LMS J. Comput. Math. 19 (Special issue A) (2016) 196–204 © 2016 Author
doi:10.1112/S1461157016000334

Roots of sparse polynomials over a finite field

Zander Kelley

Abstract

For a t-nomial $f(x)=\sum_{i=1}^{t} c_i x^{a_i} \in \mathbb{F}_q[x]$, we show that the number of distinct, nonzero roots of f is bounded above by $2(q-1)^{1-\varepsilon}C^{\varepsilon}$, where $\varepsilon = 1/(t-1)$ and C is the size of the largest coset in $\mathbb{F}_q^*$ on which f vanishes completely. Additionally, we describe a number-theoretic parameter depending only on q and the exponents a_i which provides a general and easily computable upper bound for C. We thus obtain a strict improvement over an earlier bound of Canetti *et al.* which is related to the uniformity of the Diffie–Hellman distribution. Finally, we conjecture that t-nomials over prime fields have only $O(t \log p)$ roots in $\mathbb{F}_p^*$ when $C=1$.

1. *Introduction*

Over the real numbers, the classical Descartes rule implies that the number of distinct, real roots of a t-nomial $f(x)=c_1x^{a_1}+\ldots+c_tx^{a_t} \in \mathbb{R}[x]$ is less than $2t$, regardless of its degree. It is a natural algebraic problem to look for analogous sparsity-dependent bounds over other fields that are not algebraically closed. In [3], Canetti *et al.* derive the following analog of Descartes' rule for polynomials in $\mathbb{F}_q[x]$.

Theorem 1.1 [3, Lemma 7]. *For $f(x)=c_1x^{a_1}+c_2x^{a_2}+\ldots+c_tx^{a_t} \in \mathbb{F}_q[x]$ (with c_i nonzero), if $R(f)$ denotes the number of distinct, nonzero roots of f in $\mathbb{F}_q$, then*

$$R(f) \leqslant 2(q-1)^{1-1/(t-1)}D^{1/(t-1)} + O((q-1)^{1-2/(t-1)}D^{2/(t-1)}),$$

where

$$D(f)=\min_i \max_{j\neq i}(\gcd(a_i-a_j, q-1)).$$

For $\vartheta \in \mathbb{F}_p^*$, the associated Diffie–Hellman distribution is defined by the random variable $(\vartheta^x, \vartheta^y, \vartheta^{xy})$ where x and y are uniformly random over $\{1,\ldots,p-1\}$. The Diffie–Hellman cryptosystem relies on the assumption that an attacker cannot easily determine ϑ^{xy} given the values of ϑ, ϑ^x, and ϑ^y. In [3], Canetti *et al.* showed that Diffie–Hellman distributions are very nearly uniform (which is an important property for the security of the cryptosystem), and the bound in Theorem 1.1 was the central tool which powered their arguments.

Since then, the bound has been a useful tool for studying various algorithmic and number-theoretic problems: in [7] it was used to study the complexity of recovering a sparse polynomial from a small number of approximate values (which is relevant to the security of polynomial pseudorandom number generators); in [8] it was used to study the singularity of generalized Vandermonde matrices over $\mathbb{F}_q$; in [1] it was used to study the solutions of exponential congruences $x^x = a \bmod p$; and in [5] it was used to study the correlation of linear recurring

Received 22 February 2016.

2010 Mathematics Subject Classification 11T06.

Contributed to the Twelfth Algorithmic Number Theory Symposium (ANTS-XII), Kaiserslautern, Germany, 29 August–2 September 2016.

Author partially supported by NSF grants DMS-1460766 and CCF-1409020.

sequences over $\mathbb{F}_2$. The main result of this paper is a new bound (Theorem 2.3) improving Theorem 1.1 by removing the asymptotic term and replacing D by a smaller, intrinsic parameter.

2. *Statement of results*

More recently in [2], Bi, Cheng, and Rojas studied the computational complexity of deciding whether a t-nomial f has a root in $\mathbb{F}_q^*$. They gave a sub-linear algorithm for this problem; we give a rough sketch here. First, they efficiently replace any instance with a t-nomial of degree bounded by $2(q-1)^{1-1/(t-1)}$ (while preserving the answer to the decision problem), and then they compute the greatest common divisor of this instance and $x^{q-1}-1$, which can be done in time proportional to the degree of the instance. As a result of their investigation, they derive the following characterization of the roots of a sparse polynomial in $\mathbb{F}_q[x]$.

