frac{\mathrm{Res}_x(f,h)}{(-1)^{m'.n'} \cdot \mathrm{lc}_x(h)^{m'-k'}} \cdot \mathrm{lc}_x(g)^{n'+k'-1} \\ &= \mathrm{lc}_x(g)^{m-k} \cdot g \cdot \mathrm{Res}_x(f,h) \cdot \mathrm{lc}_x(g)^{n'+k'-1} \quad [\text{as } m-k = m'-k'] \\ &= g \cdot \mathrm{Res}_x(f,h) \cdot \mathrm{lc}_x(g)^{m-k+n'+k'-1} \\ &= g \cdot \mathrm{Res}_x(f,h) \cdot \mathrm{lc}_x(g)^{m'+n'-1} \quad [\text{as } m-k+k' = m-d = m'] \end{aligned}$$

This completes the proof of induction step, as well as that of the theorem. □

We conclude this section making an important observation that any subresultant (and hence the Resultant) of two sparse polynomials of individual degree at most d is a sum of at most $d+1$ determinants of $2d \times 2d$ matrices where each entry is a coefficient of a sparse polynomial and, hence is itself a (somewhat) sparse polynomial of a small individual degree.

Observation 3.19. *Let $A, B \in \mathbb{F}[y, x_1, \ldots, x_n]$ be two s-sparse polynomials with individual degrees at most d. Then for any j, $S_y(j, A, B)$ is an $(2ds)^{2d+1}$-sparse polynomial with individual degrees at most $2d^2$.*

4 PIT for $\Sigma^{[2]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$ Circuits

In this section we prove our main result Theorem 1. We refer the reader to Appendix A for the formal definition of an algebraic circuit and PIT algorithm. For the purpose of black-box PIT algorithm, we require the notion of hitting set generators (HSG) or simply generator in short.

Definition 4.1 (Generator)**.** *Let $\mathcal{C}$ be a class of n-variate polynomials. Consider $\mathcal{G} = (\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_n) : \mathbb{F}^k \to \mathbb{F}^n$, an n-tuple of k-variate polynomials where for each $i \in [n]$, $\mathcal{G}_i \in \mathbb{F}[t_1, t_2, \ldots, t_k]$. Let $f(x_1, \ldots, x_n)$ be an n-variate polynomial. We define action of $\mathcal{G}$ on polynomial f by $f(\mathcal{G}) = f(\mathcal{G}_1, \ldots, \mathcal{G}_n) \in \mathbb{F}[t_1, \ldots, t_k]$. We call $\mathcal{G}$ a k-seeded generator for class $\mathcal{C}$ if for every non-zero $f \in \mathcal{C}$, $f(\mathcal{G}) \not\equiv 0$. Degree of generator $\mathcal{G}$ is defined as $\deg(\mathcal{G}) \stackrel{\Delta}{=} \max\{\deg(\mathcal{G}_i)\}_{i=1}^n$.*

For a polynomial-time PIT algorithm, k is kept constant. A generator $\mathcal{G}$ acts as a variable reduction map which converts an input polynomial $f \in \mathbb{F}[x_1, \ldots, x_n]$ to $f(\mathcal{G}) \in \mathbb{F}[t_1, \ldots, t_k]$ such that $f \equiv 0$ if and only if $f(\mathcal{G}) \equiv 0$. Let D be the degree of $\mathcal{G}$ and d be the individual degree of f. Then $\mathcal{G}$ gives us a brute-force hitting-set of size $(ndD)^k$ (Lemma A.2). In other words, we get a polynomial-time black-box PIT algorithm for f when k is constant, $\mathcal{G}$ can be designed in polynomial time and its degree is also polynomially bounded.

4.1 The $\Sigma^{[k]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$ Model

A size s, depth-4 $\Sigma\Pi\Sigma\Pi$ circuit computes a polynomial of the form $f = \sum_{i=1}^{k}\prod_{j=1}^{m_i} f_{ij}$, where f_{ij} are s-sparse polynomials for each $i \in [k], j \in [m_i]$. For $s = \text{poly}(n)$, [LST22] gives the first deterministic sub-exponential time PIT for constant-depth (depth-4 also) which runs in $(sn)^{O(n^\mu)}$-time, where $\mu > 0$ is any real number. While a polynomial-time PIT algorithm for general depth-4 circuit continues to be elusive, various restricted versions of this model have been attacked. One such restriction is to make the top fan-in k constant. For $k = 2$, even white-box PIT for $\Sigma^{[2]}\Pi\Sigma\Pi$ circuits is still open. A more restricted model is the class of $\Sigma^{[k]}\Pi\Sigma\Pi^{[d]}$ circuits, where the top fan-in k and the bottom fan-in d are constants. For a size-s circuit of this class, f_{ij}'s are s-sparse polynomials of constant total degree at most d. Even this restricted model seems to be quite non-trivial. Only very recently, [DDS21] gave a quasi-polynomial-time black-box PIT algorithm for this model. For $k = 3$ and $d = 2$ (f_{ij}'s are quadratic polynomials), [PS21] give a polynomial-time black-box PIT algorithm. For $k = 3$ and $d > 2$, coming up with a polynomial PIT algorithm remains an open question.

We now introduce, what we call the $\Sigma^{[k]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$ model. In the $\Sigma^{[k]}\Pi\Sigma\Pi^{[d]}$ model, the sparse polynomials f_{ij}'s have constant total degree $\leq d$. We relax this restriction to f_{ij}'s being constant *individual* degree $\leq d$ polynomials in $\Sigma^{[k]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$ model. This is a more general model, since f_{ij}'s can now have much higher total degree, like $\Omega(n)$. In Section 4.2, we give a deterministic polynomial-time black-box PIT algorithm for this model when $k = 2$ and d is any constant. We also note that our PIT algorithm works for any field $\mathbb{F}$, while the works of [PS21, DDS21] do have certain field restrictions.

4.2 The PIT Algorithm

For a polynomial f and an irreducible polynomial u, let $\mathrm{e}_u(f)$ denote the highest power of u in f. In other words, $f = u^{\mathrm{e}_u(f)} \cdot g$, such that $u \nmid g$. If $u \nmid f$, then $\mathrm{e}_u(f) = 0$. We define a polynomial Φ with respect to two non-zero polynomials P, Q as follows:

Definition 4.2. *Let $P, Q \in \mathbb{F}[x_1, \ldots, x_n]$ be two non-zero polynomials. Define a non-zero polynomial $\Phi_{P,Q} \in \mathbb{F}[x_1, \ldots, x_n]$ as:*

$$\Phi_{P,Q} \stackrel{\Delta}{=} \prod_{\substack{u,v \mid PQ \\ i\in[n]}} \mathrm{Res}_{x_i}(u,v) \cdot \prod_{i\in[n]} \mathrm{lc}_{x_i}(P) \cdot \prod_{i\in[n]} \mathrm{lc}_{x_i}(Q),$$

where u, v are irreducible factors of P or Q such that $(\mathrm{e}_u(P), \mathrm{e}_v(P))$ interlaces with $(\mathrm{e}_u(Q), \mathrm{e}_v(Q))$. Moreover, we only consider non-zero multiplicands.

The next Lemma shows that a non-zero of Φ preserves non-similarity of polynomials.

Lemma 4.3. *Let $P, Q \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be two polynomials such that $P \nsim Q$ and let $\mathbf{a} \in \mathbb{F}^n$ such that $\Phi_{P,Q}(\mathbf{a}) \neq 0$. Then, there exists an $i \in [n]$ such that $P(x_i, \mathbf{a}_{-i}) \nsim Q(x_i, \mathbf{a}_{-i})$.*

Proof. By our premise, we have $P \nsim Q$. By uniqueness of factorization, without loss of generality, there exists an irreducible factor u of P, appearing with higher power in P than in Q. That is, $\mathrm{e}_u(P) > \mathrm{e}_u(Q)$. Let $k = \mathrm{e}_u(P)$ and $\ell = \mathrm{e}_u(Q)$. We have $k > \ell \geq 0$. We get $\ell = 0$, when u does not divide Q. Let $P = u^k \cdot G$ and $Q = u^\ell \cdot H$, for some polynomials G, H such that u does not divide either of them. Define set $T \stackrel{\Delta}{=} \{v \mid v \text{ is an irreducible factor of } H \text{ and } \mathrm{e}_v(P) \geq \mathrm{e}_v(Q)\}$. Then let

$$P = u^k \cdot \left(\prod_{v \in T} v^{\mathrm{e}_v(P)}\right) \cdot G'$$
$$Q = u^\ell \cdot \left(\prod_{v \in T} v^{\mathrm{e}_v(Q)}\right) \cdot H',$$

where G', H' are the product of remaining polynomials from G and H respectively.

We choose i to be any element in $\mathrm{var}(u)$, i.e. u depends on x_i. Note that $\mathrm{lc}_{x_i}(P) \not\equiv 0$, since u is a factor of P which depends on x_i. Since lc is multiplicative, we get that $\mathrm{lc}_{x_i}(P(x_i, \mathbf{a}_{-i})) = \mathrm{lc}_{x_i}(u(x_i, \mathbf{a}_{-i}))^k \cdot \mathrm{lc}_{x_i}(G(x_i, \mathbf{a}_{-i}))$. From our premise, we also know that $\Phi_{P,Q}(\mathbf{a}) \neq 0$. Then by the definition of $\Phi_{P,Q}$, we get that $\mathrm{lc}_{x_i}(P(x_i, \mathbf{a}_{-i})) \neq 0$, which implies that $\mathrm{lc}_{x_i}(u(x_i, \mathbf{a}_{-i})) \neq 0$. Together with the fact that u has x_i-degree at least one, we conclude that $u(x_i, \mathbf{a}_{-i})$ also has x_i-degree at least one. Suppose for the sake of contradiction that $P(x_i, \mathbf{a}_{-i}) \sim Q(x_i, \mathbf{a}_{-i})$. Then, we get

$$u(x_i, \mathbf{a}_{-i})^k \cdot \left(\prod_{v \in T} v^{\mathrm{e}_v(P)}(x_i, \mathbf{a}_{-i})\right) \cdot G'(x_i, \mathbf{a}_{-i}) \sim u(x_i, \mathbf{a}_{-i})^\ell \cdot \left(\prod_{v \in T} v^{\mathrm{e}_v(Q)}(x_i, \mathbf{a}_{-i})\right) \cdot H'(x_i, \mathbf{a}_{-i})$$
$$\implies u(x_i, \mathbf{a}_{-i})^{k-\ell} \cdot \left(\prod_{v \in T} v^{\mathrm{e}_v(P) - \mathrm{e}_v(Q)}(x_i, \mathbf{a}_{-i})\right) \cdot G'(x_i, \mathbf{a}_{-i}) \sim H'(x_i, \mathbf{a}_{-i}).$$

Since $k > \ell$ and $\forall v \in T : \mathrm{e}_v(P) \geq \mathrm{e}_v(Q)$, LHS is a proper polynomial in the above equation. Moreover, $u(x_i, \mathbf{a}_{-i})$ divides LHS. Now since LHS $\sim H'(x_i, \mathbf{a}_{-i})$ and $u(x_i, \mathbf{a}_{-i})$ depends on x_i, we deduce that $H'(x_i, \mathbf{a}_{-i})$ also depends on x_i. By uniqueness of factorization, we also deduce that $u(x_i, \mathbf{a}_{-i})$ divides $H'(x_i, \mathbf{a}_{-i})$.

