Weil Image Sums
(and some related problems)

September 26, 2011
ver. 1.04

Joshua E. Hill
e-mail: hillje@math.uci.edu
Department of Mathematics,
University of California,
Irvine, CA 92697-3875

Abstract

We investigate a series of related problems in the area of incomplete Weil sums where the sum is run over a set of points that produces the image of the polynomial. We establish a bound for such sums, and establish some numerical evidence for a conjecture that this sum can be bounded in a way similar to Weil's bounding theorem.

To aid in the average case, we investigate the problem of the cardinality of the value set of a positive degree polynomial (degree $d > 0$) over a finite field with p^m elements. We provide naïve algorithms for calculating this cardinality, and analyze their computational complexity. We then show a connection between this cardinality and the number of points on a family of varieties in affine space. We couple this with Lauder and Wan's p-adic point counting algorithm, resulting in a nontrivial algorithm for calculating this cardinality. The computational complexity of this algorithm is polynomial in both p and m, but is exponential in d.

We seek to develop a corresponding set of bounds for incomplete exponential sums based on the image set's algebraic structure, including algorithms to explicitly calculate and estimate these sums. We also seek to further study the cardinality of the image set of polynomials in finite fields, including algorithms for calculating and estimating this value.

The analysis of these algorithms' computational complexity may better establish the computational complexity class for these problems.

Contents

1

Introduction

Exponential sums have been a useful tool in combining the areas of Number Theory and Algebraic Geometry, particularly in connection to Dirichlet L-functions and zeta-functions (e.g., [Davenport, 2000], [Hardy and Wright, 2000, §5.6]). We examine a generalization of a particular form of exponential sum, the Weil Sum. Weil sums have several additional applications in analytic number theory (e.g., [Davenport and Erdös, 1952], [Bateman et al., 1950]) and computer science, particularly in certain problems involving graph theory (in particular Paley graphs, as in [Chung, 1989] and [Chung et al., 1989]) and cryptography (in particular, for random number generators, as in [Babai et al., 1992], [Chor and Goldreich, 1988], and [Zuckerman, 1990]).

In a finite field with q elements, denoted $\mathbb{F}_q$ (with $q = p^m$, p prime), take a positive degree polynomial $f \in \mathbb{F}_q[x]$ with degree d.

The current discussion is restricted to additive characters into the complex units (i.e., a group homomorphism from the additive group of the field to the complex units: $\mathbb{F}_q \to \mathbb{C}^*$). These are all of the form:

$$\psi_\gamma(\alpha) = e^{\frac{2\pi i}{p}\mathrm{Tr}(\gamma\alpha)}$$

for some $\gamma \in \mathbb{F}_q$ and with the trace being the absolute trace

$$\mathrm{Tr}(\alpha) = \mathrm{Tr}_{\mathbb{F}_p}^{\mathbb{F}_q}(\alpha) = \sum_{j=0}^{m-1} \alpha^{p^j}.$$

A Weil sum is (for some fixed γ and f) a sum of the form

$$\sum_{\beta \in \mathbb{F}_q} \psi_\gamma(f(\beta)).$$

An "incomplete" exponential sum is an exponential sum taken over some proper subset of $\mathbb{F}_q$. We'll be taking this sum over some minimal set that produces the full image of f. To this end, define:

$$V_f = \{f(\beta) \mid \beta \in \mathbb{F}_q\}$$
$$S_{f,\gamma} = \sum_{\alpha \in V_f} \psi_\gamma(\alpha)$$

We denote the maximum such sum for non-trivial characters as

$$|S_f| = \max_{\gamma \in \mathbb{F}_q^*} |S_{f,\gamma}|.$$

We start with an analog of Weil's bounding theorem (Theorem 1).

Conjecture 1 (Wan)**.** *For all polynomials of degree d, with $p \nmid d$:*

1.1 *There is a real number c_d so that we have $|S_f| \leq c_d \sqrt{q}$.*[1]

1.2 $c_d \leq c\sqrt{d}$.

1.3 $c \leq 1$.

One immediate question for any sum over the image is how to efficiently establish the image set. In the case where we want to work with the average case, it would also be helpful to determine the cardinality of V_f. As it happens, this value is rather difficult to efficiently calculate. We propose and analyze several algorithms for establishing $\#\left(V_f\right)$ exactly; we first provide two classes of naïve approaches to the problem, and then provide an algorithm based on the p-adic point counting algorithm developed in [Lauder and Wan, 2008].

We seek to develop a corresponding set of bounds for incomplete exponential sums based on the algebraic structure of the image set and to estimate the cardinality of polynomials over finite fields; these may lead to further refinement for bounds on these exponential sums. In all cases, we intend to develop explicit algorithms to perform these tasks. These algorithms' computational complexity will be analyzed; these results may better establish the computational complexity class for these problems.

