

Deterministically Testing Sparse Polynomial Identities of Unbounded Degree*

Markus Bläser[†] Moritz Hardt[‡] Richard J. Lipton[§] Nisheeth K. Vishnoi[¶]

May 1, 2008

Abstract

We present two deterministic algorithms for the arithmetic circuit identity testing problem. The running time of our algorithms is polynomially bounded in s and m , where s is the size of the circuit and m is an upper bound on the number monomials with non-zero coefficients in its standard representation. The running time of our algorithms also has a logarithmic dependence on the degree of the polynomial but, since a circuit of size s can only compute polynomials of degree at most 2^s , the running time does not depend on its degree. Before this work, all such deterministic algorithms had a polynomial dependence on the degree and henceforth an exponential dependence on s .

Our first algorithm works over the integers and it requires only black-box access to the given circuit. Though this algorithm is quite simple, the analysis of why it works relies on Linnik's Theorem, a deep result from number theory about the size of the smallest prime in an arithmetic progression. Our second algorithm, unlike the first, uses elementary arguments and works over any integral domains, but it uses the circuit in a non-black-box manner. In both cases the running time has a logarithmic dependence on the largest coefficient of the polynomial.

1 Introduction

Testing polynomial identities is a fundamental problem in the theory of computation. Formally, the problem is equivalent to testing whether a multivariate polynomial $f(x_1, \dots, x_n)$ is identically zero over a specified domain such as integers or a finite field. Many algorithmic problems such as finding perfect matchings in graphs [CRS95, Lov79, MVV87], checking the equivalence of read-once branching programs [BCW80], multi-set equality testing [BK95] and primality testing [AB03] reduce to this problem. Equally many applications have been discovered in complexity theory; to only mention two celebrated results, both PSPACE=IP [Sha92, LFKN92] and NP=PCP [AS98, ALM⁺98] involved polynomial identity testing as a subroutine.

The earliest work in testing polynomial identities is due to DeMillo and Lipton [DL78], Zippel [Zip79] and Schwartz [Sch80]. These results give efficient randomized algorithms for checking polynomial identities by showing how to efficiently, using randomness, sample from of a set of $(2d)^n$ points such that no n -variate polynomial of total degree d vanishes on more than half of

*This paper builds upon the conference paper [LV03] by two of the authors.

[†]Saarland University, Saarbrücken, Germany. mblaeser@cs.uni-sb.de

[‡]Princeton University, Princeton, NJ, USA. mhardt@cs.princeton.edu. Research was done while the author visited Carnegie Mellon University hosted by Steven Rudich.

[§]Georgia Institute of Technology, Atlanta, GA, USA. rjl@cc.gatech.edu

[¶]IBM India Research Lab, New Delhi, India. nisheeth.vishnoi@gmail.com

the points in this set. This led to a very important question, open till today, of whether there are efficient deterministic algorithms for checking polynomial identities. More significantly, this problem has arguably become a canonical representative for the more general question of the power of randomization in efficient computation; namely, does $P=BPP$?

To make the complexity of the polynomial identity testing problem precise, one needs to fix a representation of the polynomial which has to be tested for being identically zero or not. If the polynomial is given as a list of its coefficients, the problem is trivial. In the interesting case, one is given an implicit representation of the polynomial such as the determinant of a matrix or an arithmetic circuit computing it. We say an algorithm is given black-box access, if it is restricted to using the input representation as an oracle only, i.e., it may only query the value of the represented polynomial at specific points. In the case of black-box access, one can show that randomness is necessary for polynomial identity testing. When it comes to arithmetic circuits, no such lower bounds are known. However, we only know of deterministic arithmetic circuit identity tests when we have an additional “promise” about the circuit or the polynomial that it represents.