THEOREM 2.1 [2, Theorem 1.1]. *For $f(x) = c_1x^{a_1} + c_2x^{a_2} + \ldots + c_tx^{a_t} \in \mathbb{F}_q[x]$, define $\delta(f) = \gcd(a_1, a_2, \ldots, a_t, q-1)$. The set of nonzero roots of f in $\mathbb{F}_q$ is the union of no more than*

$$2\left(\frac{q-1}{\delta}\right)^{1-1/(t-1)}$$

cosets of two subgroups $H_1 \subseteq H_2$ of $\mathbb{F}_q^*$, where

$$|H_1| = \delta \quad \text{and} \quad |H_2| \geqslant \delta^{1-1/(t-1)}(q-1)^{1/(t-1)}.$$

This result does not immediately yield any bound on the number of roots $R(f)$ since there is no upper bound given for the size of the H_2-cosets. However, if for some reason we were assured that the set of roots was a union of only H_1-cosets, we could conclude that

$$R(f) \leqslant \delta \cdot 2\left(\frac{q-1}{\delta}\right)^{1-1/(t-1)} = 2(q-1)^{1-1/(t-1)}\delta^{1/(t-1)},$$

which is an improvement on Theorem 1.1 since it can be easily checked that $\delta(f) \leqslant D(f)$ always.

THEOREM 2.2. *For $f(x) = c_1x^{a_1} + \ldots + c_tx^{a_t} \in \mathbb{F}_q[x]$ (with c_i nonzero), define*

$$S(f) := \{k \mid (q-1) : \text{for all } i, \text{there is a } j \neq i \text{ with } a_i \equiv a_j \bmod k\}.$$

If f vanishes completely on a coset of size k, then $k \in S(f)$.

Proof. For some generator g of $\mathbb{F}_q^*$, let $\alpha\langle g^{(q-1)/k}\rangle$ denote a coset of the unique subgroup of order k in $\mathbb{F}_q^*$, and let $\beta = \alpha^k$. The members of this coset are exactly the roots of the binomial $x^k - \beta$. So, f vanishes completely on this coset if and only if $(x^k - \beta) \mid f$, or equivalently if $f \equiv 0 \bmod (x^k - \beta)$.

To see when this happens, we view f in the ring $\mathbb{F}_q[x]/\langle x^k - \beta\rangle$. In this ring, we have the relation $x^k \equiv \beta$, so if each a_i has remainder r_i mod k, then

$$f \equiv c_1\beta^{\lfloor a_1/k\rfloor}x^{r_1} + \ldots + c_t\beta^{\lfloor a_t/k\rfloor}x^{r_t} \mod (x^k - \beta).$$

Now f might be identically zero (in this ring) since the r_i are not necessarily distinct. However, there is one obvious barrier to this: if just one r_i is unique, then f in particular contains the nonzero monomial $(c_i\beta^{\lfloor a_i/k\rfloor})x^{r_i}$. $f \equiv 0$ requires that each remainder r_i has at least one 'partner' $r_j = r_i$ so that monomials can cancel. Therefore $(x^k - \beta) \mid f$ implies that, for each $i \in \{1, 2, \ldots, t\}$, there is some $j \neq i$ with $a_i \equiv a_j$ mod k. □

https://doi.org/10.1112/S1461157016000334 Published online by Cambridge University Press

Thus $S(f)$ lists the sizes of cosets on which f might possibly vanish completely. For example, if $a_1 = 0$ and the other exponents $a_{i>1}$ are all prime to $q-1$ then $S(f) = \{1\}$, and so it is structurally impossible for f to vanish completely on any nontrivial coset, regardless of choice of coefficients $c_i \in \mathbb{F}_q^*$. On the other hand, whenever $k \in S(f)$, there is some choice of $c_i \in \mathbb{F}_q^*$ so that f does indeed vanish completely on a given coset of size k.

