On Zero-Testing and Interpolation of k-sparse Multivariate Polynomials over Finite Fields

Michael Clausen
Fakultät für Informatik
Universität Karlsruhe

Andreas Dress *
Fakultät für Mathematik
Universität Bielefeld

Johannes Grabmeier
IBM Deutschland GmbH
Wissenschaftliches Zentrum Heidelberg

Marek Karpinski * †
Institut für Informatik
Universität Bonn

May 24, 1988

*Part of this work was prepared during a visit at the Science Center Heidelberg, IBM Germany which is gratefully acknowledged.

†Supported in part by Leibniz Center fo Research in Computer Science and the DFG Grant KA 673/2-1

Abstract

Given a *black box* which will produce the value of a k-sparse multivariate polynomial for any given specific argument, one may ask for optimal strategies (1) to distinguish such a polynomial from the zero polynomial, (2) to distinguish any two such polynomials from each other and (3) to (uniformly) reconstruct the polynomial from such an information source. While such strategies are known already for polynomials over fields of characteristic zero, the equally important, but considerably more complicated case of a finite field K is studied in the present paper. The result is that the time complexity of such strategies depends critically on the degree m of the extension field of K from which the arguments are to be chosen; e.g. if m equals the number n of variables, then (1) can be solved by $k+1$ and (2) as well as (3) by $2k+1$ queries, while in case $m = 1$ essentially $2^{\log n \log k}$ queries are needed.

Introduction

The question of how to interpolate polynomials has a long history in mathematics. The interpolation formulae by Newton and Lagrange for polynomials in one indeterminate over fields of characteristic 0 of degree d laid the foundation of numeric interpolation. Many generalizations, e.g. allowing coefficients from finite fields and more than one indeterminate, related subjects, e.g. the Chinese Remainder Theorem, and applications, e.g. various multiplication algorithms having better asymptotic behaviour than the school multiplication method, have been studied.

In addition, the question of how to specify appropriate data structures to store polynomials efficiently, plays an important role, wherever polynomials occur in algorithms. The methods of sparse representation, i.e. representing a polynomial by a list of records containing a non-zero coefficient and the corresponding exponent, or by straight-line programs - see [IM83] or [K85] - proved to be very successful. In this context, the problem of finding conversion algorithms from one representation to another one, gave further motivation to study interpolation of polynomials from a slightly different point of view. Here, rather than the degree, the number of terms of a polynomial is of importance.

Closely related to the interpolation problem is the somewhat easier problem to decide by appropriate evaluations in a minimal number of steps whether a *k-sparse* polynomial in n indeterminates is the zero polynomial. Here we assume the polynomial to be given by a straight-line program or, more generally, as an oracle, i.e. a black box with as many inputs as there are indeterminates and one output. For any evaluation point as input, it produces as its output the value of the polynomial at that point. Schwartz has constructed a randomized NC-algorithm for this problem in [S80]. The corresponding interpolation problem was solved also by randomized algorithms by Zippel [Z79] and Kaltofen [K85]. Unfortunately, for finite fields their results need strong restrictions on the degree of the polynomial, compare also [BT88].

A crucial step for the construction of deterministic algorithms for these problems was the work of Grigoriev and Karpinski [GK87] on finding matchings for bipartite graphs. In [T87] and [BT88] Ben-Or and Tiwari employed their ideas to use n different primes $(p_0, ..., p_{n-1})$ to solve the '$f = 0$?'-problem over fields of characteristic 0 using only k queries, namely $(p_0^i, ..., p_{n-1}^i)$ for $0 \leq i < k$, and the interpolation problem over fields of characteristic 0, using only the $2k$ queries $0 \leq i < 2k$. The crucial point is the uniqueness of the prime factorization of integers. A first application of these algorithms was an algorithm for computing the sparse representations for all k-sparse irreducible factors of such polynomials, see [KT88].

In our paper we consider these problems for k-sparse multivariate polynomials over finite fields with essentially no restriction on the degree of the polynomials. Sections 2 and 3 are devoted to the '$f = 0$?'-problem. In Theorem 2.4 test sets in extension fields $GF(q^m)$ of $GF(q)$ are constructed for any given m, the asymptotic behaviour of their cardinality being $O((n/m)^{\log k})$ for small m. If the degree of the extension field equals the number of indeterminates, we find a test set of cardinality $k + 1$ in Theorem 2.3 which is proved to be optimal in case $n = 1$ in section 3.

In the next section various lower bounds for the necessary number of queries are determined. We show that $\sum_{i=0}^{\lfloor \log_2 k \rfloor} \binom{n}{i}$ is a lower bound for the important case where no proper field extensions are allowed, which turns out to be optimal for the field with two elements, see Corollary 2.6 and Theorem 3.2.

Section 4 is devoted to the interpolation problem. As an application of

the results of section 2 and section 3 we describe a method to construct test sets A which distinguish any two given k-sparse polynomials. However, we do not know whether these test sets contain enough elements such that a non-adaptive interpolation algorithm can be derived. Even less do we know whether such an algorithm can be found in NC.