A natural and well-studied promise is that the polynomial is sparse: the number of monomials in it with non-zero coefficients is upper bounded by a parameter m . This parameter, which could be exponential in the number of variables in general, would be a measure of sparsity of the polynomial and we allow the runtime of our identity tests to depend on it. Also, in this paper we will consider polynomials represented by arithmetic circuits. Note that an arithmetic circuit of size s can compute a polynomial of degree at most 2^s even it computes a sparse polynomial. We present two deterministic algorithms for the identity testing problem whose running time is a polynomial in s and m . Crucially, the running time of our algorithms is logarithmic in the degree of the polynomial and, hence by the discussion above, is independent of the degree. Before this work, all such deterministic algorithms had a polynomial dependence on the degree.

1.1 Related Work

Interpolating sparse polynomials. Originally, Zippel [Zip79] was interested in Sparse Polynomial Interpolation, a more difficult question that asks to recover the entire list of coefficients of the given polynomial. Clearly, the sparsity of the polynomial is crucial in order to obtain efficient algorithms for this problem. Zippel’s work was followed by many results [BO88, GKS90, CDGK91, GKS94, Wer94] in interpolation. These results essentially imply polynomial identity tests with a runtime polynomial in n , m and d where n denotes the number of variables of the input polynomial, m denotes its number of nonzero terms and d its degree. There is also some dependence on the characteristic of the underlying field and the bit size of the points queried varies. A survey of these results can be found in [KS01].

Randomized complexity. A sequence of results [CK97, LV98, AB03, KS01, Bog05, BHS08] has studied (and determined in several cases) how much randomness is required in testing polynomial identities when only black-box access to the polynomial is granted. Most of these results are not concerned with sparse polynomials in our sense, but Klivans and Spielman [KS01] give an upper bound of $O(\log(mnd))$ random bits where m is the number of nonzero terms of the given polynomial, d is the degree and n is the number of variables. Trivial deterministic simulation of this algorithm yields a runtime polynomial in m , n and d . However, as before, the polynomial dependence on d implies an exponential worst-case running time for arithmetic circuits.

In the case of finite fields, Karpinski and Shparlinski [KS96] give a randomized and a deterministic algorithm with a runtime polynomial in m , n and $\log d$. The main issue here is that their deterministic algorithm has a linear dependence on the characteristic of the field. This makes their

algorithm inefficient in the case of finite fields of large prime order. Their randomized algorithm has a better dependence on the characteristic. However, it requires a large number of random bits, namely polynomial in $m, n, \log d$. For comparison, notice this is exponentially more than [KS01].

Arithmetic circuits. In recent years, there has been progress on deterministic identity testing for depth-3 circuits of constant top fan-in [KS06, DS07, Shp07]. We can represent a polynomial with m terms by an arithmetic circuit of size $m + 1$ and depth 2. Hence, depth-3 circuits can be seen as a generalization of sparse polynomials. Previously, Agrawal [Agr05] gave a “hitting set generator” against the class of depth-2 arithmetic circuits. This result implies a deterministic polynomial identity test for sparse polynomials. However, all of these results require the input degree to be polynomially bounded.

When it comes to general circuits, no deterministic sub-exponential time arithmetic circuit identity test is known. In fact, Kabanets and Impagliazzo [KI04] show that such an algorithm would imply non-trivial circuit lower bounds. Further connections along these lines were shown by [Agr05].

1.2 Our Results

We state our two results here informally and give a brief overview of the proof techniques used. For a multivariate polynomial $f(x_1, \dots, x_n)$, let m be an upper bound on the number of monomials in it with non-zero coefficients, d be its maximum degree, and H be the largest coefficient in its standard representation (or the size of the field in the case of finite fields).

Theorem 1 (Black-Box) *Given a multivariate polynomial $f(x_1, \dots, x_n)$ with integer coefficients as a black-box, there is a deterministic algorithm which decides whether f is identically zero in time $O((mn \log(d + 1) + m \log H)^{O(1)})$.*

Theorem 2 (Arithmetic Circuit) *Given an arithmetic circuit of size s representing a multivariate polynomial $f(x_1, \dots, x_n)$ over any integral domain, there is a deterministic algorithm which decides whether f is identically zero in time $O((nm \log d)^2 s \log H)$.*

Both the algorithms and their proofs are similar in spirit and we now briefly illustrate the key ingredients and the places where they differ. The formal description and proofs appear in Sections 3 and 4.