When $\max(S) < \delta^{1-1/(t-1)}(q-1)^{1/(t-1)}$, Theorem 2.2 can be combined with Theorem 2.1 to get a bound on $R(f)$ by ruling out the possibility of H_2-cosets. If $\max(S)$ is any larger, Theorem 2.1 is no longer helpful; the most we can conclude is that

$$R(f) \leqslant |H_2| \cdot 2\left(\frac{q-1}{\delta}\right)^{1-1/(t-1)} \leqslant \max(S) \cdot 2\left(\frac{q-1}{\delta}\right)^{1-1/(t-1)},$$

which is worse than trivial: $R(f) < q-1$. However, $S(f)$ turns out also to be independently useful for deriving sparsity-dependent bounds.

Theorem 2.3. *Let $f(x) = c_1x^{a_1} + \ldots + c_tx^{a_t} \in \mathbb{F}_q[x]$ (with c_i nonzero), let $\delta(f)$ be defined as above, and let $C(f)$ denote the size of the largest coset in $\mathbb{F}_q^*$ on which f vanishes completely. If $R(f)$ denotes the number of distinct, nonzero roots of f in $\mathbb{F}_q^*$, then we have*

$$R(f) \leqslant 2(q-1)^{1-1/(t-1)}C^{1/(t-1)},$$

and furthermore if $C < \delta^{1-1/(t-1)}(q-1)^{1/(t-1)}$, then

$$R(f) \leqslant 2(q-1)^{1-1/(t-1)}\delta^{1/(t-1)}.$$

This result is a strict improvement on Theorem 1.1, since, as we will see, $D(f)$ is in particular an upper bound for $S(f)$ and therefore also for $C(f)$. In fact, we can get another easily computable upper bound for $S(f)$ that is in general tighter than $D(f)$.

Proposition 2.4. *For $f(x) = c_1x^{a_1} + \ldots + c_tx^{a_t} \in \mathbb{F}_q[x]$ define the parameters*

$$\begin{aligned}
\delta(f) &= \gcd(a_1, a_2, \ldots, a_t, q-1),\\
D(f) &= \min_i \max_{j\neq i}(\gcd(a_i - a_j, q-1)),\\
Q(f) &= \gcd_i \operatorname{lcm}_{j\neq i}(\gcd(a_i - a_j, q-1)),\\
K(f) &= \min_i \max_{j\neq i}(\gcd(a_i - a_j, Q)).
\end{aligned}$$

These parameters relate to $S(f)$ as follows.

- *$\delta(f) \in S(f)$.*
- *For all $k \in S(f)$, $k \mid Q(f)$.*
- *$D(f)$, $Q(f)$, and $K(f)$ are all upper bounds for $S(f)$, and $K(f) \leqslant \min(D(f), Q(f))$.*

3. *Optimality of the bound*

Since the polynomial which defines a given function on $\mathbb{F}_q^*$ is unique only up to equivalence mod $x^{q-1} - 1$, we restrict our attention to polynomials with degree less than $q-1$. Thus we fix the following notation:

- $\mathcal{F}(q) = \{f \in \mathbb{F}_q[x] : \deg f < q-1\}$, $\quad\mathcal{F}_1(q) = \{f \in \mathcal{F}(q) : C(f) \leqslant 1\}$;
- $\mathcal{F}(q,t) = \{f \in \mathcal{F}(q) : f \text{ has } t \text{ terms}\}$, $\quad\mathcal{F}_1(q,t) = \{f \in \mathcal{F}_1(q) : f \text{ has } t \text{ terms}\}$;
- $\mathcal{F}(q,t,r) = \{f \in \mathcal{F}(q,t) : R(f) = r\}$, $\quad\mathcal{F}_1(q,t,r) = \{f \in \mathcal{F}_1(q,t) : R(f) = r\}$.

https://doi.org/10.1112/S1461157016000334 Published online by Cambridge University Press

Recall that $C(f) \leqslant 1$ indicates that f does not vanish on any entire coset of any nontrivial subgroup of $\mathbb{F}_q^*$.

In this section, we consider the possibility that the bound in Theorem 2.3 can be improved. Consider the binomial $f(x) = x^{(q-1)/2} + 1$. When q is odd, this binomial vanishes at every nonsquare in $\mathbb{F}_q^*$, and consequently $R(f) = C(f) = (q-1)/2$. More generally, when $q-1$ is divisible by t, the t-nomial $f(x) = (x^{q-1} - 1)/(x^{(q-1)/t} - 1)$ vanishes on $t-1$ cosets of size $(q-1)/t$, and so $R(f) = (q-1)(1-1/t)$. These examples show that there is no hope of improving the bound in Theorem 2.3 by removing the dependence on $C(f)$. However, the proportion of polynomials with $C(f) > 1$ is small, so it may be worthwhile to search for improved bounds for $f \in \mathcal{F}_1(q)$.