Let $H' = v_1^{e_1} \cdot \ldots \cdot v_m^{e_m}$ be the irreducible factorization of H', for some $m \geq 1$ and $e_j \geq 1$ for all $j \in [m]$, where each v_j is irreducible. Here $e_j = \mathrm{e}_{v_j}(Q)$ for each $j \in [m]$. Then $H'(x_i, \mathbf{a}_{-i}) = v_1(x_i, \mathbf{a}_{-i})^{e_1} \cdot \ldots \cdot v_m(x_i, \mathbf{a}_{-i})^{e_m}$, where $v_j(x_i, \mathbf{a}_{-i})$'s may not be irreducible anymore due to substitution. Recall that u does not divide H and hence it does not divide H' either. Since u is irreducible, we get that $\gcd_{x_i}(u, v_j) = 1$, for all $j \in [m]$. At the same time, recall that $u(x_i, \mathbf{a}_{-i})$ divides $H'(x_i, \mathbf{a}_{-i})$. Since $H'(x_i, \mathbf{a}_{-i})$ depends on x_i, this implies that $u(x_i, \mathbf{a}_{-i})$ shares a non-trivial factor with some $v_j(x_i, \mathbf{a}_{-i})$ which depends on x_i. Thus, there exists some $j \in [m]$ such that $\gcd_{x_i}(u(x_i, \mathbf{a}_{-i}), v_j(x_i, \mathbf{a}_{-i})) \neq 1$. By definition of H', $v_j \notin T$ and hence $\mathrm{e}_{v_j}(P) < \mathrm{e}_{v_j}(Q) = e_j$ for all $j \in [m]$. Recall that $\mathrm{e}_u(P) = k > \ell = \mathrm{e}_u(Q)$. Hence $(\mathrm{e}_u(P), \mathrm{e}_{v_j}(P))$ interlaces with $(\mathrm{e}_u(Q), \mathrm{e}_{v_j}(Q))$. By our premise, $\Phi_{P,Q}(\mathbf{a}) \neq 0$. Then by the definition of $\Phi_{P,Q}$, we get that $\mathrm{Res}_{x_i}(u, v_j)(\mathbf{a}_{-i}) \neq 0$. By further applying Corollary 3.4, we deduce that $\gcd_{x_i}(u(x_i, \mathbf{a}_{-i}), v_j(x_i, \mathbf{a}_{-i})) = 1$, which gives us a contradiction. Hence, $P(x_i, \mathbf{a}_{-i}) \nsim Q(x_i, \mathbf{a}_{-i})$. □

The following is a technical lemma that will be used subsequently.

Lemma 4.4. *Let $u, v \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be two coprime and irreducible polynomials such that $\mathrm{var}(u) \cap \mathrm{var}(v)$ is non-empty. And suppose we have two polynomials $g = u^{a_1} \cdot v^{b_1}$ and $h = u^{a_2} \cdot v^{b_2}$, for some non-negative integers a_1, b_1, a_2, b_2. Define $z \triangleq a_2 + b_2$ and $t \triangleq a_1 + b_1$. For any $i \in \mathrm{var}(u) \cap \mathrm{var}(v)$, let $W \triangleq \gcd_{x_i}(g^z, h^t)$. Finally, let E be the following matrix:*

$$E \triangleq \begin{bmatrix} a_1 & a_2 \\ b_1 & b_2 \end{bmatrix}.$$

Then

$$\frac{g^z}{W} = \begin{cases} u^{\det(E)} & \text{if } \det(E) \geq 0 \\ v^{-\det(E)} & \text{otherwise.} \end{cases} \qquad \frac{h^t}{W} = \begin{cases} v^{\det(E)} & \text{if } \det(E) \geq 0 \\ u^{-\det(E)} & \text{otherwise.} \end{cases}$$

Proof. We have that:

$$g^z = (u^{a_1} \cdot v^{b_1})^z = u^{a_1a_2 + a_1b_2} \cdot v^{a_2b_1 + b_1b_2}$$
$$h^t = (u^{a_2} \cdot v^{b_2})^t = u^{a_1a_2 + a_2b_1} \cdot v^{a_1b_2 + b_1b_2}$$

If $\det(E) \geq 0$, then $a_1b_2 \geq a_2b_1$ and consequently $W = u^{a_1a_2 + a_2b_1} \cdot v^{a_2b_1 + b_1b_2}$. In that case, $g^z/W = u^{a_1b_2 - a_2b_1} = u^{\det(E)}$ and $h^t/W = v^{a_1b_2 - a_2b_1} = v^{\det(E)}$. Otherwise, if $\det(E) < 0$, then $a_2b_1 > a_1b_2$ and consequently $W = u^{a_1a_2 + a_1b_2} \cdot v^{a_1b_2 + b_1b_2}$. Then $g^z/W = v^{a_2b_1 - a_1b_2} = v^{-\det(E)}$ and $h^t/W = u^{a_2b_1 - a_1b_2} = u^{-\det(E)}$. □

We now show that under certain non-degeneracy condition, a resultant of two factors of sparse polynomials is itself a factor of a (somewhat) sparse polynomial.

Lemma 4.5. *Let $u \nsim v \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be two irreducible polynomials. Suppose there exist s-sparse, individual degree-d polynomials $g, h \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ such that the matrix*

$$E \triangleq \begin{bmatrix} \mathrm{e}_u(g) & \mathrm{e}_u(h) \\ \mathrm{e}_v(g) & \mathrm{e}_v(h) \end{bmatrix}$$

has full rank. Then for any $i \in [n] : \mathrm{Res}_{x_i}(u, v)$ is a factor of a non-zero $(sd)^{\mathcal{O}(d^3)}$-sparse, $\mathcal{O}(d^4)$-individual degree polynomial.

Proof. Consider any $i \notin \mathrm{var}(u) \cup \mathrm{var}(v)$. Then $\mathrm{Res}_{x_i}(u, v)$ is defined to be 1, which is trivially a factor of any sparse polynomial. Now consider any $i \in \mathrm{var}(u) \setminus \mathrm{var}(v)$. Then by definition, $\mathrm{Res}_{x_i}(u, v) = v^{\deg_{x_i}(u)}$. Note that both $e_v(g)$ and $e_v(h)$ cannot be zero, as E has full rank. Therefore, v is factor of g or h, which are both s-sparse. Similarly, u is also a factor of g or h which implies $\deg_{x_i}(u) \leq d$. We deduce that $\mathrm{Res}_{x_i}(u, v)$ is a factor of an s^d-sparse polynomial. Similarly, we get the same conclusion for any $i \in \mathrm{var}(v) \setminus \mathrm{var}(u)$. We are now left with $i \in \mathrm{var}(u) \cap \mathrm{var}(v)$, for which we shall prove below.

Let us write $g = u^{\mathrm{e}_u(g)} \cdot v^{\mathrm{e}_v(g)} \cdot A$ and $h = u^{\mathrm{e}_u(h)} \cdot v^{\mathrm{e}_v(h)} \cdot B$, for some polynomials $A, B \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ co-prime to both u and v. Let $g' = u^{\mathrm{e}_u(g)} \cdot v^{\mathrm{e}_v(g)}$ and $h' = u^{\mathrm{e}_u(h)} \cdot v^{\mathrm{e}_v(h)}$. Further, let $z = \mathrm{e}_u(h) + \mathrm{e}_v(h)$ and $t = \mathrm{e}_u(g) + \mathrm{e}_v(g)$. Consider polynomials $g^z = (g')^z \cdot A^z$ and $h^t = (h')^t \cdot B^t$. Since both g, h have individual degree d, we know that $\mathrm{e}_u(g), \mathrm{e}_v(g), \mathrm{e}_u(h), \mathrm{e}_v(h) \leq d$ and hence $s, t \leq 2d$. Pick any $i \in \mathrm{var}(u) \cap \mathrm{var}(v)$ and consider $\gcd_{x_i}(g^z, h^t)$. Define $W \triangleq \gcd_{x_i}((g')^z, (h')^t)$ and $Y \triangleq \gcd_{x_i}(A^z, B^t)$. Since g', h' are co-prime to both A and B, we deduce that

$$\gcd_{x_i}(g^z, h^t) = W \cdot Y.$$

By our premise, we have $\det(E) \neq 0$. Without loss of generality, let us assume $\det(E) > 0$. The other case follows similarly. Using Lemma 4.4, we get that $(g')^z / W = u^{\det(E)}$ and $(h')^t / W = v^{\det(E)}$. Therefore, we can write

$$g^z = W \cdot Y \cdot u^{\det(E)} \cdot \frac{A^z}{Y},$$
$$h^t = W \cdot Y \cdot v^{\det(E)} \cdot \frac{B^t}{Y}.$$

Note that A^z / Y and B^t / Y are proper polynomials by definition of Y. Let $\ell = \deg_{x_i}(\gcd(g^z, h^t))$. By Theorem 3.9, there exists $k \geq 0$ such that:

$$S_{x_i}(\ell, g^z, h^t) = W \cdot Y \cdot \mathrm{Res}_{x_i}\left(u^{\det(E)} \cdot \frac{A^z}{Y},\ v^{\det(E)} \cdot \frac{B^t}{Y}\right) \cdot \mathrm{lc}_{x_i}(W \cdot Y)^k.$$

Since both g^z and h^t are s^{2d}-sparse with individual degree at most $2d^2$, by Observation 3.19, $S_{x_i}(\ell, g^z, h^t)$ is $(sd)^{\mathcal{O}(d^3)}$-sparse with individual degree at most $8d^4$. Furthermore, observe that by definition:

$$\gcd_{x_i}\left(u^{\det(E)} \cdot \frac{A^z}{Y},\ v^{\det(E)} \cdot \frac{B^t}{Y}\right) = 1 \implies \mathrm{Res}_{x_i}\left(u^{\det(E)} \cdot \frac{A^z}{Y},\ v^{\det(E)} \cdot \frac{B^t}{Y}\right) \not\equiv 0 \implies S_{x_i}(\ell, g^z, h^t) \not\equiv 0.$$

Finally, since $\det(E)$ is a positive integer, we use Lemma 3.5 to deduce that

$$\mathrm{Res}_{x_i}(u, v) \mid \mathrm{Res}_{x_i}\left(u^{\det(E)} \cdot \frac{A^z}{Y},\ v^{\det(E)} \cdot \frac{B^t}{Y}\right).$$

Hence, we conclude that $\mathrm{Res}_{x_i}(u, v)$ is a factor of a non-zero $(sd)^{\mathcal{O}(d^3)}$-sparse sub-resultant polynomial. □

Using the above, we conclude that while the multiplicands in the polynomial $\Phi_{P,Q}$ may not themselves be sparse, they are factors of (some) sparse polynomials. Consequently, $\Phi_{P,Q}$ can be hit by a hitting set generator for sparse polynomials.

Lemma 4.6. *Let both $P, Q \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be products of s-sparse, individual degree-d polynomials and let $\mathcal{G}$ be a generator for $(sd)^{\mathcal{O}(d^3)}$-sparse, $\mathcal{O}(d^4)$-individual degree polynomials. Then $\Phi_{P,Q}(\mathcal{G}) \not\equiv 0$.*

Proof. Let $P = \prod_{j\in[r]} g_j$ and $Q = \prod_{k\in[m]} h_k$, where g_j, h_k are s-sparse polynomials of individual degree d, for all $j \in [r], k \in [m]$. By definition, $\Phi_{P,Q}$ has two types of multiplicands. We will show that $\mathcal{G}$ hits both types.