The first step in both the algorithms is to “convert” the multivariate polynomial $f(x_1, \dots, x_n)$ to a univariate polynomial $g(X)$. This involves just a computation of a transformation from a query point (a_1, \dots, a_n) for f to a query point a for g . This is standard and can be done efficiently even when one is given black-box access to the polynomial. The key point about this transformation is that $f(a_1, \dots, a_n)$ is non-zero if and only if $g(a)$ is non-zero. For such a property to hold the degree of g must be at least $(d + 1)^n - 1$ and, indeed, there is a transformation, namely $x_i \mapsto X^{(d+1)^{i-1}}$, which ensures that the degree of g is at most $d \sum_{i=1}^{n-1} (d + 1)^{i-1} = (d + 1)^n - 1$. Henceforth, we may assume that we are given a univariate polynomial. In this setting, our first algorithm can be summarized as follows.

Given black-box access to $g(X)$ with integer coefficients, declare ‘ g is identically zero’ if and only if g is zero at all points modulo p for sufficiently many prime numbers p .

To analyze this algorithm we need to show that there is a way to pick these primes so that a non-zero polynomial $g(X)$ does not vanish modulo many of them. There are two ways in which a

non-zero $g(X) = \sum_{i=0}^D c_i X^i$ can vanish modulo a prime p . It may happen that some non-zero c_i is divisible by p . Since the number of distinct prime divisors of any integer can be at most logarithmic in its magnitude, the number of distinct prime divisors of all the coefficients is bounded $m \log H$. More interestingly, it may happen that some $X^i \equiv X^j$ where X takes value in $\{0, 1, \dots, p-1\}$. This, by Fermat's Little Theorem, implies that $(p-1)|(i-j)$. In general it is difficult to bound the number of primes p such that $p-1$ divides an integer, but here is where we appeal to Linnik's Theorem about the smallest prime in an arithmetic progression.

We pick p carefully from an arithmetic progression $\{kq+1|k \geq 1\}$, where q is also a prime. Now, whenever $(p-1)|(i-j)$, q serves as a prime factor for $(i-j)$ and this essentially allows us to upper bound the number of primes which will be "bad" for the algorithm mentioned above. The key questions that remain unanswered are why should there exist primes in arbitrary arithmetic progressions and why should we be able to find them efficiently? Dirichlet generalized the Prime Number Theorem to prove that every arithmetic progression in which the terms have gcd 1 must have the "right density" of primes. But his theorem does not say anything about how to efficiently find such a prime. This is implied by Linnik's Theorem which shows that the smallest prime in an arithmetic progression with difference q is bounded by $q^{O(1)}$. Hence, modulo some details, to complete the proof of Theorem 1 it suffices to pick sufficiently many primes from arithmetic progressions whose difference itself is a prime number. From the discussion above $(mn \log(d+1) + m \log H)^{O(1)}$ such primes suffice.

The second algorithm observes that if one is allowed non-black-box access to the polynomial $g(X)$, it is no longer necessary to pick the primes in a clever manner as above. It also has the advantages that it works over any integral domain and it does not need hard

Given a circuit computing $g(X)$, declare ' g computes the zero polynomial' if and only if $g(X)$ vanishes modulo $X^p - 1$ for sufficiently many prime numbers p .

It is easy to see that in this case, the only way a non-zero polynomial will become zero after reducing it modulo $X^p - 1$ is when $p|(i-j)$. Since, as above, the degree of g is upper bounded by $(d+1)^n$, picking $O(mn \log d)$ prime numbers suffices. The key difference from the previous algorithm is that checking whether $g(X) \equiv 0 \pmod{(X^p - 1)}$ requires to "look" inside the circuit. We don't know how to do this in a black-box setting. Also, this gives rise to a running time which depends on the logarithm of the largest coefficient in the case of infinite fields (and the logarithm of the size of the field in the finite field case).