Proposition 3.1.

$$\frac{|\mathcal{F}(q) \backslash \mathcal{F}_1(q)|}{|\mathcal{F}(q)|} = O\left(\frac{1}{q}\right).$$

Proof. If $f \in \mathcal{F}(q)$ vanishes on a nontrivial coset, then it vanishes on a coset of prime order. Thus we can bound the number of such $f \in \mathcal{F}(q)$ by counting polynomials of the form

$$(x^p - \beta) \sum_{j=0}^{q-2-p} c_j x^j,$$

where p divides $q-1$ and β lies in the subgroup of $\mathbb{F}_q^*$ of size $(q-1)/p$. Thus the proportion of f with $C(f) > 1$ is bounded by

$$\frac{1}{|\mathcal{F}(q)|} \sum_{p \mid q-1} \left(\frac{q-1}{p}\right) q^{q-1-p} \leqslant \sum_{p \mid q-1} q^{1-p} \leqslant q^{-1} + \sum_{\substack{p \mid q-1 \\ p>2}} q^{1-p} = \frac{1}{q} + O\left(\frac{\log q}{q^2}\right).$$

Note that we have used the well-known fact that the number of distinct prime factors of an integer n is bounded by $\log n$. □

In [4], the authors investigate the existence of sparse polynomials with many roots. When q is a tth power, they give the t-nomial $f(x) = 1 + \sum_{i=1}^{t-1} x^{(q^{i/t}-1)/(q^{1/t}-1)} \in \mathbb{F}_q[x]$, which has $R(f) \geqslant q^{1-2/t}$. Furthermore, when t is prime they show that $D(f) \leqslant t/2$. In the case where q is an odd square, the authors of [6] give the trinomial $f(x) = x^{q^{1/2}} + x - 2$ which has $R(f) = q^{1/2}$, and it is shown that $C(f) = 1$. Thus, both examples are able to attain a large number of roots in $\mathbb{F}_q^*$ without vanishing on a large coset, and they show that the $O(q^{1-1/(t-1)})$ bound from Theorem 2.3 is nearly optimal in the general setting. However, these examples both share a special property: they vanish on entire translations of a subspace of $\mathbb{F}_q$. We are unaware of any example of a sparse polynomial which has a large number of roots in 'general position'. Consequently, we propose that a much better bound is possible for the special case of prime fields $\mathbb{F}_p$, which have no proper subfields.

Let $R_{p,t} = \max\{R(f) : f \in \mathcal{F}_1(p,t)\}$. Obviously $R_{p,1} = 0$ and $R_{p,2} = 1$, because monomials have no roots in $\mathbb{F}_p^*$, and a binomial defines a coset in $\mathbb{F}_p^*$ if it has a root at all. We have checked by computer that the following inequalities hold:

- $R_{p,3} < 1.8 \log p$ for $p \leqslant 139\,571$;
- $R_{p,4} < 2.5 \log p$ for $p \leqslant 907$;
- $R_{p,5} < 2.9 \log p$ for $p \leqslant 101$.

Therefore, the current bound of $R_{p,t} = O(p^{1-1/(t-1)})$ appears to be far from optimal for t-nomials over $\mathbb{F}_p$ which do not vanish on a nontrivial coset.

https://doi.org/10.1112/S1461157016000334 Published online by Cambridge University Press

It is easy to see that the proportion of polynomials $f \in \mathcal{F}(p)$ which have $R(f) = r$ is bounded by $1/r!$. Indeed, simply count the proportion of polynomials of the form

$$\left(\prod_{i=1}^{r}(x-\alpha_i)\right)\left(\sum_{i=0}^{p-2-r} c_i x^i\right),$$

with $\alpha_i \in \mathbb{F}_p^*$ distinct, which gives

$$\frac{\binom{p-1}{r}p^{p-1-r}}{|\mathcal{F}(p)|} = \binom{p-1}{r}\frac{1}{p^r} \leqslant \frac{1}{r!}.$$

With this in mind, we propose that the observed logarithmic behavior of $R_{p,t}$ can be explained by the following heuristic. Let $t(f)$ denote the number of nonzero terms of f. Then $R(f)$ and $t(f)$ are statistically independent properties of a random $f \in \mathcal{F}_1(p)$. This heuristic does not hold precisely, but it motivates the following conjecture, which captures the sentiment while allowing for some error.