[1] Conjecture 1.1 is true so long as we have $q \gg d$ as a consequence of [Cohen, 1970], the Chebotarev density theorem, and unpublished results of Wan and Lenstra.

2 Literature Survey

2.1 WEIL SUMS

There is a *vast* body of literature on exponential sums; these have been studied in various contexts for a very long time. Gauss used them in his *Disquisitiones Arithmeticae* [Gauss, 1801] (indeed all of the sums we examine can be thought of as specializations of these Gauss sums). Exponential sums of the form we consider were considered by Hermann Weyl, where the function f was required to be smooth ([Weyl, 1914] and [Weyl, 1916]).

Conjecture 1 is analogous to a theorem of André Weil [Weil, 1948]:

Theorem 1 (Weil). *Let $f(x) \in \mathbb{F}_q[x]$ be of degree $d > 1$ with $p \nmid d$ and let ψ be a non-trivial additive character of $\mathbb{F}_q$. Then*

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

There are many results that address Weil sums of certain polynomials of special forms that lead to various specialized bounds (e.g., $ax^n + b$, quadratics, p-linear polynomials, etc.; see [Lidl and Niederreiter, 1997, Chp. 5, §4] for an excellent introduction). More recently, the value of Weil sums has been evaluated for polynomials of the form $ax^{p+1} + bx$ (in [Carlitz, 1980]) and $ax^{p^\alpha+1} + bx$ (in [Coulter, 1998a] and [Coulter, 1998b]).

Bombieri and Sperber examine sums of the same form as we are examining (in [Bombieri and Sperber, 1995]), but instead sum over the points in quasi-projective varieties.

There have also been several bounds and explicit values calculated in the instance where $q = p$, including incomplete sums over intervals (e.g., [Korobov, 1958]).

2.2 CARDINALITY OF IMAGE SETS

There are a few trivial bounds that can be immediately established; there are only q elements in the field, so $\#\left(V_f\right) \leq q$ (where $\#(\cdot)$ denotes the cardinality). Additionally, any polynomial of degree d can have at most d roots, thus for all $a \in V_f$, $f(x) = a$ is satisfied at most d times. This is true for every element in V_f, so $\#\left(V_f\right) d \geq q$, whence

$$\left\lceil \frac{q}{d} \right\rceil \leq \#\left(V_f\right) \leq q$$

(where $\lceil \cdot \rceil$ is the ceiling function).[1]

Both of these bounds can be achieved: if $\#\left(V_f\right) = q$, then f is called a "permutation polynomial" and if $\#\left(V_f\right) = \lceil q/d \rceil$, then f is said to have a "minimal value set".

One way of exploring the behavior of $\#\left(V_f\right)$ is to look at asymptotic results that apply for many or most polynomials. Initial results by Uchiyama in [Uchiyama, 1954] showed that if

$$f^*(u, v) = \frac{f(u) - f(v)}{u - v} \tag{2.1}$$

is absolutely irreducible, then $\#\left(V_f\right) > \frac{q}{2}$ for sufficiently large p; [Carlitz, 1955] then showed that the requirement that (2.1) be absolutely irreducible could not be dropped. In [Uchiyama, 1955], the average value for $\#\left(V_f\right)$ was established in terms of a value

$$\mu_d = 1 - \frac{1}{2!} + \frac{1}{3!} - \cdots + \frac{(-1)^{d-1}}{d!}.$$

This is clearly just a power series expansion of $\left(1 - e^{-1}\right)$, so as $d \to \infty$, μ_d quickly converges to this value. The *average* value across all polynomials was then seen to be

$$\#\left(V_f\right) \sim \mu_d q + O(1).$$

In [Birch and Swinnerton-Dyer, 1959], Birch and Swinnerton-Dyer made this estimate more concrete for a class of polynomials that they somewhat optimistically called "general polynomials" (that is those polynomials such that the Galois group of $f(x) - t$ over $\bar{\mathbb{F}}_q(t)$ is the symmetric group on d elements). So long as f is a general polynomial, we have $\mu = \mu_d$, and

$$\#\left(V_f\right) = \mu q + O_d\left(q^{1/2}\right).$$

[1] This lower bound is commonly written $\lfloor (q-1)/d \rfloor + 1$, likely in reference to [Carlitz et al., 1961], where this choice was convenient due to their specific naming conventions.