A virtue of our algorithms is their conceptual simplicity. Furthermore, both algorithms have a logarithmic runtime dependence on the degree of the input polynomial. In the arithmetic circuit model, this implies that we do not need to place any restriction on the degree of the input. Even though sparse polynomials have been extensively studied, the natural case of unbounded degrees in the arithmetic circuit model has not been treated previously. Now we proceed to formal proofs. We start with some preliminaries in Section 2.

2 Preliminaries

In this paper, K denotes an integral domain. We will refer to the *degree* of a polynomial $f \in K[x_1, \dots, x_n]$ as is its maximum degree, that is, the smallest integer d such that the exponent of every variable in f is bounded by d .

We consider division-free arithmetic circuits over K with fan-in 2 at each multiplication and

addition gate. Multiplication and addition by constants is allowed. The size of a circuit is the number of multiplication and addition gates. Over fields of characteristic zero, we add the logarithm of the magnitude of each constant to the size of the circuit. We denote the class of such arithmetic circuits of size s by $\mathcal{C}(s)$. We let $\mathcal{C}(s, m, d)$ denote the restriction of $\mathcal{C}(s)$ to circuits that compute polynomials with at most m nonzero terms and degree bounded by d . We say an algorithm is given “black-box” access to a circuit if the algorithm is not given the circuit itself but, instead, it is given an oracle which on input of a specific point will output the value of the circuit at that point. In the case of circuits with integer coefficients, the algorithm may request the value of the circuit modulo a positive integer. We use the notation $O^\sim(t)$ to suppress poly-logarithmic factors of t , i.e., $O^\sim(t) = O(t \cdot \log^{O(1)} t)$.

Multivariate to univariate reductions. As remarked earlier, our algorithms will be designed for univariate polynomials and extended to the multivariate case using the following (folk) substitution attributed to Kronecker.

Lemma 1 *Let $f \in K[x_1, \dots, x_n]$ be a polynomial of degree at most d . Then the substitution $x_i \mapsto X^{(d+1)^{i-1}}$ maps f to a univariate polynomial $g \in K[X]$ of degree at most $(d+1)^n$ such that any two distinct monomials w and w' in f map to distinct monomials in g . In particular, if f is not identically zero in $K[x_1, \dots, x_n]$, then g is not identically zero in $K[X]$.*

2.1 Number-Theoretic Preliminaries

Fact 2 *An integer $k \geq 1$ has at most $\log k$ distinct prime divisors.*

Proof. Let p_1, \dots, p_t be the distinct prime divisors of k . We have $k \geq \prod_{i=1}^t p_i \geq 2^t$. Hence, $t \leq \log k$. \square

We will need the following variant of the Prime Number Theorem.

Theorem 3 *The k -th prime number is of order $\Theta(k \log k)$.*

Let $p(k)$ denote the smallest prime number in the arithmetic progression $\{jk + 1 \mid j \geq 1\}$. Linnik’s Theorem gives an *unconditional* upper bound on $p(k)$.

Theorem 4 (Linnik’s Theorem) *There is a constant $L > 1$ (called Linnik’s constant) such that $p(k) < k^L$ for every sufficiently large $k \geq k_0$.*

The best known value for L is 5.5 due to Heath and Brown. Schinzel, Sierpinski, and Kanold have conjectured the value to be 2. For a detailed discussion of these facts the reader is referred to the book by Ribenboim [Rib96].

We are interested in the value of $p(q)$ where q itself is a prime number. Notice, it is not immediately clear whether or not two distinct primes $q_1 \neq q_2$, could have $p(q_1) = p(q_2)$. However, we can show that no more than L primes can map to the same smallest prime. In particular, if we take a set of $5t$ distinct primes $\{q_i \mid 1 \leq i \leq 5t\}$, then the set $\{p(q_i) \mid 1 \leq i \leq 5t\}$ has cardinality at least t .

