

Efficient Approximation Algorithms for Sparse Polynomials over Finite Fields

Marek Karpinski¹

Igor Shparlinski²

TR-94-029

August, 1994

Abstract

We obtain new lower bounds on the number of non zeros of sparse polynomials and give a fully polynomial time (ϵ, δ) approximation algorithm for the number of non-zeros of multivariate sparse polynomials over a finite field of q elements and degree less than $q - 1$. This answers partially to an open problem of D. Grigoriev and M. Karpinski. Also, probabilistic and deterministic algorithms for testing identity to zero of a sparse polynomial given by a "black-box" are given. Finally, we propose an algorithm to estimate the size of the image of a univariate sparse polynomial.

¹Department of Computer Science, University of Bonn, 53117 Bonn, Email marek@cs.uni-bonn.de, and the International Computer Science Institute, Berkeley, California. Supported in part by Leibniz Center for Research in Computer Science, by the DFG Grant KA 673/4-1, and the ESPRIT BR Grant 7097.

²School of MPCE, Macquarie University, Sydney, NSW 2109, Australia, Email igor@mpce.mq.edu.au

Introduction

In the recent paper [3] (improving and generalizing some previous results of [4], [6], [9]) lower bounds have been obtained for the number of zeros and non-zeros of a t -sparse multivariate polynomial over a finite field \mathbb{F}_q of q elements. As it was mentioned in [3] there is no real chance to improve essentially the lower bound for the number of zeros (as it would imply a randomized subexponential algorithm for the famous 3-SAT problem) but an analogous question for the number of non zeros was posed as an open problem.

Here we show that indeed for a very wide class of polynomials a lower bound of the conjectured in [3] type holds. We show that, roughly speaking, the density of non-zeros is at least t^{-1} rather than $t^{-\log q}$ as in [3]. Note that for arbitrary polynomials (i.e. when degrees up to $q - 1$ are allowed) the mentioned lower bound $t^{-\log q}$ cannot be possibly sharpened (see the remark after Theorem 2 of [3]). Of course the obtained improvement immediately leads to an improved Monte-Carlo approximation algorithm for the number of non-zeros of a t -sparse polynomial and to an RNC-algorithm for testing identity to zero of a sparse polynomials given by a "black-box". It gives a polynomial time approximation algorithm even in case of growing q (the previous one needs q to be fixed as contains $\log q$ in the exponent).

Then we consider a related question about the zero-testing of t -sparse multivariate polynomials over \mathbb{F}_q in the black-box model of [2], [5], [16]. For the case on a non-prime field we obtain several improvements of previously known results. It is hoped they can be applied to the more general problem of polynomial interpolation.

Finally we show that in some cases, the image size of a univariate t -sparse polynomials can be estimated quite quickly. For example, for any $A > 0$ say one can check if it is less than $\log^A q$ in polynomial time.

1 Counting Non Zeros of Sparse Polynomials

Let $f(x_1, \dots, x_m) \in \mathbb{F}_q[x_1, \dots, x_m]$ be a t -sparse polynomial (i.e. a polynomial containing exactly t monomials). Denote by $R(f)$ the number of non-zeros of f over F_q^* that is the number of $(a_1, \dots, a_m) \in [\mathbb{F}_q^*]^m$ such that $f(a_1, \dots, a_m) \neq 0$.

Theorem 1 *Let f be a t -sparse polynomial in m variables over \mathbb{F}_q with $t \geq 1$ and $\deg f \leq q - 2$ then $R(f) \geq (q - 1)^m / t$.*

Proof. Let us use induction in m . For $m = 1$, let us consider a polynomial

$$f(x) = \sum_{i=1}^t c_i x^{n_i},$$

with $c_i \in F_q^*$, $i = 1, \dots, t$, and $0 \leq n_1 < \dots < n_t \leq q - 2$. Let θ be a primitive root of F_q . Then $R(f)$ is the number of $u = 0, \dots, q - 2$ such that $f(\theta^u) \neq 0$. Evidently, it is enough to show that for any integer h , among the following t elements

$$f(\theta^u), \quad u = h, \dots, h + t - 1,$$

there exists at least one non-zero.