Conjecture 3.2. There exists a constant $\gamma > 0$ such that

$$\frac{|\mathcal{F}_1(p,t,r)|}{|\mathcal{F}_1(p,t)|} \leqslant \left(\frac{1}{r!}\right)^{\gamma}$$

for all p prime, $t \in \mathbb{N}$, and $r \in \mathbb{N}$.

We have checked by computer that the inequality in Conjecture 3.2 holds with $\gamma = 1/2$ for all $r \in \mathbb{N}$ in the following cases:

- $t = 3$, $p \leqslant 30\,977$;
- $t = 4$, $p \leqslant 907$;
- $t = 5$, $p \leqslant 101$.

Theorem 3.3. *If Conjecture 3.2 is true, then* $R_{p,t} = O(t \log p)$.

Proof. Suppose Conjecture 3.2 is true. Then we have

$$|\mathcal{F}_1(p,t,r)| \leqslant |\mathcal{F}_1(p,t)|\left(\frac{1}{r!}\right)^{\gamma} \leqslant p^{2t}/(r!)^{\gamma}.$$

If $p^{2t}/(r!)^{\gamma} < 1$ then the set $\mathcal{F}_1(p,t,r)$ is empty, so we must have $p^{2t}/(R_{p,t}!)^{\gamma} \geqslant 1$, or equivalently, $\log(R_{p,t}!) \leqslant \log(p^{2t/\gamma})$. Applying Stirling's approximation, we get

$$R_{p,t} \leqslant R_{p,t}\log R_{p,t} \sim \log(R_{p,t}!) \leqslant (2/\gamma)t\log p = O(t\log p). \qquad \square$$

For a more detailed account of the computational and heuristic support for the conjectural logarithmic bound in the case of trinomials, see [4, 6].

4. *Proofs*

The general strategy employed here (and in [2, 3]) for obtaining sparsity-dependent bounds on $R(f)$ can be loosely sketched as follows. Consider integers e prime to $q-1$, which have the property that the map $x \mapsto x^e$ is a bijection on $\mathbb{F}_q^*$. Since $x \mapsto x^e$ simply permutes the elements of $\mathbb{F}_q^*$, we have $R(f(x)) = R(f(x^e))$. Furthermore, $f(x^e)$ is equivalent (as a mapping on $\mathbb{F}_q^*$) to any $g(x) = c_1x^{b_1} + \ldots + c_tx^{b_t}$ with $b_i \equiv ea_i \bmod (q-1)$. Thus the basic idea is to find some

https://doi.org/10.1112/S1461157016000334 Published online by Cambridge University Press

e so that the remainders of $ea_i \bmod (q-1)$ are all small, yielding a g of small degree, and so $R(f) = R(g) \leqslant \deg(g)$.

The following lemma, a fact about the geometry of numbers, will be our main tool for achieving the desired degree reduction.

LEMMA 4.1. *Fix the natural numbers* $a_1, a_2, \ldots, a_t, N$. *If* $n \leqslant N/\gcd(a_1, a_2, \ldots, a_t, N)$, *there is an* $e \in \{1, 2, \ldots, n-1\}$ *and a* $v \in N\mathbb{Z}^t$ *so that*

$$0 < \max_{1 \leqslant i \leqslant t} |ea_i + v_i| \leqslant N/n^{1/t}.$$

Proof. Consider the vectors $l_i = i(a_1, \ldots, a_t) = (ia_1, \ldots, ia_t) \in (\mathbb{R}/N\mathbb{Z})^t$ for $i \in \{1, 2, \ldots n\}$. Let $\|\cdot\|_\infty$ denote the standard infinity norm on $\mathbb{R}^t$. We wish to view these vectors geometrically as points in $\mathbb{R}^t$, but they are only defined up to equivalence in $(\mathbb{R}/N\mathbb{Z})^t$, so define