Lemma 3 *Let $q_1 < \dots < q_v$ and p be primes such that $p(q_i) = p$ for all $i \in \{1, \dots, v\}$ and suppose $q_1 > k_0$ where k_0 is the constant from Linnik’s Theorem above. Then, $v < L$.*

Proof. By hypothesis, there is a positive integer r such that $p - 1 = rq_1q_2 \dots q_v$. Hence, $q_1^v \leq \prod_{i=1}^v q_i < p$. As p is the smallest prime in the arithmetic progression $\{1 + kq_1 \mid k \geq 1\}$, Linnik’s Theorem implies $p < q_1^L$. Hence, $q_1^v < q_1^L$ and thus $v < L$. \square

3 A Black-Box Polynomial Identity Test Over the Integers

In this section we prove Theorem 1. The key is the following lemma.

Lemma 4 *Let $f = \sum_{i=0}^d c_i X^i$ be a univariate polynomial such that each coefficient $c_i \in \mathbb{Z}$ with $|c_i| \leq H$. Suppose that at least one, but no more than m coefficients of f are nonzero. Then, there are less than $5 \log H + m \log d$ distinct prime numbers q such that f vanishes on every point modulo $p(q)$.*

Proof. Let i be the smallest index such that the coefficient $c_i \neq 0$.

Claim 5. The set $\{q \text{ prime} \mid c_i \equiv 0 \pmod{p(q)}\}$ has cardinality at most $5 \log H$.

This claim follows directly from Lemma 3 once we observe there are at most $\log c_i \leq \log H$ prime numbers q such that $c_i \equiv 0 \pmod{q}$.

Claim 6. The set $\{q \text{ prime} \mid X^j \equiv X^i \pmod{p(q)} \text{ for some } j > i\}$ has cardinality less than $m \log d$.

In order to prove Claim 6, suppose there is an index $j > i$ such that $X^i \equiv X^j \pmod{p(q)}$. Equivalently, $(p(q) - 1) \mid (j - i)$ (by Fermat's Little Theorem). Since $p(q)$ lies on the arithmetic progression $\{1 + kq \mid k > 0\}$, we conclude $q \mid (j - i)$. However, $j - i \leq d$ and hence $j - i$ has at most $\log d$ prime divisors. On the other hand, there are at most $m - 1$ choices for $j > i$. Our second claim is proven.

It remains to note, whenever f vanishes modulo a prime number q , then q is contained at least one of the two sets above. \square

Algorithm A

Input: Black-box access to a circuit C on n inputs, parameters $R, d \in \mathbb{N}$.

Output: Whether or not C computes the identically zero polynomial over \mathbb{Z} .

1. Determine the first R distinct prime numbers p_1, \dots, p_R .
2. For all $p \in \{p_1, \dots, p_R\}$ and all $k \in \{0, \dots, p-1\}$, query the black-box at the point (r_1, \dots, r_n) where r_j denotes the remainder of $k^{(d+1)^{j-1}}$ modulo p .
3. Output “ C is zero” if all queries returned are zero, otherwise output “ C is nonzero”.

Theorem 5 *Given black-box access to a circuit $C \in \mathcal{C}(s, m, d)$ over \mathbb{Z} , we can decide whether C computes the identically zero polynomial in time polynomial in $ms \log(d + 1) + \log H$ where H is an upper bound on the absolute value of each coefficient computed by C . Moreover the bit lengths of the queries to the black-box are logarithmic in $m, s, \log d$ and $\log H$.*

Proof. Suppose the given circuit C does not compute the identically zero polynomial over \mathbb{Z} . We claim that Algorithm A will reject C for the parameter setting $R = (5 \log H + mn \log d)^{5.5}$. Indeed, for this choice of R it follows from Lemma 4 that there exists a prime number p among the first R primes such that $C'(X)$ does not vanish over $\mathbb{Z}/p\mathbb{Z}$. Here, $C'(X)$ denotes the circuit obtained from C by applying the Kronecker substitution (Lemma 1). Specifically, there exists a $k \in \{0, \dots, p-1\}$ such that $C'(k) \not\equiv 0 \pmod{p}$. Hence, the query computed by Algorithm A for the given choice of p and k will return a nonzero value.