Indeed if all of them equal zero then we get that the following system of linear equations

$$\sum_{i=1}^t z_i \theta^{jn_i} = 0, \quad j = 0, \dots, t - 1,$$

has a non-zero solution $z_i = c_i \theta^{hn_i}$, $i = 1, \dots, t$. On the other hand the system has a Vandermonde matrix with different entries $\theta^{n_1}, \dots, \theta^{n_t}$ (as $n_i \not\equiv n_j \pmod{q-1}$ for $1 \leq i < j \leq t$). The obtained contradiction proves the estimate for $m = 1$.

Now, let us consider the general case. We represent a t -sparse polynomial in the form

$$f(x_1, \dots, x_m) = \sum_{i=1}^s f_i(x_1, \dots, x_{m-1}) x_m^{n_i}$$

where $f_i(x_1, \dots, x_{m-1})$, $i = 1, \dots, s$, are some non-zero polynomials over F_q and $0 \leq n_1 < \dots < n_t \leq q - 2$. It is evident that at least one of the coefficient polynomials is r -sparse with $r \leq t/s$. Therefore, from the induction conjecture we get that there are at least $(q - 1)^{m-1}/r$ vectors $(a_1, \dots, a_{m-1}) \in [F_q^*]^{m-1}$ such that $f(a_1, \dots, a_{m-1}, x_m)$ is a non-zero polynomial thus it is a k sparse polynomial with $1 \leq k \leq s$ and therefore has at least $(q - 1)/k \geq (q - 1)/s$ non-zeros. Thus $R(f) \geq (q - 1)^{m-1} r^{-1} (q - 1) k^{-1} \geq (q - 1)^m t^{-1}$. \square

Now let us consider the total number of non zeros $T(f)$ of f over F_q that is the number of $(a_1, \dots, a_m) \in [F_q]^m$ such that $f(a_1, \dots, a_m) \neq 0$.

As in [3], [8], denote by $G(f)$ the set of $(a_1, \dots, a_m) \in [F_q]^m$ for which at least one of the monomials containing in the representation of f is not equal to zero.

Theorem 2 *Let f be a t -sparse polynomial in m variables over F_q with $t > 0$ and $\deg f \leq q - 2$ then $T(f) \geq |G(f)|/t$.*

Proof. For an n -dimensional $(0, 1)$ vector $\lambda = (\lambda_1, \dots, \lambda_m) \in \{0, 1\}^m$ denote by $G_\lambda(f)$ the subset of $G(f)$ containing vectors having zero coordinates on the same positions where λ has, that is

$$G_\lambda(f) = \{(a_1, \dots, a_m) \in G(f) \mid a_i = 0 \Leftrightarrow \lambda_i = 0, i = 1, \dots, m\}.$$

If $\text{wt}(\lambda)$ denote the Hamming weight of λ (i.e. the number of nonzero coordinates) then either $|G_\lambda(f)| = (q-1)^{\text{wt}(\lambda)}$ or $G_\lambda(f) = \emptyset$.

Let us denote by f_λ the polynomial in $\text{wt}(\lambda)$ variables obtained from f by specialization to zero all variables x_i having the index i , $1 \leq i \leq m$ such that $\lambda_i = 0$. Evidently,

$$T(f) = \sum_{\lambda \in \{0,1\}^m} R(f_\lambda), \quad |G(f)| = \sum_{\lambda \in \{0,1\}^m} |G_\lambda(f)|$$

Now it is enough to prove that $R(f_\lambda) \geq |G_\lambda(f)|/t$ for all $\lambda \in \{0, 1\}^m$. Indeed, if $G_\lambda(f) = \emptyset$ then it is evident, otherwise f_λ is a non-zero s -sparse polynomial with $1 \leq s \leq t$ and thus applying Theorem 1 we get the desired inequality. \square