They also proved that μ depends only on d and two Galois groups:

$$G(f) = \mathrm{Gal}\left(f(x) - t/\mathbb{F}_q(t)\right)$$
$$G^+(f) = \mathrm{Gal}\left(f(x) - t/\bar{\mathbb{F}}_q(t)\right)$$

Cohen refined this in [Cohen, 1970] and provided an explicit statement for μ in terms of Galois groups. Let K be the splitting field for $f(x) - t$ over $\mathbb{F}_q(t)$ and $k' = K \cap \bar{\mathbb{F}}_q$. Finally, define:

$$G^*(f) = \left\{\sigma \in G(f) \mid K_\sigma \cap k' = \mathbb{F}_q\right\}$$
$$G_1(f) = \{\sigma \in G(f) \mid \sigma \text{ fixes at least one point}\}$$
$$G_1^*(f) = G_1(f) \cap G^*(f)$$

Cohen found that we then have $\mu = \#\left(G_1^*\right)/\#\left(G^*\right)$. This provides a wonderful combinatorial explanation of μ_d, as the proportion of nonderangements in S_d!

In [Voloch, 1989], Voloch showed that for general q, the Galois group condition described in [Birch and Swinnerton-Dyer, 1959] implies that the surface $f^*(x, y) = 0$ meets the smoothness requirement $d^2y/dx^2 \neq 0$, which he demonstrated was sufficient to provide a lower bound on $\#\left(V_f\right)$:

$$\#\left(V_f\right) \geq \frac{2q^2}{(d+1)q + (d-1)(d-2)}.$$

The problem of establishing $\#\left(V_f\right)$ has been studied in various forms for at least the last 115 years, but exact formulations for $\#\left(V_f\right)$ are known only for polynomials in very specific forms. The behavior of $\#\left(V_f\right)$ when f is constant or degree 1 is clear ($\#\left(V_f\right) = 1$ and $\#\left(V_f\right) = q$, respectively). Kantor partially solved the cubic case (mod 3) in [Kantor, 1915], and then Uchiyama in [Uchiyama, 1954] completely characterized $\#\left(V_f\right)$ for f of degree 2 ($p \neq 2$) or 3 ($p \neq 2, 3$).

For higher degree polynomials, exact formulae for $\#\left(V_f\right)$ are only known for polynomials in a few special forms. The special case of the p-linear polynomial is fairly straight forward: for linear operators, the size of the image is just the ratio of the total size of the space divided by the kernel of the map. Dickson Polynomials of the first kind have been well studied, and their image set is completely understood (this class includes the cyclic polynomial X^d. See [Chou et al., 1988] for details). In [Cusick, 1998], Cusick determines the exact value for $\#\left(V_f\right)$ for $f(X) = X^k(1 + X)^{2^m-1}$ in $\mathbb{F}_{2^{2m}}$, for $k = \pm 1, \pm 2,$ or 4 and then in [Cusick, 2005] for $f(X) = (X + 1)^d + X^d + 1$ for particular values of d over $\mathbb{F}_{2^m}$.

More is known about polynomials that fall into the special cases that we have already introduced: permutation polynomials (including exceptional polynomials) and polynomials with minimal value sets. There are few permutation polynomials known (indeed, permutation polynomials are asymptotically fairly sparse. A randomly selected polynomial is a permutation polynomial with probability e^{-q} for large q. [von zur Gathen, 1991a])

Dickson classified all permutation polynomials of degree less than or equal to six in his thesis (published in [Dickson, 1896/97]). Additional classes of permutation polynomials include certain parameter sections of Dickson Polynomials of the first and second kind, reversed Dickson Polynomials, Linearized Polynomials, and polynomials of the form $x^{(q+1)/2} + ax$. See [Lidl and Niederreiter, 1997, chp. 7] for a wonderful introduction on this topic.

Hayes moved the question of characterizing permutation polynomials into the realm of algebraic geometry in [Hayes, 1967] by noting that f is a permutation polynomial if and only if the variety defined over $\bar{\mathbb{F}}_q^2$ by $f^*(X, Y)$ only has $\mathbb{F}_q$-rational points on the diagonal $X = Y$. This approach became the study of exceptional polynomials, those polynomials such that the factorization of $f^*(X, Y)$ into irreducibles in $\mathbb{F}_q[X, Y]$ contains no absolutely irreducible terms (that is, each irreducible term in the factorization must *not* be irreducible in $\bar{\mathbb{F}}_q[X, Y]$). The characteristic of being an exceptional polynomial was recognized quite early as being very closely related to that of being a permutation polynomial. In [Cohen, 1970], Cohen proved that almost all exceptional polynomials were also permutation polynomials, and Wan removed the last special cases in [Wan, 1993]. A consequence of the Lang-Weil bound is that if $p \nmid d$, $d > 1$ and $q > d^4$, then any permutation polynomial of degree d is also an exceptional polynomial. Thus, for sufficiently large fields, the notions of permutation polynomial and exceptional polynomial are largely the same.