Finally, we shall show that $1 + 2k - \lfloor \frac{2k-1}{q} \rfloor$ evaluations over $GF(q^n)$ enable us to reconstruct f, where $f \in GF(q)[X_0, \ldots, X_{n-1}]$ is a polynomial satisfying $\deg_{X_i}(f) < q$ for all i. Furthermore, this algorithm is in NC modulo the problem of finding an NC-algorithm to calculate discrete logarithms. To do this we combine three tools in order to recover f: generalized Newton identities, uniqueness of the q-adic representation of the exponents of non-zero elements in $GF(q^n)$ with respect to a primitive element, and finally, the Frobenius automorphism $y \mapsto y^q$ of $GF(q^n)$ which keeps fixed all elements of $GF(q)$.

In [GKS88] closely related problems have been studied by Grigoriev, Karpinski and Singer. There it was shown that for given n, k and q one can find test sets for the '$f = 0$?'-problem of order $k(1 + (n-1)\binom{k}{2})$, provided that one works over a *slight* extension field $GF(q^m)$ of $GF(q)$ with $m = 2\log_q(kn)$. Furthermore an NC-interpolation algorithm is developped in this situation. This contrasts in a rather intriguing way to our lower bound $\sum_{i=0}^{\lfloor \log_2 k \rfloor} \binom{n}{i}$ for the number of necessary queries in case $m = 1$.

Our results may have useful applications in the area of learning machines, which we would like to investigate in subsequent papers.

1 Notations

The most general setting of the questions we are interested in are the following ones: For any two sets X and Y and any subset $\mathcal{P} \subseteq X^Y$ of mappings from Y into X one may ask for minimal *test sets* A of Y which will allow to distinguish different mappings in $\mathcal{P}$. Hence we define

$$\mathcal{B}(\mathcal{P}) := \{B \mid B \subseteq Y,\ \forall f, g \in \mathcal{P}\ \exists b \in B\ \ (f(b) \neq g(b))\}$$

and for $f \in \mathcal{P}$ we define

$$\mathcal{A}(\mathcal{P}, f) := \{A \mid A \subseteq Y,\ \forall g \in \mathcal{P} \setminus \{f\}\ \exists a \in A\ \ (g(a) \neq f(a))\}$$

and

$$c(\mathcal{P}, f) := min\{\#A \mid A \in \mathcal{A}(\mathcal{P}, f)\}.$$

If $X = K$ is an arbitrary field and $\mathcal{P}$ a linear subspace, then $c(\mathcal{P}, 0) = dim\mathcal{P}$, hence for arbitrary $\mathcal{P}$ we conclude $c(\mathcal{P}, 0) \leq dim\ span_K\mathcal{P}$. Therefore, w.l.o.g. one may restrict one's attention to those subsets $\mathcal{P} \subseteq K^X$ which span the whole space K^X.

In this note we consider the following special case: For a finite field $GF(q)$ of prime power order q the ring of (polynomial) maps from $GF(q)^n$ into $GF(q)$ is isomorphic to $GF(q)[X_0, \ldots, X_{n-1}]$, the polynomial ring in n indeterminates, modulo the ideal generated by $X_0^q - X_0, \ldots, X_{n-1}^q - X_{n-1}$. We identify its elements with the polynomials $f \in GF(q)[X_0, \ldots, X_{n-1}]$ satisfying $\deg_{X_i}(f) < q$ for all i.

Let $\mathcal{P}_k^n(q)$ denote the set of all such polynomials f which in addition are k-sparse, i.e. the positive integer k is an upper bound for the number of non-zero coefficients of f. For given q we want to discuss upper and lower bounds for the number

$$c_k^n(q) := c(\mathcal{P}, 0)$$

where $\mathcal{P}$ consists of all polynomials in $\mathcal{P}_k^n(q)$, considered as maps from $GF(q)^n$ into $GF(q)$. In this case we also write $\mathcal{A}_k^n(q)$ for $\mathcal{A}(\mathcal{P}, 0)$ and $\mathcal{B}_k^n(q)$ for $\mathcal{B}(\mathcal{P})$. More generally, for given q and m we shall consider

$$c_k^n(q, m) := c(\mathcal{P}_m, 0)$$

where $\mathcal{P}_m$ again consists of all polynomials in $\mathcal{P}_k^n(q)$, but now considered as maps from $GF(q^m)^n$ into $GF(q^m)$. In this case we also write $\mathcal{A}_k^n(q, m)$ for $\mathcal{A}(\mathcal{P}_m, 0)$ and $\mathcal{B}_k^n(q, m)$ for $\mathcal{B}(\mathcal{P}_m)$.

2 Test Sets and Upper Bounds

To derive upper bounds for $c_k^n(q, m)$ by constructing evaluation sets in $\mathcal{A}_k^n(q, m)$ the following observation appears to be crucial.