$$\|l\|_N = \min_{v \in N\mathbb{Z}^t} \|l + v\|_\infty,$$

which gives the smallest norm of any representative of the equivalence class $l + N\mathbb{Z}^t$ viewed as a point in $\mathbb{R}^t$ (equivalently, $\|l\|_N$ gives the distance from l to the nearest lattice point in $N\mathbb{Z}^t$). Suppose that

$$d = \min_{i \neq j} \|l_j - l_i\|_N.$$

Since the vectors are all at least d apart, the sets

$$B_i = \{x \in (\mathbb{R}/N\mathbb{Z})^t : \|x - l_i\|_N < d/2\}$$

are disjoint, so each l_i sits in its own personal box of volume d^t. We may choose to represent these n disjoint sets uniquely in the fundamental domain $[0, N)^t$, which has volume N^t. Therefore we have a total volume of $n \cdot d^t$ sitting in a volume of N^t; we conclude that $d \leqslant N/n^{1/t}$.

Note that the modular definition of distance is crucial here; consider instead n points in $[0, N]^t$ that are d-separated only in the standard l_∞ metric. A volume-packing argument becomes more complicated in this case because the box around a point near the boundary may lie partly outside $[0, N]^t$ (it does not ‘wrap around’), and so some points do not absorb a full d^t worth of volume from $[0, N]^t$.

To finish, we find i, j (with $1 \leqslant i < j \leqslant n$) so that $\|l_j - l_i\|_N = d$ and set $l_e = l_{(j-i)} = (j-i)(a_1, \ldots, a_t) = l_j - l_i$. We have

$$\|l_e\|_N = \min_{v \in N\mathbb{Z}^t} \|(ea_1, \ldots, ea_t) + v\|_\infty \leqslant N/n^{1/t},$$

and e satisfies $1 \leqslant e \leqslant n - 1$. The subgroup of $(\mathbb{Z}/N\mathbb{Z})^t$ generated by $(a_1, \ldots, a_n)$ has order $N/\gcd(a_1, \ldots, a_t, N) \geqslant n$. Since $0 < e < n$, $e(a_1, \ldots, a_n) \not\equiv (0, \ldots, 0) \in (\mathbb{Z}/N\mathbb{Z})^t$, which verifies that $\|l_e\|_N > 0$. □

Lemma 4.1 and its proof are extremely similar in spirit to the argument used by Canetti *et al.* in [3]. They also viewed the n vectors as points in $[0, N)^t$, but to find a pair of nearby points they partitioned the hypercube into less than n equal-sized sub-cubes and appealed to the pigeonhole principle. Here we were able to avoid this discretization of space which led to the small asymptotic term appearing in Theorem 1.1, which turns out to be unnecessary.

Proof of Theorem 2.3. The second claim is immediate from Theorem 2.1, since there can be no H_2-cosets of roots. We now prove the first claim.

https://doi.org/10.1112/S1461157016000334 Published online by Cambridge University Press

Let $f(x) = c_1x^{a_1} + c_2x^{a_2} + \ldots + c_tx^{a_t} \in \mathbb{F}_q[x]$ with c_i nonzero, and let C denote the size of the largest coset in $\mathbb{F}_q^*$ on which f vanishes completely. For our purposes, we may assume that $a_1 = 0$, since otherwise we can write

$$\begin{aligned} f(x) &= x^{a_1}\tilde{f}(x),\\ \tilde{f}(x) &= c_1 + c_2x^{a_2-a_1} + \ldots + c_tx^{a_t-a_1}, \end{aligned}$$

showing that f has a root at zero, but its nonzero roots are just the roots of $\tilde{f}$, so $R(f) = R(\tilde{f})$ and $C(f) = C(\tilde{f})$. Therefore we continue assuming that $a_1 = 0$.

Consider $\delta(f) = \gcd(a_2, \ldots, a_t, q-1)$. The nonzero roots of $f(x) = c_1 + c_2x^{a_2} + \ldots + c_tx^{a_t}$ are in one-to-one correspondence with the solutions of the system

$$\begin{gathered} c_1 + c_2y^{a_2/\delta} + \ldots + c_ty^{a_t/\delta} = 0 \quad y \in \langle g^\delta \rangle,\\ x^\delta = y \quad x \in \mathbb{F}_q^*. \end{gathered}$$

If f has no roots in $\mathbb{F}_q^*$ then our bound is of course true, so suppose this system has at least one solution (y_0, x_0). Then in fact the system has at least δ solutions and f vanishes on the coset $\{x : x^\delta = y_0\}$. This allows us to conclude that $C \geqslant \delta$, and so $((q-1)/C) \leqslant (q-1)/\gcd(a_2, \ldots, a_t, q-1)$.