For the first type, let u, v be any irreducible factors of P or Q such that $(\mathrm{e}_u(P), \mathrm{e}_v(P))$ interlaces with $(\mathrm{e}_u(Q), \mathrm{e}_v(Q))$. We wish to show that $\mathrm{Res}_{x_i}(u, v) \not\equiv 0 \implies \mathrm{Res}_{x_i}(u, v)(\mathcal{G}) \not\equiv 0$. Observe that,

$$\mathrm{e}_u(P) = \sum_{j=1}^{r} \mathrm{e}_u(g_j) \qquad \mathrm{e}_u(Q) = \sum_{k=1}^{m} \mathrm{e}_u(h_k)$$
$$\mathrm{e}_v(P) = \sum_{j=1}^{r} \mathrm{e}_v(g_j) \qquad \mathrm{e}_v(Q) = \sum_{k=1}^{m} \mathrm{e}_v(h_k).$$

By definition, each of these e-values is a non-negative integer. Therefore by Lemma 2.3, there exists $j \in [r], k \in [m]$ such that $E \triangleq \begin{bmatrix} \mathrm{e}_u(g_j) & \mathrm{e}_u(h_k) \\ \mathrm{e}_v(g_j) & \mathrm{e}_v(h_k) \end{bmatrix}$ has full rank. Then by Lemma 4.5, for any $i \in [n]$: $\mathrm{Res}_{x_i}(u, v)$ is factor of some $(sd)^{\mathcal{O}(d^3)}$-sparse, $\mathcal{O}(d^4)$-individual degree polynomial. Since, $\mathcal{G}$ is a generator for such polynomials, we deduce that $\mathrm{Res}_{x_i}(u, v)(\mathcal{G}) \not\equiv 0$.

For the second type, by multiplicative property of lc, we know that for any $i \in [n]$,

$$\mathrm{lc}_{x_i}(P) = \prod_{j\in[r]} \mathrm{lc}_{x_i}(g_j) \quad \text{and} \quad \mathrm{lc}_{x_i}(P)(\mathcal{G}) = \prod_{j\in[r]} \mathrm{lc}_{x_i}(g_j)(\mathcal{G}).$$

Note that $\mathrm{lc}_{x_i}(g_j)$ is also s-sparse with individual degree d. Hence, $\mathrm{lc}_{x_i}(g_j)(\mathcal{G}) \not\equiv 0$, for all $j \in [r], i \in [n]$. This implies $\mathrm{lc}_{x_i}(P)(\mathcal{G}) \not\equiv 0$, for all $i \in [n]$ (whenever $\mathrm{lc}_{x_i}(P) \not\equiv 0$). Similarly, we can show $\mathrm{lc}_{x_i}(Q)(\mathcal{G}) \not\equiv 0$, for all $i \in [n]$. We conclude that $\Phi_{P,Q}(\mathcal{G}) \not\equiv 0$. □

For a generator $\mathcal{G} = (\mathcal{G}_1, \ldots, \mathcal{G}_n)$, we define $\mathcal{G}_{-i} \triangleq (\mathcal{G}_1, \ldots, \mathcal{G}_{i-1}, \mathcal{G}_{i+1}, \ldots, \mathcal{G}_n)$. For a polynomial $f \in \mathbb{F}[x_1, x_2, \ldots, x_n]$, we define $f(x_i, \mathcal{G}_{-i}) \triangleq f(\mathcal{G}_1, \ldots, \mathcal{G}_{i-1}, x_i, \mathcal{G}_{i+1}, \ldots, \mathcal{G}_n)$. By combining the result with Lemma 4.3 we obtain the following:

Corollary 4.7. *Let $P, Q \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be products of s-sparse, individual degree-d polynomials such that $P \nsim Q$. Let $\mathcal{G} = (\mathcal{G}_1, \ldots, \mathcal{G}_n) : \mathbb{F}^t \to \mathbb{F}^n$ be a generator for $(sd)^{\mathcal{O}(d^3)}$-sparse, $\mathcal{O}(d^4)$-individual degree polynomials. Then there exists an $i \in [n]$ such that $P(x_i, \mathcal{G}_{-i}) \nsim Q(x_i, \mathcal{G}_{-i})$.*

Proof. By Lemma 4.6, we get that $\Phi_{P,Q}(\mathcal{G}) \not\equiv 0$. From Lemma A.2, we know that there exists a set $W \subseteq \mathbb{F}$ of large enough size such that $\mathcal{G}(W^t) \subseteq \mathbb{F}^n$ is a hitting set for $\Phi_{P,Q}$. In particular, there exists $\mathbf{b} \in W^t$ such that for $\mathbf{a} \triangleq \mathcal{G}(\mathbf{b})$, we have $\Phi_{P,Q}(\mathbf{a}) \neq 0$. By Lemma 4.3, there exists an $i \in [n]$ such that $P(x_i, \mathbf{a}_{-i}) \nsim Q(x_i, \mathbf{a}_{-i})$. Now suppose $P(x_i, \mathcal{G}_{-i}) \sim Q(x_i, \mathcal{G}_{-i})$. Then have that

$$P\big(\mathcal{G}_1(\mathbf{b}), \ldots, \mathcal{G}_{i-1}(\mathbf{b}), x_i, \mathcal{G}_{i+1}(\mathbf{b}), \ldots, \mathcal{G}_n(\mathbf{b})\big) \sim Q\big(\mathcal{G}_1(\mathbf{b}), \ldots, \mathcal{G}_{i-1}(\mathbf{b}), x_i, \mathcal{G}_{i+1}(\mathbf{b}), \ldots, \mathcal{G}_n(\mathbf{b})\big).$$

This implies that $P(x_i, \mathbf{a}_{-i}) \sim Q(x_i, \mathbf{a}_{-i})$, which is a contradiction. Hence, $P(x_i, \mathcal{G}_{-i}) \nsim Q(x_i, \mathcal{G}_{-i})$. □

Finally, we can prove the main result of the section. Theorem 1 follows from this Lemma.

Lemma 4.8. *There exists a deterministic algorithm that given n, d, s and a black-box access to a $\Sigma^{[2]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$ circuit of size s determines if $C \equiv 0$, in time $\text{poly}((sd)^{d^3}, n)$. Algorithm 1 provides the outline.*

Algorithm 1: Black-box PIT algorithm for class $\Sigma^{[2]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$

Input: A black-box access to a polynomial $f(x_1, \ldots, x_n)$ computed by $\Sigma^{[2]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$ circuit

Output: "ZERO", if f is identically zero and "NON-ZERO", otherwise.

Call Lemma A.3 to get generator $\mathcal{G}$ of seed-length 1 for n-variate polynomials of sparsity $\leq (sd)^{\mathcal{O}(d^3)}$ and individual degree $\leq \mathcal{O}(d^4)$.

for $i \leftarrow 1$ *to* n **do**

  Compute the bivariate polynomial $f(x_i, \mathcal{G}_{-i})$.

  Call Lemma A.1 to do brute-force black-box PIT for $f(x_i, \mathcal{G}_{-i})$.

  if $f(x_i, \mathcal{G}_{-i}) \not\equiv 0$ **then return** "NON-ZERO".

end

return "ZERO".

Proof. We now analyze the correctness and runtime complexity of Algorithm 1.

Correctness: Note that $f \equiv 0 \implies f(x_i, \mathcal{G}_{-i}) \equiv 0$ trivially, for all $i \in [n]$. Thus, the algorithm outputs "ZERO" in this case, as desired. Now suppose $f \not\equiv 0$. Let $f = P + Q$, where both P, Q are product of s-sparse, individual degree d polynomials. If $P \nsim Q$, then Corollary 4.7 implies that there exists an $i \in [n]$ such that $P(x_i, \mathcal{G}_{-i}) \nsim Q(x_i, \mathcal{G}_{-i})$. In particular, $P(x_i, \mathcal{G}_{-i}) \neq -Q(x_i, \mathcal{G}_{-i})$. Since $f(x_i, \mathcal{G}_{-i}) = P(x_i, \mathcal{G}_{-i}) + Q(x_i, \mathcal{G}_{-i})$, we deduce that $f(x_i, \mathcal{G}_{-i}) \not\equiv 0$ in this case. Now suppose $f \not\equiv 0$ but $P \sim Q$. Let $P = cQ$, for some $c \in \mathbb{F}$. Then $f = (c+1)Q$, where $c \neq -1$ and $Q \not\equiv 0$. This means that there exists an $i \in [n]$ such that $\text{lc}_{x_i}(Q) \not\equiv 0$. Using Lemma 4.6, we know that $\Phi_{P,Q}(\mathcal{G}) \not\equiv 0$ and thus by definition of $\Phi_{P,Q}$, we get $\text{lc}_{x_i}(Q)(\mathcal{G}) \not\equiv 0$. This implies that $Q(x_i, \mathcal{G}_{-i}) \not\equiv 0$. We conclude that $f(x_i, \mathcal{G}_{-i}) = (c+1)Q(x_i, \mathcal{G}_{-i}) \not\equiv 0$. Thus whenever $f \not\equiv 0$, the algorithm outputs "NON-ZERO".

Time complexity: By Lemma A.3, degree of generator $\mathcal{G}$ is $\text{poly}((sd)^{d^3}, n)$. Note that f has individual degree at most sd and thus $f(x_i, \mathcal{G}_{-i})$ has individual degree $\leq sd.\deg(\mathcal{G})$. Then by Lemma A.1, testing non-zeroness of the bivariate polynomial $f(x_i, \mathcal{G}_{-i})$ takes only $\text{poly}((sd)^{d^3}, n)$ time. The n iterations only add a factor of n. □

5 Exact Power Testing

In this section we prove Theorem 3. A polynomial $f \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ is an *exact power* if there exists (another) polynomial $g \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ and $e \in \mathbb{N}$ such that $f = g^e$.

We first show an $s^{O(d)}$ sparsity bound for g, when f is an s-sparse, reverse-monic polynomial of individual degree d (Lemma 5.4). Moreover, we also get an algorithm to compute g for this case in Algorithm 2. For the general case, when f is not reverse-monic, we use a standard transformation to convert it into a reverse-monic polynomial $\hat{f}$ (Definition 5.10). This transformation preserves the exact power structure of f, i.e. $f = g^e$ implies $\hat{f} = h^e$, for some suitable polynomial h. One can then invoke Lemma 5.4 to get a nice sparsity bound for h. Unfortunately, we do not recover a sparsity bound for g from this. The main reason is that this transformation is not exactly reversible, as there is no 1-1 correspondence between monomials of f and $\hat{f}$; there is an s^d-sparsity blow-up in $\hat{f}$. Intuitively, there is a 'loss of information' in this transformation. However, we still make use of the 'available information' to get a recursive algorithm for determining whether f is an exact power in Algorithm 3.

5.1 Preliminaries

Notations: Let $f \in \mathbb{F}[x_1, \ldots, x_n]$. For some $i \in [n]$, let x_i be a variable in the support of f. For the sake of convenience, we slightly abuse the notation $\mathbf{x}$ to denote the set $\{x_1, \ldots, x_n\} \setminus \{x_i\}$ throughout this section. Let $f = \sum_{j=0}^{d} f_j \cdot x_i^j$ such that $\forall j$, $f_j \in \mathbb{F}[\mathbf{x}]$ and $f_d \not\equiv 0$. The leading coefficient of f w.r.t x_i is defined as $\mathrm{lc}_{x_i}(f) \triangleq f_d$. Polynomial f is called monic w.r.t variable x_i, if $\mathrm{lc}_{x_i}(f) = 1$. We say that f is x_i-reverse-monic if $f|_{x_i=0} = f_0 = 1$. We say that f is reverse-monic if there exists some variable $x_i \in \mathrm{supp}(f)$ such that f is x_i-reverse monic.