Theorem 3 *Let f be a t -sparse polynomial in m variables over \mathbb{F}_q with $\deg f \leq q-2$ then for any $\epsilon > 0$ and $\delta > 0$ there exists a randomized algorithm using $O(m \log q)$ random bits and $O(\epsilon^{-2} m t \log q \log(1/\delta))$ arithmetical operations in \mathbb{F}_q and computing an approximation T to $T(f)$ such that*

$$\Pr\{|T - T(f)| < \epsilon T(f)\} > 1 - \delta.$$

Proof. Using the estimate of Theorem 2, and the efficient construction of the set $G(f)$ from [8] one gets the pointed out algorithm (see [6]). \square

Also, as in [8], one can get a parallel version of the last algorithm.

Unfortunately, our restriction on the degree $\deg f \leq q-2$ does not allow us to consider a more interesting and important question about the number of zeros of a polynomial (the standard reduction use an auxiliary polynomial $F = f^{q-1} - 1$, see [4], [3], [8]). Moreover, as we have mentioned, it was shown in [3] that without some other restriction on f their estimate cannot be improved. On the other hand we conjecture that in this special case of polynomial of the shape $f^{q-1} - 1$ the result of [3] is not sharp.

2 Zero Testing of Sparse Polynomials

A t -sparse polynomial f of degree $\deg f \leq d$ in m variables over \mathbb{F}_q is said to be given by a "black-box" if in any point over any extension \mathbb{F}_{q^l} , $l = 1, 2, \dots$, we can compute it in time $(lt \log d \log q)^{O(1)}$ (but we do not know its coefficients). A typical example is a polynomial given by the determinant of a matrix with polynomial entries.

From the definition above, the question about constructing the corresponding extensions \mathbb{F}_{q^l} of \mathbb{F}_q naturally arises. We do not consider this question here in details and in all algorithms below we assume that we are given the corresponding extension but it is easy reformulate all of them in a form taking into account the costs of such construction (without losing the main features of the algorithm).

Indeed, during recent years a very substantial progress in this area has been achieved (for a survey see Chapter 2 of [14]). Say for the field of Theorem 4 below we may use the probabilistic algorithm from [12] with expected number of

$$O(l^2 \log^2 l \log \log l + l \log q \log l \log \log l)$$

arithmetical operations over \mathbb{F}_q , as for Theorem 5 we may apply the deterministic algorithm from [11] using

$$l^4 p^{1/2} (\log l \log q)^{O(1)}$$

arithmetical operations over \mathbb{F}_q where p is the characteristic of \mathbb{F}_q (thus it is a polynomial time algorithm for fields of small characteristic). There are many other fast algorithms as well.

Moreover, as in fact everywhere, we need only to have an extension of degree l exceeding some lower bound L_0 (rather than satisfying the following stronger condition $l = L_0$) we can use an algorithm of [1] that for any L_0 in polynomial time $(L_0 \log q)^{O(1)}$ constructs an extension \mathbb{F}_{q^l} of degree l with

$$L_0 \leq l \leq cL_0 \log q$$

where $c > 0$ is an effectively computable absolute constant.

Theorem 4 *Let f be a t -sparse polynomial in m variables over \mathbb{F}_q with $0 \leq t \leq \tau$ and $\deg f \leq d$ and given by a "black-box". Then for any $\delta > 0$ there is an algorithm for testing identity to zero of f , using $O(m\tau \log q \log(1/\delta))$ random bits, $O(\tau \log(1/\delta))$ parallel evaluations of f over \mathbb{F}_{q^l} with any $l \geq \lceil \log(d+2)/\log q \rceil$ and having the probability of the correct answer at least $1 - \delta$.*

Proof. Choosing $N = \lfloor 4\tau \log(1/\delta) \rfloor + 1$ points from $[\mathbb{F}_{q^l}^*]^m$ at random in parallel and using the estimate of Theorem 1 (as $d \leq q^l$), one gets the desired algorithm (see [6]). \square

Taking into account the remarks about constructing finite fields, in particular, the estimate (2), we get the following probabilistic polynomial-time test.

