

Elementary bounds on character sums with polynomial arguments

Topics in Finite Fields (Fall 2023)

University of Toronto

Swastik Kopparty

Last modified: Tuesday 17th October, 2023

Let $P(X) \in \mathbb{F}_q[X]$ be a polynomial of degree at most d . What can we say about the distribution of values of $P(x)$ as x varies in \mathbb{F}_q ? It is reasonable to expect the values to be spread out all over \mathbb{F}_q , landing in any set S with probability about $|S|/q$.

One important way of measuring the uniformity of distribution¹ is through the character sums:

$$\left| \sum_{x \in \mathbb{F}_q} \psi(P(x)) \right|,$$

$$\left| \sum_{x \in \mathbb{F}_q} \chi(P(x)) \right|.$$

Over this course, we will see several theorems showing that these sums are all small (unless there is an obvious reason for them not to be!), and several ways of taking advantage of these theorems.

1 Single Monomials

Suppose $P(X) = X^d$. Let ψ be a nontrivial additive character. What can we say about the absolute value of the sum:

$$S = \sum_{x \in \mathbb{F}_q} \psi(P(x)).$$

Without loss of generality, we can assume d divides $q-1$, otherwise we can replace d with $d' = \text{GCD}(d, q-1)$. Looking more closely, we see that we have basically already found a bound for this sum. Since the map $x \mapsto x^d$ is d -to-1 on \mathbb{F}_q^* , we conclude that:

$$S = 1 + d \cdot \sum_{y \in H} \psi(y),$$

where H is the multiplicative subgroup of \mathbb{F}_q^* consisting of nonzero perfect d' th powers. Thus, from the bound we derived using Gauss sums:

$$|S| \leq 1 + d\sqrt{q}.$$

Again, for small degrees d , this is the kind of cancellation that we would expect from a sum of q random points on the unit circle.

Looking at the whole argument in totality, what we just did is highly highly dependent on $P(X)$ being a single monomial. The multiplicativity of the monomial brings the characters of the multiplicative group into play. For a general polynomial $P(X)$, we have no such option.

We now give a different argument for bounding $|S|$. This argument will be amenable to generalization to all polynomials.

¹Note that we are not saying that if x is picked uniformly from \mathbb{F}_q , the distribution of $P(x)$ is uniform over \mathbb{F}_q (or even close to uniform in, say, the ℓ_1 distance).

What we are saying is more like this: for some natural collection \mathcal{T} of test functions $T : \mathbb{F}_q \rightarrow \mathbb{C}$, the expected value of $T(P(x))$ for x picked uniformly from \mathbb{F}_q is approximately equal the the expected value of $T(y)$ for y picked uniformly from \mathbb{F}_q . For example, one could take \mathcal{T} to be the set of all indicator functions of arithmetic progressions in \mathbb{F}_q , or to be the set of all additive characters of \mathbb{F}_q .

1.1 $P(X) = X^2$

First let $P(X) = X^2$. Let $S = \sum_{x \in \mathbb{F}_q} \psi(x^2)$. If q is even, then S is easily seen to be 0, so we assume q is odd.

We proceed by direct expansion:

$$\begin{aligned} |S|^2 &= \left(\sum_{x \in \mathbb{F}_q} \psi(x^2) \right) \cdot \left(\sum_{y \in \mathbb{F}_q} \overline{\psi(y^2)} \right) \\ &= \sum_{x, y \in \mathbb{F}_q} \psi(x^2 - y^2), \\ &= \sum_{x, y \in \mathbb{F}_q} \psi((x - y)(x + y)). \end{aligned}$$

Now we note that as (x, y) varies over \mathbb{F}_q^2 , $(x - y, x + y)$ also varies over \mathbb{F}_q^2 (this is where we used the oddness of q). Thus:

$$\begin{aligned} |S|^2 &= \sum_{a, b \in \mathbb{F}_q} \psi(ab) \\ &= \sum_{a \in \mathbb{F}_q} \begin{cases} q & a = 0 \\ 0 & a \neq 0 \end{cases} \\ &= q. \end{aligned}$$

Thus $|S| = \sqrt{q}$.

1.2 $P(X) = X^d$

Now let us generalize this to general $d|(q-1)$.