Definition 5.1 (Chapter 6 [GG99])**.** *Let R be a unique factorization domain and K be its field of fractions. Let $g \in K[y]$ be a polynomial over K such that $g = \sum_{i=0}^{m} (g_i/b) \cdot y^i \in K[y]$, where $b \in R$ is the common denominator. The content of g is defined as* $\mathrm{cont}(g) = \gcd(g_0, \ldots, g_m)/b$*. We define primitive part of g as* $\mathrm{pp}(g) = g/\mathrm{cont}(g)$*. Observe that* $\mathrm{cont}(g) \in K$*, while* $\mathrm{pp}(g) \in R[y]$*.*

Examples: For $R = \mathbb{Z}$ and $K = \mathbb{Q}$, consider $g = 3y + 9/2 \in \mathbb{Q}[y]$. Then $g = (6y + 9)/2$ and $\mathrm{cont}(g) = \gcd(6, 9)/2 = 3/2 \in \mathbb{Q}$. And $\mathrm{pp}(g) = g/\mathrm{cont}(g) = 2y + 3 \in \mathbb{Z}[y]$. Let us now consider a bi-variate example. Let $R = \mathbb{F}[x]$ and $K = \mathbb{F}(x)$. Consider $g = (x^2 - 1) \cdot y + (x - 1)/(x + 1) \in \mathbb{F}(x)[y]$. Then, $\mathrm{cont}(g) = \gcd((x^2 - 1)(x + 1), (x - 1))/(x + 1) = (x - 1)/(x + 1) \in \mathbb{F}(x)$. And $pp(g) = (x + 1)^2 \cdot y + 1 \in \mathbb{F}[x][y]$.

The following lemma will be useful to us later on.

Lemma 5.2. *Let R be a unique factorization domain and K be its field of fractions. Let $f \in R[y]$ and $g \in K[y]$ such that $f = g^e$. Then, $g \in R[y]$.*

Proof. By definitions of content and primitive parts in Definition 5.1, we know that $\text{cont}(g) \in K$, while $\text{pp}(g) \in R[y]$ for $g = \text{cont}(g) \cdot \text{pp}(g)$.

Gauss's Lemma states that the product of two primitive polynomials is also primitive. From this, one can derive that for two polynomials $g, h \in K[y]$, $\text{cont}(gh) = \text{cont}(g) \cdot \text{cont}(h)$ and $\text{pp}(gh) = \text{pp}(g) \cdot \text{pp}(h)$. In particular, $\text{cont}(g^e) = \text{cont}(g)^e$ and $\text{pp}(g^e) = \text{pp}(g)^e$. Since $f = g^e$, we get that

$$\text{cont}(f) = \text{cont}(g)^e. \tag{5.3}$$

Since $f \in R[y]$, we know that $\text{cont}(f) \in R$ by definition. We will now use this to prove that $\text{cont}(g) \in R$ also. This will suffice to prove that $g = \text{cont}(g) \cdot \text{pp}(g) \in R[y]$. Note that we can write $\text{cont}(g)$ in the simplest form as $\text{cont}(g) = \frac{a}{b}$, where $a, b \in R$ and $\gcd(a, b) = 1$. Let $d \stackrel{\Delta}{=} \text{cont}(f) \in R$. Using (5.3), we get that $d = \left(\frac{a}{b}\right)^e = \frac{a^e}{b^e}$. Now, let p be an irreducible factor of b in R. Since $d \in R$, p must divide the numerator a^e. If p divides a^e, then p must divide a also. This contradicts with the fact that $\gcd(a, b) = 1$. This means, that the denominator b must be one and hence, $\text{cont}(g) \in R$. Thus, $g \in R[y]$. □

5.2 Exact Power Testing

We start with the case when our input polynomial f is reverse-monic w.r.t. some variable. We generalize it to the general case in the section after.

5.2.1 Reverse Monic Case

We use Newton's Binomial Theorem to get the sparsity bound for $f^{1/e}$ below. This tool has been used before to bound the size of e-th roots for the classes of algebraic circuits, formulas and ABPs. See for example [Dut18, ST21]. Although, these works assume $\text{char}(\mathbb{F})$ to be 0 or a non-divisor of e, we prove our result below for fields of arbitrary characteristic.

Lemma 5.4. *Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be an s-sparse polynomial of individual degree d which is x_i-reverse monic, for some $i \in [n]$. If $f = g^e$ for some polynomial $g \in \mathbb{F}[x_1, \ldots, x_n]$ and $e \in \mathbb{N}$, then g is $s^{d/e+1}$-sparse.*

Proof. We can write g as $g = f^{1/e} = (1 + (f - 1))^{1/e}$. By Newton's Binomial Theorem, this gives

$$g = (1 + (f - 1))^{1/e} = \sum_{i=0}^{\infty} \binom{1/e}{i} (f - 1)^i. \tag{5.5}$$

We first focus on the case when $\text{char}(\mathbb{F})$ is either zero or it does not divide e. In that case $1/e$ is well defined in $\mathbb{F}$ and so are all the binomial coefficients appearing in (5.5). Since f is reverse-monic in x_i, $f|_{x_i=0} = 1$. This means that $(f - 1)$ has x_i-degree ≥ 1. Since g has x_i-degree $= d/e$, (5.5) becomes a finite sum modulo the ideal $\langle x_i \rangle^{d/e+1}$. Thus,

$$g = \sum_{i=0}^{d/e} \binom{1/e}{i} (f - 1)^i \mod \langle x_i \rangle^{d/e+1}. \tag{5.6}$$

Since $\|f-1\| \leq s$, it is easy to see that $G = \sum_{i=0}^{d/e} \binom{1/e}{i} (f-1)^i$ is $s^{d/e+1}$-sparse. Since $g = G \bmod \langle x_i \rangle^{d/e+1}$ and going mod $\langle x_i \rangle^{d/e+1}$ can only decrease sparsity, we get that g is also $s^{d/e+1}$-sparse.

Now we handle the case when char($\mathbb{F}$) divides e. Let p be the characteristic of $\mathbb{F}$, for some prime p. Let $e = p^k \cdot q$, where p^k is the highest power of p which divides e, for some integer $k \geq 1$ and $p \nmid q$. Then by the famous *Frobenius endomorphism*, we know that:

$$\begin{aligned} g(x_1, \ldots, x_n)^p &= g(x_1^p, \ldots, x_n^p) \\ g(x_1, \ldots, x_n)^{p^k \cdot q} &= \left(g(x_1^{p^k}, \ldots, x_n^{p^k}) \right)^q . \end{aligned} \tag{5.7}$$

Since $f = g^e = g^{p^k \cdot q}$, we can use the variable transformation $y_j \leftarrow x_j^{p^k}$, for each $j \in [n]$ to get that

$$f = g(y_1, \ldots, y_n)^q.$$

Observe that x_i-degree in every non-zero monomial of f is a multiple of $e = p^k \cdot q$, therefore f is a proper polynomial in $\mathbb{F}[y_1, \ldots, y_n]$. Moreover f is still s-sparse as the transformation does not affect sparsity. We further note that if f was reverse-monic in x_i, it will also be reverse-monic in y_i. Thus, we have reduced to the previous case, since p does not divide q. Moreover, the individual degree of f is now reduced, specifically y_i-degree of f is $d' = d/p^k$. This implies that $g(y_1, \ldots, y_n)$ has sparsity $\leq s^{d'/q+1} = s^{d/e+1}$. Since this transformation does not affect sparsity, we deduce that our original g is also $s^{d/e+1}$-sparse. □

Remark 5.8. Lemma 5.4 is also true for a monic f of sparsity s such that $f = g^e$. This is because we can convert a monic f into a reverse-monic $\hat{f}$ by the reversal transformation, $\hat{f} = \mathrm{rev}_i^d[f]$ (see Definition 6.3). Observe that if f is monic in x_i, then $\hat{f}$ is reverse-monic in x_i. By definition, this transformation is invertible. In fact, $f = \mathrm{rev}_i^d[\hat{f}]$, thus given $\hat{f}$, we can recover f. We also get that $\|f\| = \|\hat{f}\| = s$. Since $\hat{f} = \hat{g}^e$ and $\hat{g}$ is $s^{d/e+1}$-sparse using Lemma 5.4, we also get that g is $s^{d/e+1}$-sparse.

In fact, Lemma 5.4 gives rise to an algorithm to compute the e^{th} root of a reverse-monic f, as shown below.

Lemma 5.9. *Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be an s-sparse polynomial of individual degree d which is x_i-reverse monic for some $i \in [n]$. If $f = g^e$ for some polynomial $g \in \mathbb{F}[x_1, \ldots, x_n]$ and $e \in \mathbb{N}$, then there is a deterministic algorithm to compute g in* $\mathrm{poly}(s^{d/e}, n, d)$ $\mathbb{F}$*-operations.*

Proof. Algorithm 2 below computes the required g when $f = g^e$, for some reverse-monic f.

Algorithm 2: To compute e^{th} root of f:

Input: Polynomial $f \in \mathbb{F}[x_1, \ldots, x_n]$ of individual degree $\leq d$, s-sparse and reverse-monic in variable x_i such that $f = g^e$.

Output: Root g.

1 Let $p \stackrel{\Delta}{=} \text{char}(\mathbb{F})$.
2 **if** $p > 0$ *and* $p \mid e$ **then**
3 Let $e = p^k \cdot q$, where p^k is the highest power of p that divides e and $p \nmid q$.
4 $f \leftarrow f\left(x_1^{1/p^k}, \ldots, x_n^{1/p^k}\right)$.
5 $d \leftarrow d/p^k$.
6 $e' \leftarrow e$ and $e \leftarrow e/p^k$. `/* Saving value of e in e' and updating e to q */`
7 **end**
8 $G \longleftarrow \sum_{i=0}^{d/e} \binom{1/e}{i} (f-1)^i$.
9 $g \longleftarrow G \pmod{x_i^{d/e+1}}$.
10 **if** $p > 0$ *and* $p \mid e'$ **then**
11 $g \leftarrow g\left(x_1^{p^k}, \ldots, x_n^{p^k}\right)$.
12 **end**
13 **return** g.

Correctness: Follows from Lemma 5.4.

Time Complexity: Steps 2 to 6, except Step 4 can be done in $O(d)$ time. Step 4 will take $\text{poly}(s, n, d)$ time as f is s-sparse. Steps 8 and 9, each take $\text{poly}(s^{d/e})$ $\mathbb{F}$-operations as $\|g\| \leq \|G\| \leq s^{d/e+1}$. Thus, total complexity is $\text{poly}(s^{d/e}, n, d)$ $\mathbb{F}$-operations. □

Remark. Using Remark 5.8, Lemma 5.9 also works for a monic f by first making it reverse-monic, computing its e^{th} root and then returning the reversal of that.

5.2.2 General Case

Now, we handle the case where input f is not monic or reverse-monic in any variable. In this case, we are not able to compute the exact root, but we can solve the decision version of this problem, that is we show how to efficiently test if $f = g^e$, for some g and $e \geq 1$.

We first give a standard trick to convert a polynomial into a reverse-monic polynomial. The properties mentioned below are fairly straightforward to prove, see for example [BSV20, Lem 5.5].

Definition 5.10 (Reverse-monic transformation)**.** *Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be an s-sparse polynomial of individual degree at most d. Pick any variable $x_i \in \text{supp}(f)$ such that $f|_{x_i=0} \not\equiv 0$. Set $x_i = y$ and let*

$f_0 \triangleq f|_{y=0}$. *We define*

$$\hat{f} = \frac{1}{f_0} \cdot f(\mathbf{x}, f_0 \cdot y).$$

This transformation has some nice properties:

1. $\hat{f}$ *is reverse-monic in* y. *Moreover* $\hat{f}$ *is a proper polynomial in* $\mathbb{F}[x_1, \ldots, x_n]$ *(not a rational function).*
2. $\|\hat{f}\| \leq s^d$.
3. *Individual degree of* $\hat{f}$ *is at most* d^2. *However,* $\deg_y(\hat{f}) = \deg_{x_i}(f) \leq d$.

We remark that it could be the case that the trailing coefficient $f_0 = 0$ for every x_i above. We show how to handle that case in Step 3 of Algorithm 3. So without loss of generality, we can always convert our polynomial f into reverse-monic $\hat{f}$.