Corollary 5 *Let f be a t -sparse polynomial in m variables over \mathbb{F}_q with $0 \leq t \leq \tau$ and $\deg f \leq d$ and given by a "black-box" then for any $\delta > 0$ there is an algorithm for testing identity to zero of f , using $(m\tau \log d \log q \log(1/\delta))^{O(1)}$ arithmetical operations over \mathbb{F}_q and having the probability of the correct answer at least $1 - \delta$.*

Note that several deterministic algorithms are known for this problem but for q growing all of them are exponential with respect to q (see [2], [5], [16]). All these algorithms are based on evaluations of the polynomial in several points over some extension \mathbb{F}_{q^l} computed from a primitive root of this field. So, in order to get an effective algorithm we should find a primitive root firstly.

All known (probabilistic and deterministic) algorithms to find a primitive root work in two steps:

- Step 1. Construct a "small" set $\mathfrak{M} \in \mathbb{F}_{q^l}$ containing a primitive root.
- Step 2. Find a primitive root testing for primitiveness every element of M .

Unfortunately, the second step needs the factorization of $q^l - 1$ that is very time consuming for large l . It was the reason why in the previous papers the authors tried to work in slight extensions of \mathbb{F}_q . Here we show that in fact we may drop Step 2 thus some recent results concernig Step 1 enable us to design fast deterministic tests.

In particular, we show below (see Theorem 5) that in the case when q is a power of a fixed prime number (say $q = 2^r$) Theorem 2.3 of [2] leads to a polynomial time algorithm and moreover it enables us to improve the results of [5] with respect to the power of q in the estimate of the number of processors (at for non-prime fields).

Theorem 6 *Let \mathbb{F}_q be of characteristic p . Then for any positive d and τ , in time*

$$p(m\tau \log d \log q)^{O(1)}$$

one can construct a field \mathbb{F}_{q^k} with $k = m \lceil \log(d+2)/\log q \rceil$ and a test-set $U \in [\mathbb{F}_{q^k}]^m$ of size $|U| = p(m \log q)^{O(1)}$ such that a t -sparse polynomial f in m variables over \mathbb{F}_q with $0 \leq t \leq \tau$ and $\deg f \leq d$ is identical to zero if and only if $f(u_1, \dots, u_m) = 0$ for any $(u_1, \dots, u_m) \in \mathfrak{U}$.

Proof. Firstly let us construct the field \mathbb{F}_{q^k} where $k = lm$, $l = \lceil \log(d+2)/\log q \rceil$ by using the algorithm of [12], in time

$$O\left(p^{1/2}(m \log d \log q)^{O(1)}\right)$$

(see (1)). Then the algorithms of [13] and of [15] (see also Chapter 2 of [14]), in time $p(m \log d \log q)^{O(1)}$ construct a set $\mathfrak{M} \in F_{q^k}$ of size $|\mathfrak{M}| = p(m \log d \log q)^{O(1)}$ and containing a primitive root θ of F_{q^k} .

It follows from Theorem 2.3 of [2] that f is identical to zero if and only if $f(0, \dots, 0) = 0$ and $f(\theta^i, \theta^{iq^l}, \dots, \theta^{iq^{l(m-1)}}) = 0$, $i = 0, \dots, \tau - 1$. Defining

$$\mathfrak{U} = \{(0, \dots, 0)\} \cup \{(\mu^i, \mu^{iq^l}, \dots, \mu^{iq^{l(m-1)}}) \mid \mu \in \mathfrak{M}, 0 \leq i \leq \tau - 1\}$$

we get the desired set. □

Corollary 7 *Let \mathbb{F}_q be a finite field of characteristic p and let f be a t -sparse polynomial in m variables over \mathbb{F}_q with $0 \leq t \leq \tau$ and $\deg f \leq d$ and given by a "black-box" then there is an algorithm for testing identity to zero of f in time $p(m \log d \log q)^{O(1)}$.*

Let us mention that the results above are new even for "large" p (say when $q = p^2$) and give improvements of the corresponding tests from [2], [5], [16]. Also they can be implemented in parallel.