Lemma 2.1 *Let $n = n_1 + n_2$, $A_{k_1}^{n_1} \in \mathcal{A}_{k_1}^{n_1}(q, m)$ and $A_{k_2}^{n_2} \in \mathcal{A}_{k_2}^{n_2}(q, m)$ for all $k_1 \cdot k_2 \leq k$. Then $\bigcup_{k_1 \cdot k_2 \leq k} A_{k_1}^{n_1} \times A_{k_2}^{n_2}$ is a member of $\mathcal{A}_k^n(q, m)$.*

Proof. Consider $0 \neq f \in \mathcal{P}_k^n(q)$ as a polynomial in the indeterminates $X_{n_1}, ..., X_{n-1}$ with polynomial coefficients in $K[X_0, \ldots, X_{n_1-1}]$. The corresponding number k_2 of terms of f is of course bounded by k, one of the k_2 non-zero polynomial coefficients, say f_β, has at most $k_1 := \lfloor \frac{k}{k_2} \rfloor$ terms. Hence there exists an element $a^{(1)} \in A_{k_1}^{n_1}$ such that $f_\beta(a^{(1)}) \neq 0$. Consequently, $f(a^{(1)}, X_{n_1}, \ldots X_{n-1})$ is a non-zero k_2-sparse polynomial in n_2 indeterminates for which we can find an $a^{(2)} \in A_{k_2}^{n_2}$ with $f(a^{(1)}, a^{(2)}) \neq 0$. ∘

Corollary 2.2 *For an arbitrary improper partition $\pi = (\pi_0, \ldots, \pi_{s-1})$ of n, i.e. $\pi_i \in \mathbb{N}$ and $\sum_{i=0}^{s-1} \pi_i = n$ — in short $\pi \models n$ — and for all $\kappa \leq k$ let $A_\kappa^{\pi_i}$ be an arbitrary set from $\mathcal{A}_\kappa^{\pi_i}(q, m)$. Then $\bigcup_{\kappa_0 \cdot \kappa_1 \cdot \ldots \cdot \kappa_{s-1} \leq k} A_{\kappa_0}^{\pi_0} \times A_{\kappa_1}^{\pi_1} \times \ldots \times A_{\kappa_{s-1}}^{\pi_{s-1}}$ is a member of $\mathcal{A}_k^n(q, m)$.*

Obviously, corresponding results also hold for arbitrary ground fields of arbitrary characteristic, zero or prime. Corollary 2.2 will be used in conjunction with the following result.

Theorem 2.3 *Let $f \in GF(q)[X_0, \ldots, X_{n-1}]$ be a k-sparse polynomial, $k \geq 2$, satisfying $\deg_{X_i}(f) < q$, for all i, and let ω be a primitive element of $GF(q^n)$. Then f is the zero-polynomial if and only if $f(0, \ldots, 0) = 0$ and $f_i := f(\omega^{iq^0}, \omega^{iq^1}, \ldots, \omega^{iq^{n-1}}) = 0$, for all i satisfying $0 \leq i < k$ and $q \nmid i$ in case $i > 0$. Any set consisting of one element which has no zero components is a test set for the case $k = 1$.*

Proof. If $f \in GF(q)[X_0, \ldots, X_{n-1}]$ satisfies $\deg_{X_i}(f) < q$, for all i, then f is a linear combination over $GF(q)$ of the q^n monomials $X^\alpha := X_0^{\alpha_0} \cdot \ldots \cdot X_{n-1}^{\alpha_{n-1}}$, where α ranges over all maps in $\mathbf{q}^n := \{0, \ldots, q-1\}^{\{0,\ldots,n-1\}}$:

$$f = \sum_{\alpha \in \mathbf{q}^n} c_\alpha X^\alpha.$$

Now assume $f(0, ..., 0) = c_{(0,\ldots,0)} = 0$ and $f_i = 0$ for all i satisfying $0 \leq i < k$ and $q \nmid i$ in case $i > 0$. By the properties of the Frobenius automorphism we have

$$f_{j \cdot q} = (f_j)^q, 1 \leq j \cdot q < k$$

and hence $f_i = 0$ for all i satisfying $0 \le i < k$. The mapping

$$\Omega: \mathrm{q}^n \setminus \{(0,\ldots,0)\} \to GF(q^n) \setminus \{0\}$$

defined by

$$\Omega_\alpha := \prod_{0 \le \nu < n} \omega^{\alpha_\nu \cdot q^\nu}$$

is bijective since $\Omega_\alpha = \omega^{(\sum \alpha_\nu q^\nu)}$, so from the q-adic expansion of the exponent we can recover α. Let A be any k-subset of $\mathrm{q}^n \setminus \{(0,\ldots,0)\}$ containing the support $\mathrm{supp}(f) := \{\alpha : c_\alpha \neq 0\}$ of f. Then

$$f_i = \sum_{\alpha \in \mathrm{q}^n \setminus \{(0,\ldots,0)\}} c_\alpha \Omega_\alpha^i = \sum_{\alpha \in A} c_\alpha \Omega_\alpha^i$$

for all $0 \le i < k$. Thus we obtain the following matrix equation

$$(\Omega_\alpha^i)_{0 \le i < k, \alpha \in A} \cdot (c_\alpha)_{\alpha \in A} = (f_i)_{0 \le i < k}.$$

The k-square matrix (Ω_α^i) is a non-singular Vandermonde matrix since the Ω_α are pairwise different. Hence f is the zero-polynomial. The case $k = 1$ is clear. □