By definition of this transformation, one can easily show that if $f = g^e$, then $\hat{f} = h^e$, for some suitable h. However, the converse is not always true. For example, consider $f(x, z) = z(x+1)^2$. It is not an exact power but if we make it reverse-monic w.r.t. x we get $\hat{f}(y, z) = 1/z \cdot f(zy, z)$. It turns out to be $\hat{f} = (zy+1)^2$, which is an exact power. For testing whether f is an exact power, we need a converse also. In the two claims below, we find the extra condition on trailing coefficient, which gives us a suitable converse. This will amount to a recursive algorithm for exact power testing in Algorithm 3.

Claim 5.11 ($\Rightarrow$). *If* $f = g^e$ *in* $\mathbb{F}[\mathbf{x}, x_i]$, *then* $\hat{f} = h^e$ *in* $\mathbb{F}[\mathbf{x}, y]$ *for some polynomial* h *and* $f_0 = g_0^e$ *in* $\mathbb{F}[\mathbf{x}]$, *where* $g_0 \triangleq g|_{x_i=0}$.

Proof. Let $f = f_k \cdot x_i^k + \ldots + f_1 \cdot x_i + f_0$ and $g = g_m \cdot x_i^m + \ldots + g_1 \cdot x_i + g_0$. If $f = g^e$, then $k = em$ and $f_0 = g_0^e$. Thus,

$$\begin{aligned}
\hat{f} &= \frac{1}{f_0} \cdot f(\mathbf{x}, f_0 \cdot y) \\
&= \frac{1}{f_0} \cdot g(\mathbf{x}, f_0 \cdot y)^e \\
&= \left(\frac{g(\mathbf{x}, f_0 \cdot y)}{g_0} \right)^e = h^e,
\end{aligned}$$

for $h \triangleq \frac{g(\mathbf{x}, f_0 \cdot y)}{g_0}$. By definition of the reverse-monic transformation, $\hat{f} \in \mathbb{F}[\mathbf{x}, y]$ is a proper polynomial in this ring. Clearly, h is in $\mathbb{F}(\mathbf{x})[y]$ by definition. Also $\hat{f} = h^e \in \mathbb{F}[\mathbf{x}, y]$, therefore h also belongs to $\mathbb{F}[\mathbf{x}, y]$, by Lemma 5.2. □

Claim 5.12 ($\Leftarrow$). *If* $\hat{f} = h^e$ *in* $\mathbb{F}[\mathbf{x}, y]$ *and* $f_0 = b^e$ *in* $\mathbb{F}[\mathbf{x}]$, *then* $f = g^e$ *in* $\mathbb{F}[\mathbf{x}, x_i]$.

Proof. Observe that,

$$\begin{aligned} f(\mathbf{x}, x_i) &= f_0 \cdot \hat{f}\left(\mathbf{x}, \frac{x_i}{f_0}\right) \\ &= f_0 \cdot h\left(\mathbf{x}, \frac{x_i}{f_0}\right)^e \\ &= \left(b \cdot h\left(\mathbf{x}, \frac{x_i}{f_0}\right)\right)^e \\ &= \left(g(\mathbf{x}, x_i)\right)^e, \end{aligned}$$

for $g \stackrel{\Delta}{=} b \cdot h(\mathbf{x}, \frac{x_i}{f_0})$. Clearly, $g \in \mathbb{F}(\mathbf{x})[x_i]$ from above. But since $f \in \mathbb{F}[\mathbf{x}][x_i]$ and $f = g^e$, this implies $g \in \mathbb{F}[\mathbf{x}, x_i]$ by Lemma 5.2. □

Let $r_{\mathbb{F}}(a, e)$ denote the time complexity of deciding whether $a = b^e$, for $a, b \in \mathbb{F}$ and for some $e \in \mathbb{N}$. Then,

- For a finite field $\mathbb{F} = \mathbb{F}_q$, $r_{\mathbb{F}} = \text{poly}(\log q)$ $\mathbb{F}$-opertions (Lemma 5.14).
- For the field of rationals $\mathbb{F} = \mathbb{Q}$, $r_{\mathbb{F}} = \text{poly}(e, \log a)$ $\mathbb{F}$-operations. For an integer (or rational number), it is easy to even compute the e-th root by binary search, or one can simply invoke univariate factorization ([LLL82]) for $x^e - a$ to compute $a^{1/e}$.

Theorem 3 follows from the next lemma.

Lemma 5.13. *Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be an s-sparse polynomial of individual degree d. There is a deterministic algorithm to test whether $f = g^e$ for some polynomial $g \in \mathbb{F}[x_1, \ldots, x_n]$ and $e \in \mathbb{N}$. The algorithm takes $\text{poly}(s^{d^2}, n, d) + r_{\mathbb{F}}(f(0, \ldots, 0), e)$ $\mathbb{F}$-operations.*

Proof. The $e = 1$ case is trivial. Run the Algorithm 3 below for each $e \in \{2, \ldots, d\}$. If any such e exists such that $f = g^e$, that is Algorithm 3 outputs YES, then f is an exact power. Otherwise, if

for every e Algorithm 3 outputs NO, then f is not an exact power.

Algorithm 3: Exact power testing

Input: An s-sparse polynomial $f \in \mathbb{F}[x_1, \ldots, x_n]$ with individual degree d and an integer $e \in \{2, \ldots, d\}$.

Output: YES, if $f = g^e$ for some polynomial g and NO, otherwise.

1 For each $i \in [n]$, check whether f is reverse-monic in variable x_i. If such an i exists, then set $\hat{f} \stackrel{\Delta}{=} f, y \stackrel{\Delta}{=} x_i$ and go to Step 8 directly, else go to Step 2.

2 Choose any $i \in [n]$. Set $f_0 \stackrel{\Delta}{=} f|_{x_i=0}$ and $x_i \stackrel{\Delta}{=} y$.

3 **if** $f_0 = 0$ **then**

4 Let k be the highest power of x_i such that x_i^k divides f.

5 If $e \not| k$ then output NO and return, otherwise set $f = f/x_i^k$ and $f_0 = f|_{x_i=0}$.

6 **end**

7 Define $\hat{f} \stackrel{\Delta}{=} \frac{1}{f_0} \cdot f(\mathbf{x}, f_0 \cdot y)$.

8 Invoke Algorithm 2 for $\hat{f}$ which is reverse-monic in variable y to get candidate root $\hat{g}$.

9 Check whether $\hat{f} = \hat{g}^e$, by multiplying out. If it is, go to Step 10, otherwise output NO.

10 For the $(n-1)$-variate polynomial $f_0 \stackrel{\Delta}{=} f|_{x_i=0} \in \mathbb{F}[\mathbf{x}]$, recursively check whether f_0 is e^{th} power of some polynomial. If it is, then output YES, otherwise output NO.

We discuss the correctness and time-complexity of Algorithm 3 below.

Correctness: If f is indeed equal to g^e for some $e \in \{2, \ldots, d\}$, then by Claim 5.11, $\hat{f} = h^e$ for some h and $f_0 = b^e$, for $b = g|_{x_i=0}$. Thus, by Lemma 5.9, Step 8 will compute the correct root $\hat{g}$ and in Step 10, the algorithm will output YES. If f is not an exact power for any $e \in [d]$, the algorithm will output NO in either Step 5 or Step 9 or Step 10. This follows due to Claim 5.12 (consider contrapositive).

Time Complexity: Step 1 takes $\text{poly}(s, n, d)$ $\mathbb{F}$-operations as we only have to check whether $f|_{x_i=0} = 1$ at most n times. In Steps 2-6, we are required to compute the trailing coefficient of f w.r.t x_i-variable, which takes $\text{poly}(s, n, d)$ $\mathbb{F}$-operations. Step 7 takes at most $\text{poly}(s^d, n, d)$ time. Step 8 takes at most $\text{poly}(\|\hat{f}\|^{d/e}, n, d)$ $\mathbb{F}$-operations by Lemma 5.9 as y-degree of $\hat{f}$ is still d. Since $\|\hat{f}\| \leq s^d$, this step takes at most $\text{poly}(s^{d^2/e}, n, d)$ $\mathbb{F}$-operations. Multiplying out in Step 9 will take take at most $\text{poly}(s^{(d^2/e)\cdot e}) = \text{poly}(s^{d^2})$ $\mathbb{F}$-operations as $\|\hat{g}\| \leq s^{d^2/e+1}$. Note that in Step 10, we recurse on f_0, which has sparsity $\leq s$, therefore there is no blow-up of sparsity in the recursion. Hence, this step takes $\text{poly}(s^{d^2}, n, d)$ $\mathbb{F}$-operations to reach the base case of deciding whether the field element $f(0, \ldots, 0)$ has an e-th root. The total complexity is thus, $\text{poly}(s^{d^2}, n, d) + r_{\mathbb{F}}(f(0, \ldots, 0), e)$ $\mathbb{F}$-operations. □

Using the standard theory of finite fields, we show how to test whether a finite field element is

an exact power. In other words, we show that $r_{\mathbb{F}} = \text{poly}(\log q)$ for a finite field F_q. This is required for the base case of Algorithm 3, when working over finite fields.

Lemma 5.14 (Folklore)**.** *For a finite field $\mathbb{F}_q$, we can decide whether an element $a \in \mathbb{F}_q$ is a k^{th} power residue, i.e. $a = b^k$ for some $b \in \mathbb{F}_q$ in* $\text{poly}(\log q)$ $\mathbb{F}_q$*-operations.*

Proof. We will focus on $\mathbb{F}_q^*$ since $0 = 0^k$ trivially. We will prove that $a = b^k$ for some $b \in \mathbb{F}_q^*$ and $k \geq 1$, if and only if $a^{\frac{q-1}{d}} = 1$ in $\mathbb{F}_q^*$, where $d = \gcd(k, q-1)$. Having proved that, we can simply test this by computing $a^{\frac{q-1}{d}}$ in $\text{poly}(\log q)$ $\mathbb{F}_q$-operations using repeated squaring. Now we prove both the directions for

$$a = b^k \Leftrightarrow a^{\frac{q-1}{d}} = 1.$$

($\Rightarrow$) Observe that $a = b^k \Rightarrow a^{\frac{q-1}{d}} = b^{\frac{k(q-1)}{d}}$. Since $d = \gcd(k, q-1)$, d divides k, hence $\frac{k}{d}$ is an integer. Thus, we get that $a^{\frac{q-1}{d}} = b^{\frac{k(q-1)}{d}} = 1$ by using the generalization of Fermat's Little Theorem, i.e. $x^{q-1} = 1$ for all $x \in \mathbb{F}_q^*$.

($\Leftarrow$) We know that $\mathbb{F}_q^*$ is a cyclic group of order $q-1$. Let g be its generator such that $a = g^r$, for some $r \in [q-2]$ (For $r = 0$, we know that $1 = 1^k$ trivially). Now, if $a^{\frac{q-1}{d}} = 1$, we get that $g^{\frac{r(q-1)}{d}} = 1$. This implies that $\frac{r}{d}$ is an integer or that d divides r. By Bezout's identity, we know that $d = \gcd(\mathrm{k}, \mathrm{q}-1) = sk + t(q-1)$ for some integers s, t. Since $d \mid r$, we get that $r = s'k + t'(q-1)$. This proves that a is a k^{th} power residue as:

$$a = g^r = g^{s'k + t'(q-1)} = g^{s'k} = b^k$$

in $\mathbb{F}_q^*$ for $b = g^{s'}$. □

6 Co-Factor Polynomial Sparsity

In this section we prove Theorem 4 as well as Corollary 1.4. In fact, we prove somewhat technically stronger versions of these theorems with more explicit parameters (not just in terms of big-Oh). We begin with some technical definitions. Some of these have been used implicitly in the previous sections.