It seems that the estimate of [10] leads to the construction of the corresponding set \mathfrak{M} (containing a primitive root) of size $p^{1/2}(m \log d \log q)^{O(1)}$ in time of the same order. In this case we would get an improvement of the zero-identity test of [5] with respect to the number of processors beginning from $q = p$ ($p^{1/2}$ rather than p).

3 Image-Size of Sparse Polynomials

Here we consider the following question. Let f be a t -sparse univariate polynomial over \mathbb{F}_q and let we are given an integer number $I > 0$. How quickly can we test whether the image-size

$$I(f) = |\{f(x) : x \in \mathbb{F}_q\}|$$

is a least I ?

The "brute force" algorithm takes time $q(\log q)^{O(1)}$. Below we show that for small I (for $I < q^{1/t}$) it can be done faster.

Theorem 8 *Let f be a t -sparse univariate polynomial over \mathbb{F}_q with $\deg f \leq q - 1$ given by a "black-box" and assume that a primitive root θ of \mathbb{F}_q is given. Then for any $I > 0$ one can test if $I(f) \geq I$ in time $I^t(\log q)^{O(1)}$.*

Proof. Let

$$f(x) = \sum_{i=1}^t a_i x^{n_i}$$

where

$$a_1, \dots, a_t \in \mathbb{F}_q^*, \quad 0 \leq n_1 < \dots < n_t \leq q - 1.$$

Let us consider the sequence

$$u(x) = \sum_{i=1}^t a_i \theta^{x n_i}, \quad x = 1, 2, \dots,$$

where θ is a fixed primitive root of F_q . Let us denote

$$D = \gcd(n_1, \dots, n_t, q - 1), \quad T = (q - 1)/D.$$

We show that the vectors

$$U(x) = (u(x), \dots, u(x + t - 1)), \quad x = 1, \dots, T,$$

are pairwise different. Indeed, if $U(x) = U(y)$, $1 \leq x < y \leq T$, then

$$\sum_{i=1}^t a_i (\theta^{y n_i} - \theta^{x n_i}) \theta^{j n_i} = 0, \quad j = 0, \dots, t - 1.$$

Because of (3), it implies

$$\theta^{y n_i} = \theta^{x n_i}, \quad i = 1, \dots, t,$$

thus

$$n_i \equiv 0 \pmod{(q - 1)/d}, \quad i = 1, \dots, t.$$

where $d = \gcd(x - y, q - 1)$. Therefore,

$$D = \gcd(n_1, \dots, n_t, q - 1) \geq (q - 1)/d > (q - 1)/T = D,$$

the obtained contradiction shows that (4) is impossible. Now as t -dimensional vectors $U(x)$, $x = 1, \dots, T$, are pairwise different, their coordinates takes at least $T^{1/t}$ different values. Also, it is easy to see that $u(x) = u(x + T)$, $x = 1, 2, \dots$

Therefore, in order to check if $I(f) > I$ it is enough to compute vectors $U(1), \dots, U(Q)$ with $Q = I^t$. If there are $1 \leq x < y \leq Q$ such that $U(x) = U(y)$ therefore $T \leq Q$ and

$$I(f) = f(0) \cup \{u(1), \dots, u(Q)\},$$

otherwise $T > Q$ and

$$|I(f)| \geq T^{1/t} > I.$$

Taking into account that the size of the set (5) can be computed in time $Q(\log q)^{O(1)}$ we get the desired result. \square

Of course using various algorithms to find primitive roots (or just "small-sized" sets containing a primitive root) one can get Theorem 6 in an unconditional form (in any case it can be done in time $q^{1/4}(\log q)^{O(1)}$). The natural question is can the same result be obtained without finding an auxiliary primitive root.