The test set given in this theorem is an element of $\mathcal{A}_k^n(q,n)$ and hence $c_k^n(q,n) \le 1 + k - \lfloor (k-1)/q \rfloor$. To state the main result of this section we need the following test sets
$T_k^n(q,m) :=$

$$\{a = (a_\nu)_{0 \le \nu < n-1} \in GF(q^m)^n : a_\nu = \epsilon_\mu \omega^{i_\mu \cdot q^\iota}, \nu = \mu \cdot m + \iota, 0 \le \iota < m,\ \epsilon_\mu \in$$

$$\{0,1\}, i_\mu = 0 \text{ or } 1 \le i_\mu < k \text{ such that } q \not| \, i_\mu, 2^{\#\{\mu : c_\mu = 0\}} \cdot \prod_{\{\mu : \epsilon_\mu \neq 0\}} (1 + i_\mu) \le k\},$$

where ω is a primitive element in $GF(q^m)$. Corollary 2.2 and Theorem2.3 together imply:

Theorem 2.4 *Let $f \in GF(q)[X_0,\ldots,X_{n-1}]$ be a k-sparse polynomial, $k \ge 2$, satisfying $\deg_{X_i}(f) < q$ for all i. Then f is the zero-polynomial if and only if it vanishes at all elements of $T_k^n(q,m)$.*

Proof. Use the partition $\pi := (m, \ldots, m, m_1)$ of n with $m_1 \leq m$ in Corollary 2.2 and choose the test sets $A^{\pi}_{\kappa_\mu}$ according to Theorem 2.3. Note further that the occurence of a block of zeros implies that the corresponding κ_μ is at least 2 while in the other cases it is at least $1 + i_\mu$. o

Corollary 2.5 *Define*

$$UB(n,k,q) := \sum_{\substack{\kappa=(\kappa_0,\ldots,\kappa_{q-1})\models n \\ 2^{\kappa_0}+\kappa_2\cdot 3^{\kappa_3}\cdot\ldots\cdot(q-1)^{\kappa_{q-1}}\leq k}} \binom{n}{\kappa} = \#T^n_k(q,1).$$

Then we have

$$c^n_k(q) \leq UB(n,k,q)$$

as well as

$$c^n_k(q,m) \leq UB(\lceil \frac{n}{m} \rceil, k, q^m).$$

Note that $UB(n,k,q) \leq (n\cdot(q-1))^{\lfloor \log_2 k \rfloor}$. More precise estimates can be derived from Ch.10 §11 in the book by MacWilliams and Sloane [MS72].

In [GKS88] it is shown by Grigoriev, Karpinski and Singer that $c^n_k(q,m) \leq k(1+(n-1)\binom{k}{2})$, once m satisfies $\lceil \frac{q^m-1}{2n(q-1)} \rceil - 1 > (n-1)\binom{k}{2}$ which is certainly true for $m \geq 2\log_q(kn)$. Using their results instead of Theorem 2.3, the above method can be applied similarly to yield

$$c^n_k(q,m) \leq \sum_{\substack{\kappa=(\kappa_1,\kappa_2,\ldots)\models n \\ 2^{\kappa_2}\cdot 3^{\kappa_3}\cdot\ldots\leq k}} \binom{\lceil \frac{n}{n_0} \rceil}{\kappa} \cdot (1+(n_0-1)(i-1)(3i-2)/1)^{\kappa_i},$$

where

$$n_0 := max\{\tilde{n} : \lceil \frac{q^m-1}{2\tilde{n}(q-1)} \rceil - 1 > (\tilde{n}-1)\binom{k}{2}\}.$$

This result is interesting for $n \gg q^{m/2}k$.

In the next two corollaries important special cases are considered.

Corollary 2.6 *Let $f \in GF(2)[X_0, \ldots, X_{n-1}]$ be a polynomial satisfying $\deg_{X_i}(f) < 2$, for all i and which is k-sparse. Then f is the zero-polynomial if and only if all $f(a) = 0$ for all $a \in A_k := T_k^n(2,1)$, the set of all elements from $GF(2)^n$ having at most $\lfloor log_2 k \rfloor$ zero positions. Hence $c_k^n(2) \leq \sum_{i=0}^{\lfloor \log_2 k \rfloor} \binom{n}{i}$ elements.*

Corollary 2.7 *Let ω be a primitive element in $GF(q)$. Then the set $\{(1, \ldots, 1)\} \cup \{a \in GF(q)^n : a_\nu \in \{0, \omega\}$ for exactly one ν and $a_\nu = 1$ elsewhere $\}$ is a test set to decide whether a binomial is 0 and hence*

$$c_2^n(q) \leq \begin{cases} 1+n, & \text{if } q = 2 \\ 1+2n, & \text{if } q \neq 2 \end{cases}.$$