6.1 Multivariate Reversal operation

Definition 6.1 (Reverse Monic, Reverse Pseudo-Monic)**.** *We say that a polynomial $f \in \mathbb{F}[x_1, \ldots, x_n]$ is* reverse monic *if there exists a variable $x_i \in \text{supp}(f)$ such that $f|_{x_i=0} = 1$. If we know the variable beforehand, we say that f is x_i-*reverse monic*. We can extend this definition to a set of variables x_I, for some arbitrary $I \subseteq [n]$. We say that f is I-*reverse monic*, if $f|_{x_I=0_I} = 1$. We say that f is* reverse pseudo-monic, *x_i-*reverse pseudo-monic, *I-*reverse pseudo-monic *respectively, when instead of* 1 *the result of setting variables to* 0 *is a non-zero field element or a single monomial.*

In other words, the constant term of a reverse monic polynomial is 1, when regarded as a polynomial in the remaining variables. The following are immediate connections between some of the previously defined concepts.

Observation 6.2. *Let $h \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a polynomial, $i \in [n]$ and $I \subseteq [n]$. Then:*

- $\text{supp}(h|_{x_i=0}) = \{\mathbf{e} \in \text{supp}(h) \mid e_i = 0\}$.
- *h is I-reverse pseudo-monic if and only if $|\text{supp}(h|_{x_I=0_I})| = 1$.*

The following transformation will be useful later on to convert specific polynomials into I-reverse pseudo-monic polynomials.

Definition 6.3 (Reversal Transformation)**.** *Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be a polynomial and let $\ell \in \mathbb{N}$. We define the* reversal operation *on f with respect to a variable x_i as follows:*

$$\text{rev}_i^\ell[f] \triangleq x_i^\ell \cdot f|_{x_i=\frac{1}{x_i}} = x_i^\ell \cdot f(x_1, \ldots, x_{i-1}, \frac{1}{x_i}, x_{i+1}, \ldots, x_n).$$

By iteration, we can extend this definition to a set of variables x_I, for some arbitrary $I = \{i_1, \ldots, i_r\} \subseteq [n]$.

$$\text{rev}_I^\ell[f] \triangleq \text{rev}_{i_1}^\ell \left[\text{rev}_{i_2}^\ell[\ldots \text{rev}_{i_r}^\ell[f]]\right].$$

Alternatively:

$$\text{rev}_I^\ell[f] \triangleq x_{i_1}^\ell \cdots x_{i_r}^\ell \cdot f(y_1, \ldots, y_n),$$
$$\text{where } y_j = \begin{cases} \frac{1}{x_j}, & \text{if } x_j \in I \\ x_j, & \text{if } x_j \notin I. \end{cases}$$

For intuition, express f as a polynomial in x_i such that $f = f_d x_i^d + f_{d-1} x_i^{d-1} + \ldots + f_1 x_i + f_0$, where each coefficient f_j is a polynomial in variables other than x_i. Then, $\text{rev}_i^d[f]$ reverses the order of coefficients in this representation. That is, $\text{rev}_i^d[f] = f_0 x_i^d + f_1 x_i^{d-1} + \ldots + f_{d-1} x_i + f_d$. In particular, if f is monic in x_i then $\text{rev}_i^d[f]$ is x_i-reverse monic.

The following lemma summarizes some of the basic, yet useful properties of the reversal transformation. Subsequently, we will use these properties implicitly.

Lemma 6.4. *Let $f, g, h \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ such that $f = g \cdot h$. Let $i \in [n]$ and suppose that $d \geq \deg_{x_i}(f)$. Then:*

1. *$\text{rev}_i^d[f]$ is a polynomial (and not a rational function).*
2. *$\deg_{x_i}(\text{rev}_i^d[f]) \leq d$.*
3. *$\|\text{rev}_i^d[f]\| = \|f\|$.*
4. *Let a, b such that $d = a + b$. Then $\text{rev}_i^d[f] = \text{rev}_i^a[g] \cdot \text{rev}_i^b[h]$.*

6.2 Unique Projections

Definition 6.5. *Let $V \subseteq \mathbb{N}^n$. A* unique projection *of V of* length k *is a set $\{(i_1, e_1), (i_2, e_2), \ldots, (i_k, e_k)\}$ such that there exists a unique vector $\mathbf{v} \in V$ satisfying $\forall j \in [k] : v_{i_j} = e_j$.*
A unique projection of a polynomial $h \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ is defined as a unique projection of $\mathrm{supp}(h)$.

In other words, there exists a unique monomial in the monomial representation of h that contains the pattern $x_{i_1}^{e_1} x_{i_2}^{e_2} \cdots x_{i_k}^{e_k}$. The following is immediate from the definition.

Observation 6.6. *Let $h \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a polynomial and let $\{(i_1, e_1), (i_2, e_2), \ldots, (i_k, e_k)\}$ be a unique projection of h. Pick $j \in [k]$ and let $\ell \geq e_j$. Then*

$$\{(i_1, e_1), (i_2, e_2), \ldots, (i_{j-1}, e_{j-1}), (i_j, \ell - e_j), (i_{j+1}, e_{j+1}), \ldots, (i_k, e_k)\}$$

is a unique projection of $\mathrm{rev}_{i_j}^{\ell}[h]$.

Subsequently, we demonstrate the usefulness of unique projections.

Lemma 6.7. *Let $h \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a polynomial with a unique projection of the form $\{(i_1, 0), (i_2, 0), \ldots, (i_k, 0)\}$ (i.e. $\forall j \in [k] : e_k = 0$). Then h is $\{i_1, i_2, \ldots, i_k\}$-reverse pseudo-monic.*

Proof. Let $I = \{i_1, i_2, \ldots, i_k\}$. By iterative application of Part 1 of Observation 6.2, we obtain that

$$\mathrm{supp}(h|_{x_I = 0_I}) = \left\{ \mathbf{e} \in \mathrm{supp}(h) \;\middle|\; \forall j \in [k] : e_{i_j} = 0 \right\}$$

As I corresponds to a unique projection, the set of the RHS contains exactly one vector and the claim follows from Part 2 of Observation 6.2. □

Lemma 6.8. *Let $h \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a multilinear polynomial and let $\{(i_1, e_1), (i_2, e_2), \ldots, (i_k, e_k)\}$ be a unique projection of h. Furthermore, let $J = \{i_j \mid e_j = 1\}$. That is, the set of all indices i_j for which $e_j = 1$. Then $\tilde{h} \stackrel{\Delta}{=} \mathrm{rev}_J^1[h]$ is $\{i_1, i_2, \ldots, i_k\}$-reverse pseudo-monic.*

Proof. First, note that since h is a multilinear polynomial, we have that $e_j = 0$ for indices $j \in \{i_1, \ldots, i_k\} \setminus J$. Subsequently, by iterative application of Observation 6.6, we obtain that $\tilde{h}$ is a multilinear polynomial with a unique projection $\{(i_1, 0), (i_2, 0), \ldots, (i_k, 0)\}$. Note that $\tilde{h}$ is a proper polynomial (and not a rational function) by iterative application of Part 1 in Lemma 6.4. The claim then follows from Lemma 6.7. □

We conclude this section by showing that every set contains a unique projection of (at most) logarithmic size and a relation of unique projections with δ-entropy polynomials that were defined in [BS21].

Lemma 6.9. *Let $V \subseteq \mathbb{N}^n$ of size $|V| \leq s$. Then V has a unique projection of length at most* $\log s + 1$.

Proof. The proof is by induction on the size of V. For the base case $|V| = 1$ there exists a unique projection of length 1. Now assume $|V| \geq 2$. Therefore, V contains at least two different vectors $\mathbf{u} \neq \mathbf{w}$. Let i be such that $u_i \neq w_i$. Let us denote $a = u_i$ and $b = w_i$. Partition V into $V_a \stackrel{\Delta}{=} \{\mathbf{v} \in V \mid v_i = a\}$ and $V_b \stackrel{\Delta}{=} \{\mathbf{v} \in V \mid v_i = b\}$. We have that $|V_a| + |V_b| \leq |V|$. Hence, wlog $1 \leq |V_a| \leq s/2$. By the induction hypothesis, V_a has a unique projection of length at most $\log(s/2) + 1 = \log s$. We now add the index i and $e_i = a$ to the set to obtain a unique projection for V of size $\log s + 1$. □

Recently, [BS21] defined a class of polynomials called 'low-entropy' polynomials and showed an $(nd)^{O(d\delta)}$ sparsity upper bound for the factors of a δ-entropy polynomial. We quickly give their definition of a δ-entropy set and then show a combinatorial connection of entropy with our notion of unique projections below. We note this connection between these two combinatorial concepts but our results are incomparable from those in [BS21].

Definition 6.10 ([BS21])**.** *A vector $\mathbf{v} \in \mathbb{N}^n$ has entropy δ if it has the same value in $(n - \delta)$ of its coordinates. A set $V \subseteq \mathbb{N}^n$ is called a δ-entropy set if for every $\mathbf{v} \in V$, $\mathbf{v}$ has entropy $\leq \delta$.*

Lemma 6.11. *Let $V \subseteq \mathbb{N}^n$ be a δ-entropy set. Then V has a unique projection of length at most $2\delta + 1$.*

Proof. Let $\mathbf{u} \in V$ be a vector with maximum entropy in V. Let $m(\mathbf{u})$ denote the majority value in $\mathbf{u}$. In other words, $\mathbf{u}$ has the largest number of non-majority values among all vectors in V. Let $u_{i_1}, \ldots, u_{i_k}$ be all the non-majority values in $\mathbf{u}$. Since $\mathbf{u}$ has entropy $\leq \delta$, the length of this sequence k is at most δ. We note that the remaining elements of $\mathbf{u}$ outside this sequence have the same value (equal to $m(\mathbf{u})$). We extend the sequence to length $k + \delta + 1$ by adding any $\delta + 1$ elements from these remaining elements of $\mathbf{u}$ to get: $u_{i_1}, \ldots, u_{i_k}, u_{i_{k+1}}, \ldots, u_{i_{k+\delta+1}}$. We claim that $\{(i_1, u_{i_1}), \ldots, (i_{k+\delta+1}, u_{k+\delta+1})\}$ is a unique projection of the set V.

We need to show that if $\mathbf{v}$ is another vector in V such that $v_{i_j} = u_{i_j}$, for each $j \in [k + \delta + 1]$ (values agree on projection), then $\mathbf{v} = \mathbf{u}$. Observe that $u_{i_{k+1}} = \ldots = u_{i_{k+\delta+1}} = m(\mathbf{u})$, by definition of k. This means $v_{i_{k+1}} = \ldots = v_{i_{k+\delta+1}} = m(\mathbf{u})$ also. We deduce that at least $\delta + 1$ coordinates of $\mathbf{v}$ have value $m(\mathbf{u})$. We also know that $\mathbf{v}$ has entropy $\leq \delta$ and note that for a $\leq \delta$-entropy vector, if any $\delta + 1$ coordinates have the same value, that value is the majority value. Hence, $m(\mathbf{v}) = m(\mathbf{u})$. Now suppose for the sake of contradiction that there exists some coordinate i_r outside the projection ($r > k + \delta + 1$) for which $v_{i_r} \neq u_{i_r}$. Since all the non-majority values of $\mathbf{u}$ have already appeared in the projection coordinates, we deduce that $u_{i_r} = m(\mathbf{u}) = m(\mathbf{v})$. This means that $v_{i_r} \neq m(\mathbf{v})$. In that case, v_{i_r} is another element in $\mathbf{v}$ apart from $v_{i_1}, \ldots, v_{i_k}$ which is distinct from $m(\mathbf{v})$. This is a contradiction to our assumption that $\mathbf{u}$ is a vector with maximum entropy. Hence, $v_{i_j} = u_{i_j}$ for all $j > k + \delta + 1$. By our premise, they also agree on the $k + \delta + 1$ projection coordinates. Hence $v_j = u_j$, for all $j \in [n]$ and thus, $\mathbf{v} = \mathbf{u}$. Moreover, since $k \leq \delta$, length of this unique projection is $k + \delta + 1 \leq 2\delta + 1$. □

6.3 Proofs of Theorem 4 and Corollary 1.4

We are now ready to state and prove the technical results of this section that will imply Theorem 4 and Corollary 1.4. In what follows, let $f, h \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be two s-sparse polynomials such that $f = gh$.