The considerations above give the bounds

$$\frac{q-1}{D} \geq |I(f)| \geq \left(\frac{q-1}{D}\right)^{1/t}.$$

Also they show that either $|I(f)|$ can be computed by the "brute force" algorithm in time $qD^{-1}(\log q)^{O(1)}$ or it is large enough. The question is: in the second case, can one use a Monte Carlo algorithm to estimate $|I(f)|$ (note that the results of [6] cannot be applied directly).

References

- [1] Adleman, L.M., Lenstra, H.W., *Finding Irreducible Polynomials over Finite Fields*, in *Proc. 18th ACM Symp. on Theory of Comp.*, 1986, pp. 350–355.
- [2] Clausen, M., Dress, A. Grabmeier, J., Karpinski, M., and Singer, M., *On Zero-Testing and Interpolation of k -Sparse Multivariate Polynomials over Finite Fields*, in *SIAM J. on Comput.* 19, 1990, pp. 1059–1063.
- [3] Grigoriev, D., Karpinski, M., *An Approximation Algorithm for the Number of Zeros of Arbitrary Polynomials over $GF(q)$* , in *Proc. 32nd Ann. Symp. on Found. of Comp. Sci.*, 1991, pp. 662–669.

- [4] Grigoriev, D., Karpinski, M., *Lower Bounds for the Number of Zeros of Multivariate Polynomials over $GF(q)$* , Research Report No.8569-CS, Univ. of Bonn, 1991.
- [5] Grigoriev, D., Karpinski, M., and Singer, M., *Fast Parallel Algorithm for Sparse Multivariate Polynomials over Finite Fields*, in *SIAM J. Comput.* 19, 1990, pp. 1059-1063.
- [6] Karp, R., Luby, M., and Madras, N., *Monte-Carlo Approximation Algorithms for Enumeration Problems*, in *J. of Algorithms* 10, 1989, pp. 429-448.
- [7] Karpinski, M., *Boolean Circuit Complexity of Algebraic Interpolation Problems*, in *Lecture Notes in Computer Science* 385, Springer-Verlag, 1989, pp. 138-147.
- [8] Karpinski, M., Lhotzky, B., *An (ϵ, δ) -Approximation Algorithm of the Number of Zeros for a Multilinear Polynomial over $GF(q)$* , Preprint TR-91-022, International Computer Science Institute (ICSI), Berkeley, California.
- [9] Karpinski, M., Luby, M., *Approximating the Number of Zeros for a $GF(2)$ polynomial*, in *J. of Algorithms* 14, 1993, pp. 280-287.
- [10] Perelmuter, G.I., Shparlinski, I., *On the Distribution of Primitive Roots in Finite Fields*, (Russian), in *Uspechi Matem. Nauk* 45 (1), 1990, pp. 185-186.
- [11] Shoup, V., *Fast Construction of Irreducible Polynomials over Finite Fields*, to appear in *J. Symb. Comp.* 3, pp. 1-17.
- [12] Shoup, V., *New Algorithms for Finding Irreducible Polynomials over Finite Fields*, in *Math. Comp.* 54 (189), 1990, pp. 435-447.
- [13] Shoup, V., *Searching for Primitive Roots in Finite Fields*, in *Math. Comp.* 58 (197), 1992, pp. 369-380.
- [14] Shparlinski, I., *Computational and Algorithmic Problems in Finite Fields*, Kluwer AP, The Netherlands, 1992.
- [15] Shparlinski, I., *On Primitive Elements in Finite Fields and on Elliptic Curves*, (Russian), in *Matem. Sbornik* 181 (9), 1990, p. 1196-1206.
- [16] Werther, K., *The Computational Complexity of Interpolation of Sparse Multivariate Polynomials over Finite Fields* Research Report No.8577-CS, University of Bonn, 1992.