3 Lower Bounds

In this section we determine lower bounds for $c_k^n(q,m)$. As every k-sparse polynomial can be split into a difference of a $\lfloor k/2 \rfloor$-sparse and a $\lceil k/2 \rceil$-sparse polynomial, a set $A \in \mathcal{A}_k^n(q,m)$ has to contain an element where these polynomials have different values. Hence the map

$$\#\mathcal{P}_{\lfloor k/2 \rfloor}^n(q) \to \#GF(q^m)^A, f \mapsto (f(a))_{a \in A}$$

must be injective. Therefore

$$\#\mathcal{P}_{\lfloor k/2 \rfloor}^n(q) \leq \#GF(q^m)^A,$$

that is

$$(1/m) \cdot \log_q\Big(\sum_{i=0}^{\lfloor k/2 \rfloor} \binom{n}{i} \cdot (q-1)^i\Big) \leq c_k^n(q,m).$$

Besides this trivial result our first aim is to show that in case $m = 1$ it is not possible to decide the question whether a k-sparse polynomial is the zero polynomial knowing only polynomially many (in k and n) evaluations. We show that the number of necessary evaluations in this case is pseudo-polynomial: $\Omega(n^{\log k}) = \Omega(k^{\log n})$ as long as k is substantially smaller than 2^n.

Theorem 3.1 *Assume $A \in \mathcal{A}_k^n(q)$, that is, A is a test set of evaluation points in $GF(q)^n$ which enables us to decide whether a k-sparse polynomial $f \in GF(q)[X_0, \ldots, X_{n-1}]$ satisfying $\deg_{X_i}(f) < q$ for all i, is the zero-polynomial. Then for every subset $T \subseteq \{0, ..., n-1\}$ such that $\#T \leq \lfloor log_2 k \rfloor$ the set A of contains an element $a^T = (a_0^T, \ldots, a_{n-1}^T)$ with $T = \{i : a_i^T = 0\}$. Hence A has at least $\sum_{i=0}^{\lfloor \log_2 k \rfloor} \binom{n}{i}$ elements, i.e.*

$$\sum_{i=0}^{\lfloor \log_2 k \rfloor} \binom{n}{i} \leq c_k^n(q).$$

Proof. For every subset $T \subseteq \{0, ..., n-1\}$ such that $\#T \leq \lfloor log_2 k \rfloor$ define a polynomial

$$p_T := \prod_{i \in T} (x_i^{q-1} - 1) \cdot \prod_{i \notin T} x_i.$$

These polynomials have the following properties:

1. p_T is k-sparse.
2. $p_T(a) = 1$ if and only if $\{i : a_i = 0\} = T$.

The first property follows from $2^{\#T} \leq 2^{\lfloor \log_2 k \rfloor} \leq k$, the second from the fact that the zeros of $x_i^{q-1} - 1$ are exactly the elements of $GF(q) \setminus \{0\}$. Hence, to distinguish such a polynomial and the zero-polynomial, there has to be an element a^T as claimed in the set A. □

In case $q = 2$ we may combine Corollary 2.6 and Theorem 3.1 to determine $c_k^n(2)$ exactly:

Theorem 3.2

$$c_k^n(2) = \sum_{i=0}^{\lfloor \log_2 k \rfloor} \binom{n}{i}.$$

The next result where upper and lower bounds coincide is the case $n = 1 = m$:

Lemma 3.3

$$c_k^1(q) = \begin{cases} min\ \{k+1, q\}, & \text{if } k \geq 2 \\ 1, & \text{if } k = 1 \end{cases}.$$

Proof. (Compare the proof of Theorem 7 in [BT88]). If $k = q$, then indeed $c_q^1(q) = q$, since any map $GF(q) \to GF(q)$ is in $\mathcal{P}_q^1(q)$. If $q > k \geq 2$ and $A \subseteq GF(q)$ has cardinality k, then if $0 \notin A$ the polynomial $f := X^{q-1} - 1$ is k-sparse and vanishes on A. If $0 \in A$ then $f := \prod_{a \in A \setminus \{0\}} (X - a)$ is a non-zero polynomial in $GF(q)[X]$ of degree at most $k-1$. f and $X \cdot f$ have at most k monomials and the latter vanishes on A. Finally, if $k = 1$, one needs precisely one evaluation to check whether $f = 0$ holds. □

The next result covers the case of binomials and the proof of the theorem may give a hint about the difficulties which may arise while trying to prove sharp lower bounds for $k \geq 3$.

Theorem 3.4 *Let $q > 2$ and let ω be a primitive element of $GF(q)$. Assume $A \in \mathcal{A}_2^n(q)$, that is A is a set of evaluating points in $GF(q)^n$ which enables us to decide whether a binomial $f = c_\alpha X^\alpha + c_\beta X^\beta$ satisfying $\deg_{X_i}(f) < q$ for all i, is the zero-polynomial. Then A contains n elements $a^{(\mu)} = (a_0^{(\mu)}, \ldots, a_{n-1}^{(\mu)}), 0 \leq \mu < n$ of the shape*

$$a_\nu^{(\mu)} = 0 \quad \text{if and only if} \quad \mu = \nu.$$

Furthermore A contains $n+1$ further different elements $a^{(\mu)}, n \leq \mu \leq 2n$ having no zero components, i.e there exist $0 \leq b_\nu^{(\mu)} < q-1$ such that $a_\nu^{(n+\mu)} = \omega^{b_\nu^{(\mu)}}, 0 \leq \mu \leq n, 0 \leq \nu < n$. In particular we have

$$c_2^n(q) = 2n + 1.$$

Proof. Assume a set $A \in \mathcal{A}_2^n(k)$ is given. We define

$$p_\nu := (X_\nu^{q-1} - 1) \cdot \prod_{\mu \neq \nu} X_\mu \in \mathcal{P}_2^n(q),\ 0 \leq \nu < n.$$