This characteristic of polynomials was used in [Ma and von zur Gathen, 1995a] to provide a ZPP (Zero-error Probabilistic Polynomial time) algorithm for testing a polynomial to determine if it is a permutation polynomial. In [Shparlinski, 1992], Shparlinski provided a fully deterministic test that determines if a given polynomial is a permutation polynomial by extending [von zur Gathen, 1991b] to an algorithm that runs in $\tilde{O}((nq)^{6/7})$ for all n and q (see section 3.2 for an explanation of the "Big-Oh" and "Soft-Oh" notations).

There are numerous results that provide bounding inequalities for $\#\left(V_f\right)$, average values for $\#\left(V_f\right)$ (summed across all polynomials up to degree $q-1$) and asymptotic results for $\#\left(V_f\right)$, but these largely do not lead to exact values for $\#\left(V_f\right)$. One notable exception is Wan's proof of Mullen's conjectured bound for non-permutation polynomials (conjectured in [Mullen, 1993], proven in in [Wan, 1993]):

$$\#\left(V_f\right) \leq \left\lfloor q - \frac{q-1}{d} \right\rfloor.$$

This bound was found to be sharp by Cusick and Müller in [Cusick and Müller, 1996] ($f(X) = (X+1)\,X^{q-1}$ achieves this bound). Thus, if any polynomial is found to have more distinct points in the image than allowed by this bound, then it must be a permutation polynomial.

A similar finding by Gomez-Calderon in [Gomez-Calderon, 1988] showed that if a low degree polynomial has a sufficiently small value set, then it must have a minimal value set. In particular, if f is a polynomial of degree $3 \leq d < p$ and

$$\#\left(V_f\right) \leq \left\lfloor \frac{q-1}{d} \right\rfloor + 2\left(\frac{q-1}{d^2}\right) - 1$$

then f has a minimal value set.

These two findings act to form "exclusion zones"; certain disallowed values for $\#\left(V_f\right)$ for polynomials of particular degrees.

Several families of polynomial with minimal value sets have been discovered. All polynomials with minimal value sets with degree $d < 2p+2$ were classified in [Carlitz et al., 1961] by Carlitz, Lewis, Mills, and Straus, and then Mills continued by further classifying all polynomials of degree $d \leq \sqrt{q}$ in [Mills, 1964].

Significant additional work in this area was performed by Javier Gomez-Calderon in his doctoral thesis [Gomez-Calderon, 1986], and then with various collaborators. In his thesis and in [Gomez-Calderon and Madden, 1988], he characterizes all polynomials of degree $d < \sqrt[4]{q}$ for which $\#\left(V_f\right) < 2q/d$; many of these polynomials result in forms based on Dickson polynomials.

As we have seen, $\#\left(V_f\right)$ is known in only very limited cases. Given these restricted results, one might get the false impression that polynomials must take only certain types of images; this is incorrect! To dispel this notion, note that one can construct a polynomial that takes any arbitrary value set by using Lagrange Interpolation.

2.3 p-ADIC POINT COUNTING

Examine the variety described by the simultaneous zeros of polynomials, $f_1, \ldots, f_m \in \mathbb{F}_q[x_1, \ldots, x_n]$ over the field $\bar{\mathbb{F}}_q$; call this variety X.

Each finite extension of $\mathbb{F}_q$ is isomorphic to a field of the form $\mathbb{F}_{q^k}$ for some k; denote the set of simultaneous zeros within the extension $\mathbb{F}_{q^k}$ as $X(\mathbb{F}_{q^k})$. We can now define the zeta function for our set of polynomials:

$$Z(X) = Z(X, T) = \exp\left(\sum_{k=1}^{\infty} \frac{\#\left(X(\mathbb{F}_{q^k})\right)}{k} T^k\right).$$

The zeta function clearly contains a profound amount of information about the polynomial set. We are point counting, that is we are looking for $\mathbb{F}_q$-rational solutions to our polynomial set, which is clearly just the value $X(\mathbb{F}_q)$! Thus, if we can calculate the zeta function, we immediately know how many $\mathbb{F}_q$-rational points are on the variety.[2] Indeed, one traditionally goes in the other direction, building up the zeta function through point counting [Wan, 1999].

From the definitions this seems like a less than useful statement, but surprises abound in mathematics! Weil conjectured that the zeta function is a rational function; this was first proven in [Dwork, 1960] using p-adic methods, and then later proven by ℓ-adic cohomological methods by Grothendieck [Grothendieck, 1964]. The common zeros of our polynomial set are not expected to form any particularly nice variety (non-singular projective, a curve, an abelian variety, etc.) so there are very few options for efficiently performing point counting or calculating $Z(X, T)$ explicitly.