Lemma 6.12. *Suppose that h is reverse monic and the individual degree of g is at most d. Then g is s^{d+2}-sparse.*

Proof. By hypothesis, let h be reverse monic w.r.t. some variable $x_i \in \text{supp}(h)$. Express h as a univariate in x_i with coefficients as polynomials in the remaining variables. Since h is x_i-reverse monic, the constant term, $h|_{x_i=0}$ is 1. Therefore, every term in $(1-h)$ has x_i-degree ≥ 1. We use this observation in a division-elimination argument as follows:

$$g = \frac{f}{h} = \frac{f}{1-(1-h)} = \sum_{j=0}^{\infty} f(1-h)^j. \tag{6.13}$$

Let x_i-degree of g be d_i. Then, we can safely truncate the infinite sum in Equation (6.13) as follows:

$$g = \sum_{j=0}^{d_i} f(1-h)^j \bmod \langle x_i^{d_i+1} \rangle. \tag{6.14}$$

Equation (6.14) helps us in bounding sparsity of g. Note that going mod $\langle x_i^{d_i+1} \rangle$ can only decrease sparsity, so we focus only on the sparsity of finite sum in (6.14). Since g is a factor of f, its individual degree d_i is also upper bounded by d. Also note that both $\|f\|, \|(1-h)\| \leq s$. Therefore, we get that $\|g\| \leq \sum_{j=0}^{d} s^{j+1} \leq s^{d+2}$. □

Generalizing this observation we obtain:

Lemma 6.15. *Suppose that h is I-reverse monic and the individual degrees of the variables of g in x_I are at most d. Then g is $s^{d|I|+2}$-sparse.*

Proof. We follow the same template as in proof of Lemma 6.12, with the change that h is reverse monic with respect to a set I of variables instead of just a single variable. Express h as a polynomial in x_I variables with coefficients as polynomials in the remaining $n - |I|$ variables. Since h is I-reverse monic, $h|_{x_I=0_I}$ (the constant term of h) is 1. Therefore, every term in $(1-h)$ has total x_I-degree ≥ 1. We then get the same Equation (6.13) for g. Let $I = \{i_1, \ldots, i_k\} \subseteq [n]$, where $k = |I|$. Let individual degree of variable x_{i_j} in g be d_j for each $j \in [k]$. Then, we can truncate the infinite sum as follows:

$$g = \sum_{j=0}^{dk} f(1-h)^j \bmod \langle x_{i_1}^{d_1+1}, \ldots, x_{i_k}^{d_k+1} \rangle. \tag{6.16}$$

By the premises, for each $j \in [k]$ each individual degree d_j is upper bounded by d. Therefore, we only need to sum up to $j = dk$ in (6.16) as the total degree in x_I variables is upper bounded by dk. Therefore, we get that $\|g\| \leq \sum_{j=0}^{dk} s^{j+1} \leq s^{dk+2} = s^{d|I|+2}$. □

The next lemma transforms a pseudo-monic polynomial into a monic polynomial while maintaining the sparsity and the multiplicative properties.

Lemma 6.17. *Let $f = gh$. Suppose that h is I-reverse pseudo-monic and the individual degrees of the variables of g in x_I are at most d. Then there exists polynomials $\tilde{f}, \tilde{g}, \tilde{h} \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ such that:*

1. *$\tilde{h}$ is I-reverse monic.*
2. *$\tilde{f} = \tilde{g}\tilde{h}$.*
3. *$\|\tilde{f}\| = \|f\|, \|\tilde{g}\| = \|g\|, \|\tilde{h}\| = \|h\|$.*
4. *The individual degrees of the variables of $\tilde{g}$ in x_I are at most d.*

Proof. Let $\alpha \triangleq h|_{x_I=0_I}$. We first define $\hat{f}, \hat{g}$ and $\hat{h}$ by setting $x_i \triangleq x_i \cdot \alpha$ for all $i \in I$, into f, g and h, respectively. Next, we set $\tilde{f} \triangleq \hat{f}$, $\tilde{g} \triangleq \hat{g} \cdot \alpha$ and $\tilde{h} = \hat{h}/\alpha$. We will now prove each part of the claim.

1. First, observe that $\tilde{h}$ is, indeed, a polynomial (and not a rational function). This is due to the fact that α divides $\hat{h}$. Next, $\tilde{h}|_{x_I=0_I} = \hat{h}|_{x_I=0_I}/\alpha = h|_{x_I=0_I}/\alpha = 1$.
2. $\tilde{f} = \hat{f} = \hat{g}\hat{h} = (\hat{g} \cdot \alpha)(\hat{h}/\alpha) = \tilde{g}\tilde{h}$.
3. Since α is a monomial or a field element there is $1-1$ correspondence between the monomials of f, g, h and $\tilde{f}, \tilde{g}, \tilde{h}$, respectively.
4. By definition, $\alpha \in \mathbb{F}[x_{[n]\setminus I}]$. Hence, multiplication or division by α does not affect the degrees of the variables in I. □

By transforming a pseudo-monic polynomial into a monic polynomial we can generalize Lemma 6.15 to the pseudo-monic case.

Corollary 6.18. *Suppose that h is I-reverse pseudo-monic and the individual degrees of the variables of g in x_I are at most d. Then g is $s^{d|I|+2}$-sparse.*

Proof. Apply Lemma 6.15 on $\tilde{f}, \tilde{g}$ and $\tilde{h}$ from Lemma 6.17. We obtain that $\tilde{g}$ and hence g is $s^{d|I|+2}$-sparse. □

Remark 6.19. In the context of exact-root sparsity, we can extend the result of Lemma 5.4 from the reverse monic to the I-reverse pseudo-monic case. It is done in the exact same fashion as we moved from Lemma 6.12 to Corollary 6.18 above. The formal statement is given in Theorem 6.20 below. In addition, we observe that if $f = g^e$ then f is I-reverse pseudo-monic iff g is I-reverse pseudo-monic (for the exact same I).

Theorem 6.20. *Let $f \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a polynomial of sparsity s and individual degree at most d such that $f = g^e$ for some (other) polynomial $g \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ and $e \in \mathbb{N}$. In addition, suppose that f is I-reverse pseudo-monic for some $I \subseteq [n]$. Then the sparsity of g is bounded by $s^{O(d \cdot |I|/e)}$.*

Theorem 4 and Corollary 1.4 follow from the next two claims.

Theorem 6.21. *Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be an s-sparse polynomial, with a multilinear factor h such that $f = g \cdot h$. Suppose that the individual degree of g is at most d and h has a unique projection of length at most k. Then g is s^{dk+2}-sparse.*

Proof. Let $\{(i_1, e_1), (i_2, e_2), \ldots, (i_k, e_k)\}$ be the guaranteed unique projection of h and let $J = \{i_j \mid e_j = 1\}$. We define:

$$\tilde{f} \triangleq \mathrm{rev}_J^{d+1}[f] \ , \ \tilde{g} \triangleq \mathrm{rev}_J^{d}[g] \text{ and } \tilde{h} \triangleq \mathrm{rev}_J^{1}[h].$$

By Lemma 6.4, we have that $\tilde{f} = \tilde{g} \cdot \tilde{h}$, where $\tilde{f}$ is an s-sparse polynomial, $\tilde{g}$ is a polynomial with individual degree at most d and $\tilde{h}$ is a multilinear polynomial. Furthermore, by Lemma 2.1, $\|\tilde{h}\| \leq \|\tilde{f}\| \leq s$. Finally, by Lemma 6.8, $\tilde{h}$ is $\{i_1, i_2, \ldots, i_k\}$-reverse pseudo-monic. Consequently, by Corollary 6.18, we obtain that $\tilde{g}$ and hence g are s^{dk+2}-sparse. □

Corollary 6.22. *Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be an s-sparse polynomial, such that $f = g \cdot h$ where h is multilinear polynomial and g is a polynomial with individual degree at most d. Then g is $s^{d(\log s+1)+2}$-sparse.*

Proof. By Lemma 2.1, $\|h\| \leq \|f\| \leq s$. Consequently, by Lemma 6.9, h has a unique projection length at most $\log s + 1$. Further using Theorem 6.21, we deduce that $\|g\| \leq s^{d(\log s+1)+2}$. □

Similarly, by plugging in Lemma 6.11 into Theorem 6.21 we obtain the following relation to low-entropy polynomials.

Corollary 6.23. *Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be an s-sparse polynomial, such that $f = g \cdot h$ where h is a δ-entropy, multilinear polynomial and g is a polynomial with individual degree at most d. Then g is $s^{d(2\delta+1)+2}$-sparse.*

Remark 6.24. By using the formal expansion:

$$\frac{1}{(1-x)^\ell} = \sum_{j=0}^{\infty} \binom{j+\ell-1}{j} x^j$$

for division elimination in the proof of Lemma 6.15, we can get somewhat stronger versions of Theorem 4 below.

Theorem 6.25. *Let $f \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a polynomial of sparsity s and individual degree at most d such that $f = gh^\ell$ for some $\ell \in \mathbb{N}$. Suppose, in addition, that h is a multilinear polynomial with a unique projection of length k. Then the sparsity of g is bounded by $s^{O((d-\ell)k)}$.*

7 Future Directions

A lot of interesting open problems arise in the context of this work:

- Design a polynomial-time PIT algorithm for $\Sigma^{[k]}\Pi\Sigma\Pi^{[\text{ind-deg } d]}$ circuits with bounded k and d, for $k \geq 3$. To the best of our knowledge, the smallest open case is $k = 3$ and $d = 1$!
- Prove a polynomial-size sparsity bound (Conjecture 1.3) even for the special cases like exact-roots, multilinear co-factors.
  - In particular, improve the sparsity bound in Corollary 1.4. Ideally, get rid of the $\log s$ term in the exponent. One can start by studying the structure of polynomials with non-constant or log-sized unique projections.
  - Likewise, generalize Theorem 6.20 to work for any general f with bounded individual degree d. The smallest open case here is $d = 4$ and $e = 2$, in other words prove that square-root is sparse.
  - Can we show that $\text{Res}_{x_i}(u, v)$ is actually sparse (or "somewhat sparse") under the premises of Lemma 4.5?

Acknowledgments

The authors would like to thank the anonymous referees for their useful comments that improved the presentation of the results.