These polynomials have the property

$$p_\nu(a) = 0 \text{ if and only if } a_\nu = 0, a_\mu \neq 0 \text{ for } \mu \neq \nu.$$

Hence the first assertion follows.
Now suppose that there are at most $\tilde{n} \leq n$ elements with no zero components. We shall construct a binomial which vanishes on A. For that purpose we construct some $\alpha = (\alpha_0, \ldots, \alpha_{n-1}) \in \{0, \ldots, q-2\}^n$ and some $c \in GF(q)$ such that $\alpha \neq (0, \ldots, 0)$ and $a^\alpha := \prod_{\nu=0}^{n-1} a_\nu^{\alpha_\nu} = c$ for all elements $a \in A$ having no zero components. We denote these elements and their exponents with respect to ω as in the theorem by $a^{(\mu+\mu')}, 0 \leq \mu < \tilde{n}$ and $b_\nu^{(\mu)}$. If $c = \omega^d$ then the last condition is equivalent to

$$\omega^{\sum_{\nu=0}^{n-1} b_\nu^{(\mu)} \cdot \alpha_\nu} = \omega^d, 0 \leq \mu < \tilde{n}$$

which is equivalent to

$$(b_\nu^{(\mu)})_{0 \leq \mu < \tilde{n}, 0 \leq \nu < n} \cdot (\alpha_\nu)_{0 \leq \nu < n} = (d, \ldots, d)^t \text{ over } \mathbb{Z}/(q-1)\mathbb{Z}.$$

If the linear map

$$(\mathbb{Z}/(q-1)\mathbb{Z})^n \to (\mathbb{Z}/(q-1)\mathbb{Z})^{\tilde{n}} \text{ defined by } (b_\nu^{(\mu)})_{0 \leq \mu < \tilde{n}, 0 \leq \nu < n}$$

is not injective, then clearly there exists some non-trivial α for $d = 0$ satisfying the above equation. If the map is injective and therefore bijective, in particular $n = \tilde{n}$, then we may choose $d = 1$ and α as the unique and necessarily non-trivial pre-image of $(1, \ldots, 1)$. In any case

$$f := X_0 \cdot X_1 \cdot \ldots \cdot X_{n-1} \cdot (X^\alpha - \omega^d)$$

will vanish on A. Hence in a set $A \in \mathcal{A}_2^n(q)$ there are at least $n+1$ elements without zero components. □

Let us finally remark that even in the case $m = 1$ our upper bound $UB(n, k, q)$ does not coincide with $c_k^n(q)$, e.g. it can be shown that $c_4^4(3) \leq 32$, while $UB(4, 4, 3) = 33$. Nevertheless it appears to be very close to $c_k^n(q)$.

4 Interpolation

We first solve the problem to distinguish two k-sparse multivariate polynomials over $GF(q)$. With the notation of Section 1 we have to construct

elements of $\mathcal{B}_k^n(q,m)$. Fortunately, all the work is reduced to the construction of the test sets in Section 2 by the following

Lemma 4.1 $\mathcal{B}_k^n(q,m) = \mathcal{A}_{2k}^n(q,m)$.

Proof. Assume $B \in \mathcal{B}_k^n(q,m)$ and $0 \neq h \in \mathcal{P}_{2k}^n(q)$, then there exist polynomials f,g in $\mathcal{P}_k^n(q)$ such that $h = f - g$. Furthermore, there exists some $b \in B$ with $f(b) \neq g(b)$, hence $h(b) = f(b) - g(b) \neq 0$ which implies $B \in \mathcal{A}_{2k}^n(q,m)$. On the other hand assume $A \in \mathcal{A}_{2k}^n(q,m)$ and f,g in $\mathcal{P}_k^n(q)$. Then $h := f - g$ is in $\mathcal{P}_{2k}^n(q)$. Furthermore, there exists an $a \in A$ with $h(a) \neq 0$, hence $f(a) \neq g(a)$ which implies $A \in \mathcal{B}_k^n(q,m)$. □

It is also clear that the lower and upper bounds carry over at once. Let us remark that for any $A \in \mathcal{A}_{2k}^n(q,m)$ the evaluation map

$$\Psi_A : \mathcal{P}_k^n(q) \to GF(q^m)^A, f \mapsto (f(a)_{a\in A})$$

is injective. In particular there exist a left inverse

$$\Phi_A : GF(q^m)^A \to \mathcal{P}_k^n(q).$$

However, it is by no means clear whether the construction of an algorithm which represents some Φ_A can be done uniformly for n,k,q and m.

In the following theorem we construct a set of $1+2k-\lfloor\frac{2k-1}{q}\rfloor$ evaluation points which enable us to reconstruct f in case $m = n$.