If we proceed as before, we get:

$$\begin{aligned} |S|^2 &= \sum_{x, y \in \mathbb{F}_q} \psi(x^d - y^d) \\ &= \sum_{x, y \in \mathbb{F}_q} \psi \left(\prod_{j=0}^{d-1} (x - \omega^j y) \right), \end{aligned}$$

where ω is a primitive d th root of 1 in \mathbb{F}_q . Now what?

The expression we just encountered has no easy simplification. Instead, we will consider a family of related expressions (the family includes our expression). We will be able to find the sum of all of these expressions!

Let $S_a = \sum_{x \in \mathbb{F}_q} \psi(ax^d)$. Consider the following sum:

$$\begin{aligned} \sum_{a \in \mathbb{F}_q} |S_a|^2 &= \sum_{a \in \mathbb{F}_q} \sum_{x, y \in \mathbb{F}_q} \psi(a(x^d - y^d)) \\ &= \sum_{x, y \in \mathbb{F}_q} \sum_{a \in \mathbb{F}_q} \psi(a(x^d - y^d)) \\ &= \sum_{x, y \in \mathbb{F}_q} q \cdot 1_{x^d = y^d} \\ &\leq q^2 d. \end{aligned}$$

This is an instructive calculation. Let us pause to see what it means. The average value of $|S_a|^2$ is dq , and thus:

1. for some a , $|S_a| \leq \sqrt{dq}$,
2. for some $a \neq 0$, $|S_a| \geq \sqrt{(d-1)q}$ (since $S_0 = q$).

On the other hand, it does not say anything directly about our desired sum, S_1 (it only says that $|S_1| \leq \sqrt{d} \cdot q$, which is trivial).

The final piece is the observation that many of the S_a are equal to S_1 . Indeed, if a is a nonzero perfect d th power, then $S_1 = S_a$. Thus:

$$\sum_{a \in \mathbb{F}_q^*} |S_a|^2 \geq \frac{(q-1)}{d} |S_1|^2,$$

and so:

$$|S_1|^2 \leq \frac{q^2(d-1)}{(q-1)/d} \leq d^2 q.$$

This gives us the desired bound $|S_1| \leq d\sqrt{q}$.

2 Mordell's bound

We now generalize the above argument to arbitrary polynomials $P(X)$ of degree $\leq d$. Let ψ be a nontrivial additive character. For a given $\mathbf{a} = (a_1, \dots, a_d) \in \mathbb{F}_q^d$, let:

$$S_{\mathbf{a}} = \sum_{x \in \mathbb{F}_q} \psi\left(\sum_{i=1}^d a_i x^i\right).$$

Let ℓ be an integer to be determined later. We expand and simplify:

$$\begin{aligned} \sum_{\mathbf{a} \in \mathbb{F}_q^d} |S_{\mathbf{a}}|^{2\ell} &= \sum_{\mathbf{a} \in \mathbb{F}_q^d} \prod_{j=1}^{\ell} \left(\sum_{x_j \in \mathbb{F}_q} \psi\left(\sum_{i=1}^d a_i x_j^i\right) \right) \cdot \prod_{k=1}^{\ell} \left(\sum_{y_k \in \mathbb{F}_q} \overline{\psi}\left(\sum_{i=1}^d a_i y_k^i\right) \right) \\ &= \sum_{\mathbf{a} \in \mathbb{F}_q^d} \sum_{x_1, \dots, x_{\ell}} \sum_{y_1, \dots, y_{\ell}} \psi\left(\sum_{i=1}^d a_i \left(\sum_{j=1}^{\ell} x_j^i - \sum_{k=1}^{\ell} y_k^i\right)\right) \\ &= \sum_{x_1, \dots, x_{\ell}} \sum_{y_1, \dots, y_{\ell}} \sum_{\mathbf{a} \in \mathbb{F}_q^d} \prod_{i=1}^d \psi\left(a_i \left(\sum_{j=1}^{\ell} x_j^i - \sum_{k=1}^{\ell} y_k^i\right)\right) \\ &= \sum_{x_1, \dots, x_{\ell}} \sum_{y_1, \dots, y_{\ell}} \prod_{i=1}^d \left(\sum_{a_i \in \mathbb{F}_q} \psi\left(a_i \left(\sum_{j=1}^{\ell} x_j^i - \sum_{k=1}^{\ell} y_k^i\right)\right) \right) \\ &= \sum_{x_1, \dots, x_{\ell}} \sum_{y_1, \dots, y_{\ell}} q^d C_d(x_1, \dots, x_{\ell}, y_1, \dots, y_{\ell}), \end{aligned}$$