For a polynomial of total degree d in n variables, Lauder and Wan described in [Lauder and Wan, 2008] an algorithm that explicitly calculates the zeta function of any such variety which runs in polynomial time so long as the characteristic is suitably small (on the order of $p = O((d \log q)^C)$ for some positive constant C; see [Wan, 2008] for details). They accomplish this by developing a toric point counting algorithm that can (by repeated application) piece together the total number of points for each $\#\left(X(\mathbb{F}_{q^k})\right)$. They then go on to determine the number of values that must be computed in order to uniquely determine the numerator and denominator of the zeta function.

[2] Similarly, if one had access to the totality of all human knowledge, one could successfully rebuild the Ancient Library of Alexandria!

2.4 ROLE OF COURSE WORK

The author's prior course work has quite a lot of relevance to this area of study, both directly and indirectly. The direct application of course series from the Algebra area of specialization include:

Algebraic Geometry introduced the notion of a variety, dimension, hypersurfaces, singularity, absolute irreducibility, certain classes of cohomology and the Riemann-Roch theorem.

Algebraic Number Theory included material on various classes of the classical exponential sums (Gauss, Jacobi, and Weil sums); these were used to help to develop certain of the results of class theory. The p-adic numbers and their absolute values were introduced, and various p-adic analysis techniques were used, including (Lang's generalization of) Hensel's lemma.

Common Tools include much of the framework of commutative algebra, category theory and some homological algebra.

3

Preliminary Results

These results are taken from joint work with Daqing Wan.

3.1 WEIL IMAGE SUM BOUNDS

In principle, we are looking at arbitrary polynomials in $\mathbb{F}_q[x]$ of the form $f(x) = a_d x^d + a_{d-1}x^{d-1} + \cdots + a_1 x + a_0$.

First note that if f happens to be a permutation polynomial, then $V_f = \mathbb{F}_q$ (and thus, $|S_f| = 0$), so we needn't consider this case.

One does not have to look at all polynomials of this form in order to find all the values of $|S_f|$. For any $\lambda \in \mathbb{F}_q$, $f(x-\lambda)$ is a polynomial that has the same image set as $f(x)$. Examine the x^{d-1} term of (the expanded) $f(x-\lambda)$; there are only two contributors to this term: the $a_d(x-\lambda)^d$ term contributes $-a_d d\lambda x^{d-1}$ and the $a_{d-1}(x-\lambda)^{d-1}$ term contributes $a_{d-1}x^{d-1}$ (by the binomial theorem). The x^{d-1} term is then $(a_{d-1} - da_d\lambda)x^{d-1}$. The choice of $\lambda = \frac{a_{d-1}}{da_d}$ then results in a polynomial with no x^{d-1} term (we know this is possible, as $f(x)$ was assumed to be of degree d). Thus we can, without loss of generality, assume that our initial polynomial had no x^{d-1} term.

We can proceed still further. Let $I_f \subset \mathbb{F}_q$ be a minimal pre-image for V_f (that is, $f(I_f) = V_f$ and $\#(I_f) = \#(V_f)$). We then have:

$$
\begin{aligned}
|S_f| &= \left| \sum_{\beta \in I_f} \psi_\gamma (f(\beta)) \right| \\
&= \left| \sum_{\beta \in I_f} \psi_\gamma \left(a_d\beta^d + a_{d-2}\beta^{d-2} + \cdots + a_1\beta + a_0 \right) \right| \\
&= \left| \sum_{\beta \in I_f} \psi_\gamma \left(a_d\beta^d + a_{d-2}\beta^{d-2} + \cdots + a_1\beta \right) \psi_\gamma(a_0) \right| \\
&= \left| \sum_{\beta \in I_f} \psi_{\gamma a_d} \left(\beta^d + \frac{a_{d-2}}{a_d}\beta^{d-2} + \cdots + \frac{a_1}{a_d}\beta \right) \right|.
\end{aligned}
$$

As we are already looking across all possible additive characters (and thus are taking the maximum across all possible γ), the a_d term in γa_d simply permutes the ordering of the characters; we have thus concluded that all possible $|S_f|$ values are encountered by simply examining polynomials of the form

$$f(x) = x^d + a_{d-2}x^{d-2} + \cdots + a_1 x \text{ with } a_i \in \mathbb{F}_q.$$

MAXIMUM BOUND FOR $|S_f|$

To determine what reasonable bounds on c_d are, we examine:

$$\Phi_d = \max_{\substack{f \in \mathbb{F}_q[x] \\ \deg f = d}} \frac{|S_f|}{\sqrt{q}}$$

$$\Phi = \max_{\substack{1 \le d < q \\ p \nmid d}} \Phi_d$$

across many choices of q; we then must have $c_d \ge \Phi_d$ for any particular choice of q.