Theorem 4.2 *Let $f \in GF(q)[X_0,\ldots,X_{n-1}]$ be a k-sparse polynomial satisfying $\deg_{X_i}(f) < q$ for all i, and let ω be a primitive element of $GF(q^n)$. Then in order to construct f it suffices to know the values $f(0,\ldots,0)$ and $f_i := f(\omega^{iq^0},\omega^{iq^1},\ldots,\omega^{iq^{n-1}}) = 0$ for all i satisfying $0 \leq i < 2k$ and $q \nmid i$.*

Proof. Assume that $f \in GF(q)[X_0,\ldots,X_{n-1}]$ satisfies $\deg_{X_i}(f) < q$ for all i. Then we have

$$f = \sum_{\alpha\in q^n} c_\alpha X^\alpha.$$

We use the notation of the proof of Theorem 2.3. In addition we can assume that

$$f(0,\ldots,0) = c_{(0,\ldots,0)} = 0,$$

otherwise we construct $f - f(0,\ldots,0)$.
For any subset A of $\mathrm{q}^n \setminus \{(0,\ldots,0)\}$ we denote by $c_i(A)$ the i-th elementary symmetric polynomial in $|A|$ indeterminates, evaluated at $(\Omega_\alpha)_{\alpha\in A}$. Now substituting Ω_α, $\alpha \in A$, for X in the polynomial

$$\prod_{\beta\in A}(X - \Omega_\beta) = \sum_{j=0}^{|A|}(-1)^{|A|-j} c_{|A|-j}(A)\cdot X^j \in GF(q^n)[X] \qquad (1)$$

yields the generalized Newton identities [MS72, p. 244]

$$0 = \sum_{j=0}^{|A|}(-1)^{|A|-j} c_{|A|-j}(A)\Omega_\alpha^j, \quad \alpha \in A.$$

Fixing an i $(0 \le i < q^n)$, multiplying the equation corresponding to α by $c_\alpha\Omega_\alpha^i$ and summing over all $\alpha \in A$ results in the following system of equations

$$0 = \sum_{j=0}^{|A|}(-1)^{|A|-j} c_{|A|-j}(A) f_{i+j}, \quad 0 \le i < q^n.$$

As $c_0 = 1$, for an arbitrary superset A of $\mathrm{supp}(f)$ the equations for $0 \le i < |A|$ are equivalent to the matrix equation

$$(f_{i+j})_{0\le i,j<|A|}\cdot\left((-1)^{|A|-j} c_{|A|-j}(A)\right)_{0\le j<|A|} = -(f_{i+|A|})_{0\le i<|A|}. \qquad (2)$$

The matrix $(f_{i+j})_{0\le i,j<|A|}$ equals $(\Omega_\alpha^i)D_A(\Omega_\alpha^i)^t$, where $D_A = \mathrm{diag}((c_\alpha)_{\alpha\in A})$ is a $|A|$-square diagonal matrix, see [LN83, 9.48, 9.49]. Hence the cardinality k of $\mathrm{supp}(f)$ equals the rank of the k-square matrix $(f_{i+j})_{0\le i,j<k}$; furthermore $(f_{i+j})_{0\le i,j<k}$ is non-singular and we can calculate the polynomial $\prod_{\alpha\in\mathrm{supp}(f)}(X - \Omega_\alpha)$ from (2) and (1) for $A = \mathrm{supp}(f)$. Finding all the roots gives $\{\Omega_\alpha : \alpha \in \mathrm{supp}(f)\}$ which enables us to recover $\mathrm{supp}(f)$. The solution of

$$(\Omega_\alpha^i)_{0\le i<k,\alpha\in A}\cdot(c_\alpha)_{\alpha\in A} = (f_i)_{0\le i<k}.$$

gives the complete polynomial f. This proves the theorem. □

Now we present and analyze the algorithm, which can be derived from the last theorem.

Interpolation Algorithm. *Let $f \in GF(q)[X_0, \ldots, X_{n-1}]$ be a k-sparse polynomial satisfying $\deg_{X_i}(f) < q$, for all i; $2k < q^n$.*

INPUT: *Oracle for f.*

step 1. *Take a primitive element ω in $GF(q^n)$.*

step 2. *Ask the oracle for the $1 + 2k \cdot \lfloor \frac{2k-1}{q} \rfloor$ values $f(0, \cdots, 0)$ and f_i, where $0 \leq i < 2k$ and $q \not| \, i$ in case $i > 0$.*

step 3. *For all $0 \leq i < 2k$ which satisfy $i = q^s \cdot i_0$, $1 \leq s$, s maximal, calculate $f_i = f_{i_0}{}^{(q^s)}$.*

step 4. *Determine $\tilde{k}$, which is the rank of the matrix $(f_{i+j})_{0 \leq i,j < k}$.*

step 5. *Solve the equation $(f_{i+j})_{0 \leq i,j < \tilde{k}} \cdot ((-1)^{\tilde{k}-j} e_{\tilde{k}-j}(\mathrm{supp}(f))_{0 \leq j < \tilde{k}} = -(f_{\tilde{k}+i})_{0 \leq i < \tilde{k}}$.*

step 6. *Find all the roots Ω_α, $\alpha \in \mathrm{supp}(f)$, of the polynomial $\sum_{i=0}^{\tilde{k}} (-1)^{\tilde{k}-i} e_{\tilde{k}-i}(\mathrm{supp}(f)) \cdot X^i$.*

step 7. *Calculate the q-adic expansion of the exponents of the Ω_α with respect to ω to get $\mathrm{supp}(f)$.*

step 8. *Solve the system of linear equations $(\Omega_\alpha^i)_{0 \leq i < \tilde{k},\ \alpha \in A} \cdot (c_\alpha)_{\alpha \in A} = (f_i)_{0 \leq i < \tilde{k}}$, for $A := \mathrm{supp}(f)$.*