where

$$C_d(x_1, \dots, x_{\ell}, y_1, \dots, y_{\ell}) = \begin{cases} 1 & \text{for each } i \leq d, \sum_{j=1}^{\ell} x_j^i = \sum_{k=1}^{\ell} y_k^i, \\ 0 & \text{otherwise.} \end{cases}$$

The following argument works only if $\text{char}(\mathbb{F}_q) > d$. In the homework, you will be asked to generalize this argument to all q .

Let $p_i(z_1, \dots, z_d) = \sum_{j=1}^d z_j^i$. By the Newton identities, the sequence $(p_i(z_1, \dots, z_d))_{i=1}^d$ determines all the d elementary symmetric functions $\sigma_1, \dots, \sigma_d$ of z_1, \dots, z_d (this is where the characteristic of \mathbb{F}_q gets used).

This in turn determines the multiset $R = \{z_1, \dots, z_d\}$, since R is the multiset of roots of the polynomial:

$$X^c + \sum_{j=1}^d (-1)^j \cdot \sigma_j \cdot X^j.$$

Thus, as long as $\ell \geq d$,

$$\begin{aligned} \sum_{\mathbf{a} \in \mathbb{F}_q^d} |S_{\mathbf{a}}|^{2\ell} &= \sum_{x_1, \dots, x_\ell} \sum_{y_1, \dots, y_\ell} q^d C_d(x_1, \dots, x_\ell, y_1, \dots, y_\ell) \\ &\leq \sum_{x_1, \dots, x_\ell} \sum_{y_1, \dots, y_{\ell-d}} q^d \cdot d! \\ &= q^{2\ell} \cdot d! \end{aligned}$$

Again, this tells us that the average value of $|S_{\mathbf{a}}|^{2\ell}$ equals $q^{2\ell-d} \cdot d!$, which suggests that the typical $|S_{\mathbf{a}}|$ is at most around $q^{1-\frac{d}{2\ell}} \cdot (d!)^{1/2\ell}$. But for any particular $|S_{\mathbf{a}}|$ we get nothing nontrivial directly from this.

The key is to note that for every $y \in \mathbb{F}_q^*$, if we let $\mathbf{a}_y = (a_1 y, a_2 y^2, \dots, a_d y^d)$, then $S_{\mathbf{a}_y} = S_{\mathbf{a}}$. Furthermore, if $\mathbf{a} \neq 0$, then there are at most d choices of $y \in \mathbb{F}_q$ for which $\mathbf{a}_y = \mathbf{a}$.

Thus, for every particular nonzero $\mathbf{a} \in \mathbb{F}_q^d$, we have

$$\frac{q-1}{d} |S_{\mathbf{a}}|^{2\ell} \leq q^{2\ell} \cdot d!,$$

and so

$$|S_{\mathbf{a}}| \leq (d! \cdot d)^{\frac{1}{2\ell}} \cdot q^{1-\frac{1}{2\ell}}.$$

Taking $\ell = d$, we get the following bound.

Theorem 1 (Mordell Bound). *Let ψ be a nontrivial additive character of \mathbb{F}_q , and let $P(X)$ be a nonzero polynomial of degree d , where $d < \text{char}(\mathbb{F}_q)$. Then:*

$$\left| \sum_{x \in \mathbb{F}_q} \psi(P(x)) \right| \leq O(d \cdot q^{1-\frac{1}{2d}}).$$

Note that once we allow $d \geq \text{char}(\mathbb{F}_q)$, it is possible that the above sum equals q . For example, if $P(X) = X^{\text{char}(\mathbb{F}_q)} - X$, then $\psi_1(P(x)) = 0$ for all $x \in \mathbb{F}_q$.