The claimed runtime follows easily. The largest point queried has bit size $\log p_R$ which is logarithmic in $m, n, \log d$ and $\log H$. We conclude by observing that $n \leq s$. \square

Remark 1 The analysis of the previous algorithm works even if instead of Linnik's Theorem we consider a bound on the average size of the smallest prime in an arithmetic progression. Baker and Harman [BH96] showed that Linnik's constant is less than 1.93 on average. This leads to significant improvements in the runtime of our algorithm by polynomial factors. We omit the details.

4 Testing Sparse Polynomial Identities in the Non-Black-Box Setting

In this section we prove Theorem 2. Unlike in the black-box case, the following easy lemma turns out to be sufficient.

Lemma 7 *Let $f \in K[X]$ be a nonzero univariate polynomial with at most m nonzero terms and degree bounded by d . Then there are less than $m \log d$ prime numbers p for which $f(X) \equiv 0 \pmod{X^p - 1}$.*

Proof. Let $f(X) = \sum_{i=0}^d c_i X^i$. Fix the smallest index i for which $c_i \neq 0$. If $f(X)$ vanishes modulo $X^p - 1$, then there must exist an index $j > i$ with $c_j \neq 0$ such that $X^j \equiv X^i \pmod{X^p - 1}$. Equivalently, $p | (j - i)$. But, $0 < j - i \leq d$ and therefore $j - i$ has at most $\log d$ prime divisors (Fact 2). On the other hand, we only have $m - 1$ different choices for the index j . \square

Testing if a polynomial $f(X)$ is identically zero modulo $X^k - 1$ is a standard operation that can be carried out when f is represented by an arithmetic circuit. Given the circuit, we can recursively compute f in its reduced form. That is, at each gate we carry out the given arithmetic operation and reduce the resulting polynomial modulo $X^k - 1$. In this fashion, we maintain a polynomial of degree at most $k - 1$. Reducing the polynomial gate-wise is justified, since we are computing a ring homomorphism. Multiplying two degree $k - 1$ polynomials can be done at the cost of $O^\sim(k)$ ring operations over any integral domain using Schönhage-Strassen multiplication. Hence, the entire operation requires $O^\sim(sk)$ ring operations where s is the size of the given circuit. For details the reader is referred to the book by v. z. Gathen and Gerhard [GG03].

Our algorithm is stated and analyzed next. It will perform the Kronecker substitution so as to apply to multivariate polynomials as well.

Algorithm B

Input: Arithmetic circuit C on n inputs, parameters $R, d \in \mathbb{N}$.

Output: Whether or not C computes the identically zero polynomial.

1. Determine R distinct prime numbers p_1, \dots, p_R .
2. For each $i \in \{1, \dots, R\}$:
 - (a) Construct the circuit $C_i(X) = C(X^{r_1}, X^{r_2}, \dots, X^{r_n})$ where r_j denotes the remainder of $(d + 1)^{j-1}$ modulo p_i .
 - (b) Test if $C_i(X) \equiv 0 \pmod{X^{p_i} - 1}$.
3. Output “ C is zero” if (2b) holds for all $i \in \{1, \dots, R\}$, otherwise output “ C is nonzero”.