OUTPUT: $(c_\alpha, \alpha)_{\alpha \in \mathrm{supp}(f)}$.

Once a primitive element ω is given, we compute the rank of the k-square matrix (f_{i+j}) within $O(k^{4.5})$ arithmetic processors and $O(\log^2 k)$ parallel time [M86]. The same bounds are valid for step 5. We use [G84] for factoring the univariate polynomial of step 6. This costs $O(\log^2 k)$ parallel time and roughly the same number of processors as above. Step 7 heavily relies on the problem to calculate discrete logarithms, see e.g. [COS86]. Step 8 is of $O(k^{4.5})$ size and $O(\log^2 k)$ parallel time.

With respect to the number of queries the algorithm is optimal in case $n = 1$ and $2k < q$. To see this let A be a subset of $GF(q)$ with at most $2k$ elements. Then similar considerations as in the proof of Lemma 3.3 enables us to construct two different k-sparse polynomials which coincide on A.

Acknowledgements

We would like to thank T. Beth, J. von zur Gathen, A. Shokrollahi and P. Tiwari for helpful discussions and comments.

References

[AL86] Adleman, L.M., Lenstra, H.K. Finding Irreducible Polynomials over Finite Fields, Proc. STOC ACM, (1986), 350–355.

[B81] Ben-Or, M. Probabilistic Algorithms in Finite Fields, Proc. 22ndIEEE FOCS (1981), 394-398.

[BT87] Ben-Or, M., Tiwari, P. A Deterministic Algorithm for Sparse Multivariate Polynomial Interpolation,Proc. Bonn Workshop on Foundations of Computing, Bonn, June 28 - July 3, 1987.

[BT88] Ben-Or, M., Tiwari, P. A Deterministic Algorithm for Sparse Multivariate Polynomial Interpolation,Proc. STOC ACM, (1988).

[COS86] Coppersmith, D., Odlyzko, A.M., Schroeppel, R. Discrete Logarithms in $GF(p)$ Algorithmica 1(1986),1–15.

[G83] von zur Gathen, J. Factoring Sparse Multivariate Polynomials, Proc. 24th IEEE FOCS (1983), 172–179.

[G84] von zur Gathen, J. Parallel Algorithm for Algebraic Problems, SIAM J. Comput., 13 (1984), 808–824.

[GK87] Grigoriev, D.Y., Karpinski, M. The Matching Problem for Bipartite Graphs with Polynomially Bounded Permanents is in NC,to appear in Proc. 28th IEEE FOCS (1987), Los Angeles, Oct. 12–14, 1987.

[GKS88] Grigoriev, D.Y., Karpinski, M., Singer, M.F. Fast Parallel Algorithms for Sparse Multivariate Polynomial Interpolation over Finite Fields preprint, 1988.

[IM83] Ibarra, O.H., Moran, S. Probabilistic Algorithms for Deciding Equivalence of Straight-Line Programs, J. ACM, 30,1, (1983), 189–192.

[K85] Kaltofen, E. Computing with Polynomials Given by Straight-Line Programs I Greatest Common Divisors, Proc. 17th ACM STOC (1985), 131–142.

[KT88] Kaltofen, E., Trager, B. Sparse Factorization and Rational Function Interpolation of Polynomials Given by Black Boxes for their Evaluations Preliminary Report(1988).

[L83] Lenstra, A.K. Factoring Multivariate Polynomials over Finite Fields, Proc. 15th ACM STOC (1983), 189-192.

[LN83] Lidl, H., Niederreiter, H. Finite Fields,Encyclopedia of Mathematics and its Applications, Vol.10, Cambridge University Press 1983.

[MS72] MacWilliams, F.J., Sloane, N.J.A. The Theory of Error Correcting Codes,North Holland (1972).

[M86] Mulmuley, K. A Fast Parallel Algorithm to Compute the Rank of a Matrix over an Arbitrary Field, Proc. STOC ACM (1986), 338-339.

[S80] Schwartz, J.T. Fast Probabilistic Algorithms for Verification of Polynomial Identities, JACM, 27, 4(1980),701-717.

[T87] Tiwari, P. Deterministic Algorithm for Multivariate Polynomial Interpolation, preliminary draft, IBM Thomas J. Watson Research Center (June, 1987).

[Z79] Zippel, R.E. Probabilistic Algorithms for Sparse Polynomials, Proc. EUROSAM'79, Springer Lec. Notes Comp. Sci., 72, (1979), 216-226.