3 Kloosterman sums

For $a, b \in \mathbb{F}_q$, and ψ and additive character of \mathbb{F}_q , the Kloosterman sums are defined by:

$$K_{a,b} = \sum_{x \in \mathbb{F}_q^*} \psi_a(x) \psi_b(x^{-1}).$$

We will now give Kloosterman's elementary bound on his sums. The argument will be analogous to Mordell's method:

$$\begin{aligned} \sum_{a,b \in \mathbb{F}_q} |K_{a,b}|^4 &= \sum_{a,b} \sum_{x_1, x_2, y_1, y_2 \in \mathbb{F}_q^*} \psi_a(x_1 + x_2 - y_1 - y_2) \psi_b(x_1^{-1} + x_2^{-1} - y_1^{-1} - y_2^{-1}) \\ &= \sum_{x_1, x_2, y_1, y_2 \in \mathbb{F}_q^*} \left(\sum_a \psi_a(x_1 + x_2 - y_1 - y_2) \right) (\psi_b(x_1^{-1} + x_2^{-1} - y_1^{-1} - y_2^{-1})) \\ &= \sum_{x_1, x_2, y_1, y_2} q^2 C(x_1, x_2, y_1, y_2) \end{aligned}$$

where $C(x_1, x_2, y_1, y_2) = 1$ if $x_1 + x_2 = y_1 + y_2$ and $x_1^{-1} + x_2^{-1} = y_1^{-1} + y_2^{-1}$, and $C(x_1, x_2, y_1, y_2) = 0$ otherwise. Observe that $C(x_1, x_2, y_1, y_2) = 1$ if and only if $\{x_1, x_2\} = \{y_1, y_2\}$. Thus:

$$\begin{aligned} \sum_{a,b \in \mathbb{F}_q} |K_{a,b}|^4 &= \sum_{x_1, x_2, y_1, y_2} q^2 C(x_1, x_2, y_1, y_2) \\ &\leq \sum_{x_1, x_2} q^2 \cdot 2 \\ &= 2q^4. \end{aligned}$$

Finally, we note that $K_{a,b} = K_{ac, bc^{-1}}$ for every $c \in \mathbb{F}_q^*$. Thus:

Theorem 2. *If a, b are both nonzero, then: $|K_{a,b}| \leq O(q^{3/4})$.*

Note that the other Kloosterman sums are easily computed: $K(0, 0) = q - 1$, $K(0, b) = K(a, 0) = -1$ if $a, b \neq 0$.

4 Statement of the Weil bounds

The Weil bounds show that additive character sums with polynomial arguments (of the type we saw above) and multiplicative character sums with polynomial arguments have very high cancellation, except when they obviously don't.

Theorem 3. *Let ψ be a nontrivial additive character of \mathbb{F}_q , and let $P(X)$ be a nonzero polynomial of degree d , such that $P(X)$ cannot be written as $Q(X)^{\text{char}(\mathbb{F}_q)} - Q(X) + c$, where $Q(X) \in \mathbb{F}_q[X]$, $c \in \mathbb{F}_q$. Then:*

$$\left| \sum_{x \in \mathbb{F}_q} \psi(P(x)) \right| \leq (d-1)\sqrt{q}.$$

Theorem 4. *Let χ be a nontrivial multiplicative character of \mathbb{F}_q , and let $P(X)$ be a nonzero polynomial of degree d , such that $P(X)$ cannot be written as $cQ(X)^e$, where $Q(X) \in \mathbb{F}_q[X]$, $c \in \mathbb{F}_q$ and $e > 0$ is the smallest integer with $\chi^e = \chi_0$. Then:*

$$\left| \sum_{x \in \mathbb{F}_q} \chi(P(x)) \right| \leq d\sqrt{q}.$$

Finally, we also have the Weil bound on the Kloosterman sums:

Theorem 5. *For a, b nonzero, we have:*

$$|K_{a,b}| \leq 2\sqrt{q}.$$

To get a flavor of how we will prove these bounds, consider the quadratic residue character χ , and consider what the statement $\left| \sum_{x \in \mathbb{F}_q} \chi(P(x)) \right| \leq t$ means. It means that $P(x)$ is a quadratic residue for between $\frac{q-t}{2}$ and $\frac{q+t}{2}$ values of $x \in \mathbb{F}_q$. This is equivalent to saying that the number of $(x, y) \in \mathbb{F}_q$ satisfying equation:

$$y^2 = P(x)$$

is $q \pm t$. Thus we want to count the number of \mathbb{F}_q^2 solutions of some given bivariate polynomial $Q(X, Y)$. This is the language in which Weil proved his bound; every *absolutely irreducible* $Q(X, Y) \in \mathbb{F}_q[X, Y]$ of degree d has $q \pm O_d(\sqrt{q})$ solutions in \mathbb{F}_q^2 . $Q(X, Y)$ is said to be absolutely irreducible if it is irreducible in $\overline{\mathbb{F}_q}[X, Y]$ (i.e., even if the coefficients of the potential factors are allowed to come from the algebraic closure $\overline{\mathbb{F}_q}$). For example, $X^2 - \alpha Y^2$ is not absolutely irreducible, even if α is a quadratic nonresidue, and $X^2 + Y^2 - 1$ is absolutely irreducible.

5 A quick application

Let q be prime, and let S, T be APs in \mathbb{F}_q . (We may also let q be general, and S, T be subspaces; then with a similar argument we get slightly better bounds, as in the section on Gauss sums).

How many $x \in S$ are there such that $x^{-1} \in T$?

Let us express this analytically, and then simplify using what we know.

$$\begin{aligned}
|\{x \in S \mid x^{-1} \in T\}| &= \sum_{x \in S} 1_T(x^{-1}) \\
&= \sum_{x \in \mathbb{F}_q} 1_S(x) 1_T(x^{-1}) \\
&= \sum_{x \in \mathbb{F}_q} \sum_{a, b \in \mathbb{F}_q} \hat{1}_S(a) \psi_a(x) \hat{1}_T(b) \psi_b(x^{-1}) \\
&= \sum_{a, b \in \mathbb{F}_q} \sum_{x \in \mathbb{F}_q} \hat{1}_S(a) \psi_a(x) \hat{1}_T(b) \psi_b(x^{-1}) \\
&= \sum_{a, b \in \mathbb{F}_q} \hat{1}_S(a) \hat{1}_T(b) \left(\sum_{x \in \mathbb{F}_q} \psi_a(x) \psi_b(x^{-1}) \right) \\
&= \hat{1}_S(0) \hat{1}_T(0) + \sum_{a, b \in \mathbb{F}_q, (a, b) \neq (0, 0)} \hat{1}_S(a) \hat{1}_T(b) \left(\sum_{x \in \mathbb{F}_q} \psi_a(x) \psi_b(x^{-1}) \right) \\
&= \frac{|S||T|(q-1)}{q^2} \pm \left(\sum_{a \in \mathbb{F}_q} |\hat{1}_S(a)| \right) \left(\sum_{b \in \mathbb{F}_q} |\hat{1}_T(b)| \right) \cdot (2\sqrt{q}) \\
&= \frac{|S||T|}{q} \pm O(\sqrt{q} \log^2 q).
\end{aligned}$$

(As usual, if S, T were subspaces, we would save the $\log^2 q$ factor.

Thus for $|S||T| \gg q^{3/2} \log^2 q$, the number of $x \in S$ with $x^{-1} \in T$ is what we would expect from random sets S, T .

One weakness of the above result is that it says nothing about very small APs / subspaces. Consider the following problem: could it be that there are intervals S of length $\log^{100} q$ such that for every $x \in S$, we also have $x^{-1} \in S$? This is consistent with the above result.

We will see in a later class (hopefully) that this cannot be: there is an absolute constant C , such that for all primes $q > C$ and all intervals $S \subseteq \mathbb{F}_q$ with $|S| > C$, we have:

$$|\{x \in S \mid x^{-1} \in S\}| \leq 0.99|S|.$$

This will follow from a more general statement: the 3-regular graph $G = (V, E)$, where $V = \mathbb{F}_q$ and $E = \{(x, y) \in \mathbb{F}_q^2 \mid x - y = \pm 1 \text{ or } xy = 1\}$ is an *expander graph*.