Therefore we can apply Lemma 4.1 to find an $e \in \{1, 2, \ldots, (q-1)/C - 1\}$ and a $v \in (q-1)\mathbb{Z}^{t-1}$ so that

$$0 < \|(ea_2, \ldots, ea_t) + v\|_\infty \leqslant (q-1) \Big/ \left(\frac{q-1}{C}\right)^{1/(t-1)}.$$

Suppose $k = \gcd(e, q-1) = 1$. Then the mapping $x \mapsto x^e$ is a bijection on $\mathbb{F}_q^*$ that simply reorders the elements of $\mathbb{F}_q^*$, thus $R(f(x)) = R(f(x^e))$. We are interested in $f(x^e)$ only as a function on $\mathbb{F}_q^*$ (rather than as a formal object in $\mathbb{F}_q[x]$), and since $\mathbb{F}_q^*$ is a group of order $q-1$, this function is not changed by shifting its exponents by $v_i \in (q-1)\mathbb{Z}$. Thus we may represent the function $f(x^e)$ as the (possibly Laurent) polynomial

$$f(x^e) = c_1 + c_2x^{ea_2+v_2} + \ldots + c_tx^{ea_t+v_t},$$

which satisfies

$$0 < M = \max_{1 \leqslant i \leqslant t} |ea_i + v_i| \leqslant (q-1)^{1-1/(t-1)}C^{1/(t-1)}.$$

Again we are only interested in nonzero roots; note that $R(f(x^e)) = R(x^Mf(x^e))$. Since $x^Mf(x^e)$ is an honest polynomial in $\mathbb{F}_q[x]$ with nonnegative exponents, we have $R(f) = R(x^Mf(x^e)) \leqslant \deg(x^Mf(x^e)) \leqslant 2M$ and we are done.

However, we might have $k = \gcd(e, q-1) > 1$. In this case $x \mapsto x^e$ is not a bijection: it takes $\mathbb{F}_q^* = \langle g \rangle$ to a smaller subgroup $\langle g^e \rangle = \langle g^k \rangle$ of size $((q-1)/k)$. However, we can still cover $\mathbb{F}_q^*$ by k cosets of this subgroup. We have

$$R(f(x^e)) = \sum_{i=0}^{k-1} \frac{1}{k} R(f(g^ix^e)),$$

since $\mathbb{F}_q^* = \bigcup_{i=0}^{k-1} g^i\langle g^e \rangle$, and $x^e = y$ has k solutions for each $y \in \langle g^e \rangle$. Now we repeat our earlier tricks and arrive at

$$R(f) \leqslant \sum_{i=0}^{k-1} \frac{1}{k} \deg(x^Mf(g^ix^e)) \leqslant 2M,$$

https://doi.org/10.1112/S1461157016000334 Published online by Cambridge University Press

except that we must be careful that no $f(g^i x^e)$ is identically zero, preventing us from using degree to bound root number. If $f(g^i x^e)$ is identically zero then f vanishes completely on the coset $g^i\langle g^e\rangle = g^i\langle g^k\rangle$ of size $((q-1)/k)$. However, since $k = \gcd(e, q-1) \leqslant e < ((q-1)/C)$, we have

$$\frac{q-1}{k} > \frac{q-1}{((q-1)/C)} = C,$$

so this is impossible by the definition of C; the cosets are too large for f to vanish on completely. □

Proof of Proposition 2.4. For $f(x) = c_1x^{a_1} + \ldots + c_tx^{a_t} \in \mathbb{F}_q[x]$, we have the following equivalent definitions for S:

$$\begin{aligned} S(f) &= \{k \mid (q-1) : \forall i, \exists j \neq i \text{ such that } a_i \equiv a_j \bmod k\} \\ &= \{k \mid (q-1) : \forall i, \exists j \neq i \text{ such that } k \mid (a_i - a_j)\} \\ &= \{k \in \mathbb{N} : \forall i, \exists j \neq i \text{ such that } k \mid \gcd(a_i - a_j,\ q-1)\} \\ &= \bigcap_{i=1}^{t} \bigcup_{j \neq i} \{k \in \mathbb{N} : k \mid \gcd(a_i - a_j,\ q-1)\}. \end{aligned}$$