Runtime Analysis Step (1) of our algorithm requires time $O^\sim(R)$ using a deterministic polynomial time primality test [AKS04]. Computing the circuit C_i requires us to compute the remainder of $d + 1$ modulo p . This can be done in time $O(\log d \log p)$ [GG03]. The remaining values $(d + 1)^{j-1}$ modulo p for $1 \leq j \leq n$ can be computed with less than n multiplications modulo p . Since $p = O^\sim(R)$ by Theorem 3, Step (2a) costs time $O^\sim(Rn \log d)$. If s denotes the size of C , then the size of C_i need not be larger than $s' = s + O(n \log R)$. Hence, the test $C_i(X) \equiv 0 \pmod{X^{p_i} - 1}$ in (2b) can be performed with $O^\sim(Rs')$ arithmetic operations as discussed earlier. This results in the over all count of $O^\sim(R^2(s + n \log R)) = O^\sim(R^2s)$ arithmetic operations.

Theorem 6 *Given an arithmetic circuit $C \in \mathcal{C}(s, m, d)$ on n inputs, we can deterministically test if C computes the identically zero polynomial at the cost of $O((nm \log d)^2 s)$ ring operations over any integral domain.*

Proof. We claim that given a circuit $C \in \mathcal{C}(s, m, d)$ and the parameter setting $R = O(mn \log d)$, Algorithm B accepts C if and only if C computes the identically zero polynomial. Clearly, the algorithm accepts C , if it computes the identically zero polynomial. On the other hand, suppose C does not compute the identically zero polynomial. Let $C'(X) = C(X, X^{d+1}, \dots, X^{(d+1)^{n-1}})$. By Lemma 1, C' computes a nonzero univariate polynomial with at most m nonzero terms and degree $(d+1)^n$. Hence, by Lemma 7, there are less than $R = O(mn \log d)$ prime numbers p for which $C'(X) \equiv 0 \pmod{X^p - 1}$. It remains to observe, for all $i \in \{1, \dots, R\}$, the circuits C' and C_i compute identical polynomials modulo $X^{p_i} - 1$. The runtime of our algorithm is dominated by the number of arithmetic operations for the given setting of R . \square

Notice, over fields of characteristic zero, the cost of an arithmetic operation depends on the size of the coefficients of the polynomial. For instance, we obtain the runtime $O((nm \log d)^2 s \log H)$ over the field of rational numbers where H is an upper bound on the largest coefficient. Over a finite field of order q , the runtime will be $O((nm \log d)^2 s \log q)$.

Remark 2 It would be interesting to know whether our analysis can be improved. The authors are not aware of a polynomial f computed by a constant depth circuit for which the univariate polynomial $f(X, X^{d+1}, \dots, X^{(d+1)^{n-1}})$ vanishes modulo $X^p - 1$ for more than a polynomial (in the size of the circuit) number of primes p . In fact, Agrawal [Agr05] conjectured that a similar test does work for any constant depth circuit.

References

- [AB03] Manindra Agrawal and Somenath Biswas. Primality and identity testing via chinese remaindering. *J. ACM*, 50(4):429–443, 2003.
- [Agr05] Manindra Agrawal. Proving lower bounds via pseudo-random generators. In *Proc. 25th FSTTCS*, pages 92–105. Springer, 2005.
- [AKS04] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. PRIMES is in P. *Ann. of Math.* (2), 160(2):781–793, 2004.
- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.
- [AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: a new characterization of NP. *J. ACM*, 45(1):70–122, 1998.
- [BCW80] Manuel Blum, Ashok K. Chandra, and Mark N. Wegman. Equivalence of free boolean graphs can be decided probabilistically in polynomial time. *IPL*, 10:80–82, 1980.
- [BH96] R. C. Baker and G. Harman. The Brun-Titchmarsh theorem on average. In *Analytic number theory, Vol. 1 (Allerton Park, IL, 1995)*, volume 138 of *Progr. Math.*, pages 39–103. Birkhäuser Boston, Boston, MA, 1996.