References

[AGS19] M. Agrawal, S. Ghosh, and N. Saxena. Bootstrapping variables in algebraic circuits. *Proceedings of the National Academy of Sciences*, 116(17):8107–8118, 2019. 6

[Alo99] N. Alon. Combinatorial Nullstellensatz. *Combinatorics, Probability and Computing*, 8:7–29, 1999. 43

[AV08] M. Agrawal and V. Vinay. Arithmetic circuits: A chasm at depth four. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 67–75, 2008. 6

[BOT88] M. Ben-Or and P. Tiwari. A Deterministic Algorithm for Sparse Multivariate Polynomial Interpolation. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC)*, pages 301–309, 1988. 2

[BS21] P. Bisht and N. Saxena. Derandomization via symmetric polytopes: Poly-time factorization of certain sparse polynomials. 2021. 6, 34, 35

[BSV20] V. Bhargava, Sh. Saraf, and I. Volkovich. Deterministic Factorization of Sparse Polynomials with Bounded Individual Degree. *J. ACM*, 67(2):8:1–8:28, 2020. 2, 3, 5, 6, 8, 11, 13, 28

[CLO15] D. A. Cox, J. Little, and D. O'Shea. *Ideals, varieties, and algorithms - an introduction to computational algebraic geometry and commutative algebra (4. ed.)*. Undergraduate texts in mathematics. Springer, 2015. 4, 14

[CR88] B. Chor and R. L. Rivest. A knapsack-type public key cryptosystem based on arithmetic in finite fields. *IEEE Transactions on Information Theory*, 34(5):901–909, 1988. 1

[DDS21] P. Dutta, P. Dwivedi, and N. Saxena. Deterministic identity testing paradigms for bounded top-fanin depth-4 circuits. In *36th Conference on Computational Complexity (CCC 2021)*, volume 5, page 9, 2021. 7, 19

[DL78] R. A. DeMillo and R. J. Lipton. A Probabilistic Remark on Algebraic Program Testing. *Inf. Process. Lett.*, 7(4):193–195, 1978. 3

[Dut18] P. Dutta. Discovering the roots: Unifying and extending results on multivariate polynomial factoring in algebraic complexity. *Master's thesis, Chennai Mathematical Institute*, 2018. 26

[For15] M. A. Forbes. Deterministic divisibility testing via shifted partial derivatives. In *FOCS*, 2015. 7

[GCL92] K. O. Geddes, S. R. Czapor, and G. Labahn. *Algorithms for computer algebra*. Kluwer, 1992. 4, 14, 16

[GG99] J. von zur Gathen and J. Gerhard. *Modern computer algebra*. Cambridge University Press, 1999. 4, 14, 25

[GJR10] E. Grigorescu, K. Jung, and R. Rubinfeld. A local decision test for sparse polynomials. *Inf. Process. Lett.*, 110(20):898–901, 2010. 2

[GK85] J. von zur Gathen and E. Kaltofen. Factoring Sparse Multivariate Polynomials. *Journal of Computer and System Sciences*, 31(2):265–287, 1985. 2, 6

[GS99] V. Guruswami and M. Sudan. Improved decoding of Reed-Solomon codes and algebraic-geometry codes. *IEEE Transactions on Information Theory*, 45(6):1757–1767, 1999. 1

[Kal89] E. Kaltofen. Factorization of polynomials given by straight-line programs. In S. Micali, editor, *Randomness in Computation*, volume 5 of *Advances in Computing Research*, pages 375–412. JAI Press Inc., Greenwhich, Connecticut, 1989. 1

[KI04] V. Kabanets and R. Impagliazzo. Derandomizing Polynomial Identity Tests Means Proving Circuit Lower Bounds. *Computational Complexity*, 13(1-2):1–46, 2004. 1

[KS01] A. Klivans and D. Spielman. Randomness efficient identity testing of multivariate polynomials. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 216–223, 2001. 5, 8, 43

[KT90] E. Kaltofen and B. M. Trager. Computing with Polynomials Given by Black Boxes for Their Evaluations: Greatest Common Divisors, Factorization, Separation of Numerators and Denominators. *J. of Symbolic Computation*, 9(3):301–320, 1990. 1

[LLL82] A.K. Lenstra, H.W. Lenstr, and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen,*, 261(4):515–534, 1982. 30

[LST22] N. Limaye, S. Srinivasan, and S. Tavenas. Superpolynomial lower bounds against low-depth algebraic circuits. In *FOCS 2021*, 2022. 7, 19

[PS21] S. Peleg and A. Shpilka. Polynomial time deterministic identity testing algorithm for Σ [3] $\Pi\Sigma\Pi$ [2] circuits via Edelstein–Kelly type theorem for quadratic polynomials. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 259–271, 2021. 7, 19

[Sch80] J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, 1980. 3

[SSS13] C. Saha, R. Saptharishi, and N. Saxena. A Case of Depth-3 Identity Testing, Sparse Factorization and Duality. *Computational Complexity*, 22(1):39–69, 2013. 6

[ST21] A. Sinhababu and T. Thierauf. Factorization of polynomials given by arithmetic branching programs. *computational complexity*, 30(2):1–47, 2021. 26

[Str73] V. Strassen. Vermeidung von Divisionen. *J. of Reine Angew. Math.*, 264:182–202, 1973. 9

[Sud97] M. Sudan. Decoding of Reed Solomon codes beyond the error-correction bound. *Journal of Complexity*, 13(1):180–193, 1997. 1

[SV10] A. Shpilka and I. Volkovich. On the Relation between Polynomial Identity Testing and Finding Variable Disjoint Factors. In *Automata, Languages and Programming, 37th International Colloquium (ICALP)*, pages 408–419, 2010. Full version at https://eccc.weizmann.ac.il/report/2010/036. 2

[SV15] A. Shpilka and I. Volkovich. Read-Once Polynomial Identity Testing. *Computational Complexity*, 24(3):477–532, 2015. 43, 44

[SV18] S. Saraf and I. Volkovich. Blackbox Identity Testing for Depth-4 Multilinear Circuits. *Combinatorica*, 38(5):1205–1238, 2018. 7

[SY10] A. Shpilka and A. Yehudayoff. Arithmetic Circuits: A survey of recent results and open questions. *Foundations and Trends in Theoretical Computer Science*, 5(3-4):207–388, 2010. 42

[Vol15] I. Volkovich. Deterministically Factoring Sparse Polynomials into Multilinear Factors and Sums of Univariate Polynomials. In *APPROX-RANDOM*, pages 943–958, 2015. 2

[Vol17] I. Volkovich. On some Computations on Sparse Polynomials. In *APPROX-RANDOM*, pages 48:1–4:21, 2017. 2, 3, 6, 7

[Zip79] R. Zippel. Probabilistic algorithms for sparse polynomials. In *Proceedings of the International Symposium on Symbolic and Algebraic Computation*, pages 216–226, 1979. 2, 3

A Basics of algebraic complexity

In this section, we formally define various algebraic models of computation. We refer the reader to the excellent survey of [SY10] for a detailed discussion.

A.1 Algebraic computational models:

An *algebraic circuit* or an arithmetic circuit is a directed acyclic graph with input leaves at the bottom and a single output node at top, where the computation is done bottom-up. The leaves are labeled with variables or field constants while the internal nodes are either addition or multiplication gates. A directed edge between two nodes $u \to v$ is labeled with field constants, which gets multiplied to the polynomial computed at node u before feeding it to node v. The in-degree of a node is called its *fan-in* and out-degree is called *fan-out*. *Size* of the circuit is simply size of the directed graph, which is the maximum among number of edges and number of nodes. *Depth* of the circuit is length of the longest path from a leaf to the output node. *Degree* of the circuit is maximum degree of a polynomial computed at any node in the circuit.

A *depth-2* $\Sigma\Pi$ *circuit* of size s computes a sum of s-many monomials. Thus, depth-2 circuits compute the class of sparse polynomials. A size s, depth-4 $\Sigma\Pi\Sigma\Pi$ circuit computes a polynomial of the form $f = \sum_{i=1}^{k} \prod_{j=1}^{m} f_{ij}$, where f_{ij} are s-sparse polynomials for each $i \in [k], j \in [m]$. The much more general class of $\mathrm{poly}(n)$-sized and $\mathrm{poly}(n)$ degree algebraic circuits is called VP, which is considered the algebraic analog of complexity class P. The class VNP is considered the algebraic analog of complexity class NP. It is the class of polynomials which can be expressed as an exponential sum of a projection of a VP circuit family.

An algebraic circuit where fan-out of each node is one is called an *algebraic formula*. The class of polynomial sized formulas is called VF.

An *algebraic branching program (ABP)* is a layered directed graph with a unique source and sink vertex. Each edge is directed from one layer to the next with a linear polynomial associated to it as its weight. The weight of a path is the product of edge weights along the path. The ABP then computes the sum of all weighted paths from source to sink, as its output polynomial. The *length* of an ABP is the length of the longest path from source to sink and *width* of an ABP is the maximum number of vertices in any layer. The *size* of ABP is the product of its length and width. The class of all polynomial sized ABPs is called VBP.

A.2 PIT-Preliminaries

The problem of PIT asks for determining whether a given input polynomial is identically zero or not. The input polynomial is given in the form of some algebraic circuit. In *white-box* PIT, one can look 'inside' the input circuit while in *black-box* PIT, the input is given as a black-box and one can only evaluate the given circuit on field points. Therefore, in black-box PIT for a class of n-variate polynomials $\mathcal{C}$, we are asked to provide a set $\mathcal{H} \in \mathbb{F}^n$ such that for any non-zero $f \in \mathcal{C}$, there exists at least one point $\boldsymbol{\alpha} \in \mathcal{H}$ such that $f(\boldsymbol{\alpha}) \neq 0$. Such a set $\mathcal{H}$ is called *hitting-set* for class $\mathcal{C}$. In general, we always have the following brute-force hitting-set, which is efficient when n is small.

Lemma A.1 ([Alo99])**.** *Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be a polynomial with $\deg_{x_i}(f) \leq d_i$, for $i \in [n]$. Let $W_i \subseteq \mathbb{F}$ be a set of size at least $d_i + 1$. If $f \not\equiv 0$, then there exists $\mathbf{a} \in W_1 \times \ldots \times W_n$ such that $f(\mathbf{a}) \neq 0$.*

There is also a notion of hitting set generator (HSG) or simply generator in short, which is equivalent to a hitting set and is easier to work with PIT algorithms. We frame the PIT result in this work using generators. Definition 4.1 gives the formal definition of a generator. The following lemma shows how to obtain hitting set from a generator. It basically follows from Lemma A.1.

Lemma A.2 (Generator $\implies$ hitting-set, [SV15])**.** *Let $\mathcal{G} = (\mathcal{G}_1, \ldots, \mathcal{G}_n) : \mathbb{F}^k \to \mathbb{F}^n$ be a generator for a circuit class $\mathcal{C}$ such that $\deg(\mathcal{G}) \triangleq D$. Let $W \subseteq \mathbb{F}$ be any set of size ndD. Then, $H \triangleq \mathcal{G}(W^k)$ is a hitting set, of size $|H| \leq (ndD)^k$, for polynomials $f \in \mathcal{C}$ of individual degrees $< d$.*

Below, we mention the sparse PIT map of [KS01] which gives efficient deterministic black-box PIT for the class of sparse polynomials and few other folklore results.

Lemma A.3 (Sparse HSG, [KS01])**.** *Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be a non-zero polynomial of individual degree at most d, such that $\|f\| \leq m$. Let p be a prime larger than $\max(d, mn + 1)$. Then, there exists a $k \in [mn + 1]$ such that the univariate polynomial $f'(y) = f(y, y^{k^1 \bmod p}, \ldots, y^{k^{n-1} \bmod p})$ is non-zero. This yields a generator $\mathcal{G} : \mathbb{F} \to \mathbb{F}^n$ of seed-length 1 for the class of m-sparse polynomials such that $\deg(\mathcal{G}) = \mathrm{poly}(m, n, d)$.*

Generator for a class of a polynomials is also a generator for the product of polynomials from that class, since the generator will hit each polynomial in the product.

Lemma A.4 ([SV15])**.** *Let $f = \prod_{i=1}^{k} f_i$ be a non-zero polynomial, where for each $i \in [k]$, $f_i \in \mathcal{C}$, for some circuit class $\mathcal{C}$. Let $\mathcal{G}$ be a generator for $\mathcal{C}$. Then, $f(\mathcal{G}) \not\equiv 0$.*

Generator for a class of polynomials is also generator for the factors of a polynomial in that class. This is because the polynomial ring is an integral domain.

Lemma A.5 (Folklore)**.** *Let $f \in \mathcal{C}$ be a non-zero polynomial in circuit class $\mathcal{C}$. Let $\mathcal{G}$ be a generator for $\mathcal{C}$. If f has a non-zero factor g, then $g(\mathcal{G}) \not\equiv 0$.*

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il