Let us momentarily turn our eyes to such sums across point sets. We seek a set that, when summed across, produces the maximum such value across all $\gamma \in \mathbb{F}_q^*$. Note that if we find a set that produces a maximum for some fixed $\gamma \neq 0$, scaled versions of this set will produce the same maximum values for every other non-trivial additive character, so it suffices to examine only the $\gamma = 1$ case.

For a set $A \subset \mathbb{F}_q$, define

$$|S_A| = \left| \sum_{\alpha \in A} \psi_1(\alpha) \right|.$$

Finally, define the maximum such sum:

$$|S_{A_q}| = \max_{A \subset \mathbb{F}_q} |S_A|.$$

Note that this is not necessarily the same as $\sqrt{q}\Phi$: every subset of $\mathbb{F}_q$ is the image set of some polynomial (such a polynomial can be constructed using Lagrange interpolation, possibly followed by reduction mod $X^q - X$), but we are not guaranteed that p will not divide such a polynomial's degree. Examining small fields will show you that some point sets cannot be the image of such a polynomial, e.g., in $\mathbb{F}_4$, no degree 3 or 1 polynomial has the same image as $x^2 + x$.

Theorem 2. *If $q = p^m$ then*

$$\left|S_{A_q}\right| = \begin{cases} 2^{m-1} & p = 2 \\ \frac{p^{m-1}}{2} \csc\left(\frac{\pi}{2p}\right) & p \text{ odd} \end{cases}$$

Proof. The absolute trace is a $\mathbb{F}_p$-linear transform $\mathbb{F}_q \to \mathbb{F}_p$ that is surjective (note that $\mathrm{Tr}(x)$ is of degree p^{m-1}, thus has at most p^{m-1} zeros, but there are p^m possible inputs, so there are non-zero values of Tr. Tr is $\mathbb{F}_p$-linear, so any input producing a non-zero output can be scaled to produce any possible output.) $\#(\ker \mathrm{Tr}) = p^{m-1}$, so each additive coset of the form $\alpha + \ker \mathrm{Tr}$ is of this same size. Thus, to find a set $\hat{A} \subset \mathbb{F}_q$ with maximal $|S_{\hat{A}}|$, it suffices to find a set $A \subset \mathbb{F}_p$ that attains $|S_{A_p}|$. To extend this to $\hat{A}$, one needs only choose all the points in the corresponding cosets, resulting in a maximal value of $p^{m-1}|S_{A_p}|$.

It is trivial to note that $|S_{A_2}| = 1$, so we need only consider the odd prime case.

For $\mathbb{F}_p$, we are selecting pth roots of unity and adding them. Amusingly, computations involving this form are easily optimized by using a Discrete Fourier Transform (DFT), but we won't need to resort to this approach.

Lemma 1. *The sum of a non-zero complex number* $\mathbf{v}$ *(of modulus* ℓ*) and a complex number of modulus* 1 *at angle* θ *to* $\mathbf{v}$ *has modulus greater than* ℓ *if and only if* $\cos\theta > -\frac{1}{2\ell}$.

Proof. Without loss of generality, we can rotate the numbers so that $\mathbf{v} = \ell$, and we can then assume that the unit modulus number is $\cos\theta + i\sin\theta$. Thus we have

$$\begin{aligned} \ell < |\cos\theta + \ell + i\sin\theta| &\iff \\ \ell^2 < \cos^2\theta + 2\ell\cos\theta + \ell^2 + \sin^2\theta &\iff \\ 0 < 2\ell\cos\theta + 1 &\iff \\ -\frac{1}{2\ell} < \cos\theta. \end{aligned}$$

□

In our case, we further know that $\ell \geq 1$, so that tells us that the sum is never made greater by adding a unit modulus number whose angle is more than $\frac{2\pi}{3}$ away, and we always get a larger modulus by adding unit modulus numbers whose angle is less than $\frac{\pi}{2}$ away.