Clearly by the second definition we have $\delta(f) = \gcd(a_1, a_2, \ldots, a_t, q-1) \in S$. From the fourth definition we can get an upper bound for S by passing to the superset

$$\bigcap_{i=1}^{t} \bigcup_{j \neq i} \{k \in \mathbb{N} : k \leqslant \gcd(a_i - a_j,\ q-1)\} \supseteq S,$$

which has maximal element

$$D = \min_i \max_{j \neq i}(\gcd(a_i - a_j, q-1)).$$

Alternatively, by considering a different lattice structure on the integers, we can pass to the superset

$$\bigcap_{i=1}^{t} \{k \in \mathbb{N} : k \mid \gcd(L_j, q-1)\} = \{k \in \mathbb{N} : k \mid Q\} \supseteq S,$$

where

$$\begin{aligned} L_i &= \operatorname{lcm}(a_i - a_1, \ldots, a_i - a_{i-1}, a_i - a_{i+1}, \ldots, a_i - a_t), \\ Q &= \gcd(L_1, \ldots, L_t, q-1) = \gcd_i \operatorname{lcm}_{j \neq i}(\gcd(a_i - a_j, q-1)). \end{aligned}$$

Since we now know that, in the end, any member of S must be a divisor of Q, we can redefine S (equivalently) using a smaller ambient space:

$$\begin{aligned} S(f) &= \{k \mid Q : \forall i, \exists j \neq i \text{ such that } k \mid (a_i - a_j)\} \\ &= \bigcap_{i=1}^{t} \bigcup_{j \neq i} \{k \in \mathbb{N} : k \mid \gcd(a_i - a_j, Q)\} \\ &\subseteq \bigcap_{i=1}^{t} \bigcup_{j \neq i} \{k \in \mathbb{N} : k \leqslant \gcd(a_i - a_j, Q)\}. \end{aligned}$$

Considering the maximal element of this last superset of S gives the final upper bound

$$K = \min_i \max_{j \neq i}(\gcd(a_i - a_j, Q)),$$

which is obviously no larger than either D or Q. □

https://doi.org/10.1112/S1461157016000334 Published online by Cambridge University Press

Acknowledgements. I would like to thank my advisor, Dr J. M. Rojas, for introducing me to this problem, and the Texas A&M Supercomputing Facility for access to computational resources.

References

1. A. Balog, K. Broughan and I. E. Shparlinski, 'On the number of solutions of exponential congruences', *Acta Arith.* 148 (2010) no. 1, 93–103.

2. J. Bi, Q. Cheng and J. M. Rojas, 'Sub-linear root detection, and new hardness results, for sparse polynomials over finite fields', *Proceedings of ISSAC (International Symposium on Symbolic and Algebraic Computation,* June 26–29, Boston*)* (ACM Press, New York, 2013) 61–68.

3. R. Canetti, J. B. Friedlander, S. Konyagin, M. Larsen, D. Lieman and I. E. Shparlinski, 'On the statistical properties of Diffie–Hellman distributions', *Israel J. Math.* 120 (2000) 23–46.

4. Q. Cheng, S. Gao, J. M. Rojas and D. Wan, 'Sparse univariate polynomials with many roots over finite fields', Preprint, 2014, arXiv:1411.6346.

5. J. Friedlander, M. Larsen, D. Lieman and I. E. Shparlinski, 'On the correlation of binary M-sequences', *Des. Codes Cryptogr.* 16 (1999) no. 3, 249–256.

6. Z. Kelley and S. Owen, 'Estimating the number of roots of trinomials over finite fields', *J. Symbolic Comput.* (special issue on the topics of MEGA 2015), to appear; arXiv:1510.01758.

7. I. E. Shparlinski, 'Sparse polynomial approximation in finite fields', *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing* (ACM Press, New York, 2001) 209–215.

8. I. E. Shparlinski, 'On the singularity of generalised Vandermonde matrices over finite fields', *Finite Fields Appl.* 11 (2005) no. 2, 193–199.

Zander Kelley
Texas A&M University
College Station, TX 77843
USA

zander_k@tamu.edu

https://doi.org/10.1112/S1461157016000334 Published online by Cambridge University Press