- [BHS08] Markus Bläser, Moritz Hardt, and David Steurer. Asymptotically optimal hitting sets against polynomials. In *Proc. 35th ICALP*. Springer, 2008.
- [BK95] Manuel Blum and Sampath Kannan. Designing programs that check their work. *J. ACM*, 42(1):269–291, 1995.
- [BO88] Michael Ben-Or. A deterministic algorithm for sparse multivariate polynomial interpolation. In *Proc. 29th STOC*, pages 301–309. ACM, 1988.
- [Bog05] Andrej Bogdanov. Pseudorandom generators for low degree polynomials. In *Proc. 37th STOC*, pages 21–30. ACM, 2005.
- [CDGK91] Michael Clausen, Andreas Dress, Johannes Grabmeier, and Marek Karpinski. On zero-testing and interpolation of k -sparse multivariate polynomials over finite fields. *Theor. Comput. Sci.*, 84(2):151–164, 1991.
- [CK97] Zhi-Zhong Chen and Ming-Yang Kao. Reducing randomness via irrational numbers. In *Proc. 29th STOC*, pages 200–209. ACM, 1997.
- [CRS95] Suresh Chari, Pankaj Rohatgi, and Aravind Srinivasan. Randomness-optimal unique element isolation with applications to perfect matching and related problems. *SIAM J. Comput.*, 24(5):1036–1050, 1995.
- [DL78] Richard A. DeMillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *IPL*, 7, 1978.
- [DS07] Zeev Dvir and Amir Shpilka. Locally decodable codes with two queries and polynomial identity testing for depth 3 circuits. *SIAM J. Comput.*, 36(5):1404–1434, 2007.
- [GG03] Joachim Von Zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. Cambridge University Press, New York, NY, USA, 2003.
- [GKS90] Dima Yu Grigoriev, Marek Karpinski, and Michael F. Singer. Fast parallel algorithms for sparse multivariate polynomial interpolation over finite fields. *SIAM J. Comput.*, 19(6):1059–1063, 1990.
- [GKS94] Dima Grigoriev, Marek Karpinski, and Michael F. Singer. Computational complexity of sparse rational interpolation. *SIAM J. Comput.*, 23(1):1–12, 1994.
- [KI04] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Comput. Complex.*, 13(1/2):1–46, 2004.
- [KS96] Marek Karpinski and Igor Shparlinski. On some approximation problems concerning sparse polynomials over finite fields. *Theor. Comput. Sci.*, 157(2):259–266, 1996.
- [KS01] Adam Klivans and Daniel A. Spielman. Randomness efficient identity testing of multivariate polynomials. In *Proc. 33th STOC*, pages 216–223. ACM, 2001.
- [KS06] Neeraj Kayal and Nitin Saxena. Polynomial identity testing for depth 3 circuits. In *Proc. 21st CCC*, pages 9–17. IEEE, 2006.
- [LFKN92] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *J. ACM*, 39(4):859–868, 1992.

- [Lov79] László Lovász. On determinants, matchings, and random algorithms. In *FCT*, pages 565–574, 1979.
- [LV98] Daniel Lewin and Salil Vadhan. Checking polynomial identities over any field: Towards a derandomization? In *Proc. 30th STOC*, pages 438–437. ACM, 1998.
- [LV03] Richard Lipton and Nisheeth Vishnoi. Deterministic identity testing for multivariate polynomials. In *Proc. SODA*, pages 756–760. ACM, 2003.
- [MVV87] Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. Matching is as easy as matrix inversion. *Combinatorica*, 7(1):105–113, 1987.
- [Rib96] Paulo Ribenboim. *The New Book of Prime Number Records*. Springer, New York, 1996.
- [Sch80] Jacob Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27:701–717, 1980.
- [Sha92] Adi Shamir. $IP = PSPACE$. *J. ACM*, 39(4):869–877, 1992.
- [Shp07] Amir Shpilka. Interpolation of depth-3 arithmetic circuits with two multiplication gates. In *Proc. 39th STOC*, pages 284–293. ACM, 2007.
- [Wer94] Kai Werther. The complexity of sparse polynomial interpolation over finite fields. *Applicable Algebra in Engineering, Communication and Computing*, 5:91–103, 1994.
- [Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *Proc. ISSAC*, pages 216–226, Berlin, 1979. Springer-Verlag.