This suggests that any maximum modulus would involve at least all the roots of unity within a sweep of π radians. Examining all the roots of unity within the first and fourth quadrants, we would be adding

$$e^{\frac{2\pi ij}{p}} \text{ with } -\left\lfloor \frac{p}{4} \right\rfloor \leq j \leq \left\lfloor \frac{p}{4} \right\rfloor.$$

This sum is calculated in the following lemma:

Lemma 2. *For positive odd p,*

$$\sum_{j=-\lfloor p/4 \rfloor}^{\lfloor p/4 \rfloor} e^{\frac{2\pi ij}{p}} = \frac{1}{2} \csc\left(\frac{\pi}{2p}\right).$$

Proof. It helps notation somewhat to let $u = \lfloor p/4 \rfloor$. First, to show that the equality is correct:

$$\begin{aligned}
\sum_{j=-u}^{u} e^{\frac{2\pi ij}{p}} &= \sum_{j=0}^{u} \left(e^{\frac{2\pi i}{p}}\right)^j + \sum_{j=0}^{u} \left(e^{-\frac{2\pi i}{p}}\right)^j - 1 \\
&= -1 + \frac{e^{\frac{2\pi i(u+1)}{p}} - 1}{e^{\frac{2\pi i}{p}} - 1} + \frac{e^{-\frac{2\pi i(u+1)}{p}} - 1}{e^{\frac{-2\pi i}{p}} - 1} \\
&= \frac{e^{\frac{2\pi i(u+1)}{p}} - e^{-\frac{2\pi iu}{p}}}{e^{\frac{2\pi i}{p}} - 1}.
\end{aligned}$$

p is a positive odd integer, and can thus be represented as $p = 4j + r$ for some non-negative integer value of j and $r = 1$ or $r = 3$; both these proceed in the same way; we examine the case where $p \equiv 1 \pmod 4$:

Let $p = 4j + 1$, then $u = j$, leaving us with

$$\frac{e^{\frac{2\pi i(j+1)}{4j+1}} - e^{-\frac{2\pi ij}{4j+1}}}{e^{\frac{2\pi i}{4j+1}} - 1}.$$

Subtracting (the exponential form of) $\frac{1}{2} \csc\left(\frac{\pi}{2+8j}\right)$ leaves us with

$$\frac{e^{\frac{2i(1+j)\pi}{1+4j}} - ie^{\frac{3i\pi}{2+8j}}}{e^{\frac{2i\pi}{1+4j}} - 1} = \frac{e^{\frac{2i(1+j)\pi}{1+4j}} - e^{\frac{2i(1+j)\pi}{1+4j}}}{e^{\frac{2i\pi}{1+4j}} - 1}.$$

The numerator is thus 0, and the denominator is clearly never 0, so the difference is 0.

Thus, in both cases we have equality, so we have the desired result. □

It turns out that this is exactly the maximal modulus.

Lemma 3. *For odd p, $\frac{1}{2}\csc\left(\frac{\pi}{2p}\right)$ is the maximal modulus of any sum of distinct pth roots of unity.*

Proof. We start with the case examined above; we take the sum and check to see if we could get a larger modulus for the sum by adding the next root of unity; as its angle is greater than $\frac{\pi}{2}$, we can examine $\theta = \frac{\pi}{2} + \alpha$, so

$$\cos\theta = \cos\left(\frac{\pi}{2} + \alpha\right) = -\sin\alpha.$$

If we have $p = 4j + r$, then

$$\begin{aligned}\alpha &= \left\lceil \frac{p}{4} \right\rceil \frac{2\pi}{p} - \frac{\pi}{2} \\ &= \frac{\pi(4\left\lceil \frac{p}{4} \right\rceil - p)}{2p} \\ &= \frac{\pi(4-r)}{2p}.\end{aligned}$$

Thus, if we can exclude any advantage in the case where $p = 4j+3$, then we can conclude that we cannot increase the modulus of the sum by adding the next value.

From lemmas 1 and 2 , we had

$$\begin{aligned} -\frac{1}{2\ell} < \cos\theta \iff \\ -\frac{1}{2\left(\frac{1}{2}\csc\left(\frac{\pi}{8j+6}\right)\right)} < -\sin\alpha \iff \\ \sin\left(\frac{\pi}{8j+6}\right) > \sin\left(\frac{\pi}{8j+6}\right).\end{aligned}$$

This is clearly false, so we find that adding the next root never increases the modulus, though it may leave it unchanged.

If we had some other set of pth roots of unity that produced the maximal modulus, we have seen that it must at least include all the roots of unity within $\pi/2$ of it on the unit circle, and cannot possibly include any points further than $2\pi/3$ away on the unit circle. We can rotate these points to the position analyzed above, thus the modulus of the sum cannot be larger than the value above. Thus we have our result. □

The set of points that, when summed, produce the maximal modulus is clearly not unique. The underlying set of roots of unity can be rotated, maintaining its modulus. In addition, we saw above that in the case where $p \equiv 3 \pmod 4$, adding the root of unity just past $\pi/2$ (on either side) from the sum does not change the total modulus, so each of these is a free choice.

Combining these results, we find the maximal modulus for the general case, as stated. □

Corollary 1. *As $p \to \infty$ along the odd primes, $|S_{A_p}| \searrow \frac{p}{\pi}$*

Proof.

$$\left|\frac{1}{2}\csc\left(\frac{\pi}{2p}\right) - \frac{p}{\pi}\right| = \left|\frac{1}{2\sin\left(\frac{\pi}{2p}\right)} - \frac{p}{\pi}\right|$$
$$= \left|\frac{p}{\pi}\left[\frac{\left(\frac{\pi}{2p}\right)}{\sin\left(\frac{\pi}{2p}\right)}\right] - \frac{p}{\pi}\right|.$$

$\frac{1/x}{\sin 1/x} \searrow 1$ as $x \to \infty$ (with $x \geq 6/\pi$), so this difference approaches 0 monotonically from above. □

As immediate corollaries to our corollary, we can directly determine bounds.

Corollary 2. *If $q = p^m$ then as $p \to \infty$ along the odd primes, we have*

$$\left|S_{A_q}\right| \searrow \frac{q}{\pi}.$$

Next, by noting that $\frac{\csc\left(\frac{\pi}{2p}\right)}{2p}$ is a decreasing function in our domain of interest, we can extract non-asymptotic bounds:

Corollary 3. *If $q = p^m$, then*

$$\left|S_{A_q}\right| \leq \begin{cases} q/2 & p = 2 \\ q/3 & p \text{ odd} \end{cases}.$$

Figure 3.1 has a graph that depicts the maximal possible exponential sum modulus as q varies.

Table 3.2 summarizes the maximal bound for several field sizes. Note that these reflect the maximum value for all possible polynomials, not just for those whose degree is relatively prime to the characteristic of the field.

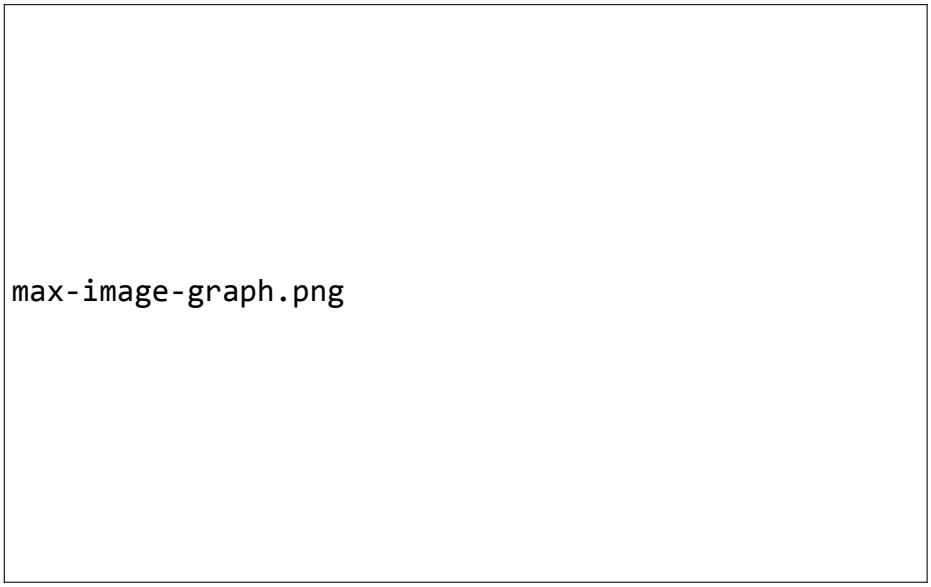

Figure 3.1: Maximum Exponential Sum Modulus Bound

Table 3.2: Maximal Additive Sums

q	$\lvert S_{A_q}\rvert$	$\Phi_d \leq$
2	1.0000	0.7070
3	1.0000	0.5770
4	2.0000	1.0000
5	1.6200	0.7240
7	2.2500	0.8490
8	4.0000	1.4100
9	3.0000	1.0000
11	3.5100	1.0600
13	4.1500	1.1500
16	8.0000	2.0000
17	5.4200	1.3100
19	6.0500	1.3900
23	7.3300	1.5300
25	8.0900	1.6200
27	9.0000	1.7300
29	9.2400	1.7100
31	9.8700	1.7700

q	$\lvert S_{A_q}\rvert$	$\Phi_d \leq$
32	16.0000	2.8300
37	11.8000	1.9400
41	13.1000	2.0400
43	13.7000	2.0900
47	15.0000	2.1800
49	15.7000	2.2500
53	16.9000	2.3200
59	18.8000	2.4500
61	19.4000	2.4900
64	32.0000	4.0000
67	21.3000	2.6100
71	22.6000	2.6800
73	23.2000	2.7200
79	25.1000	2.8300
81	27.0000	3.0000
83	26.4000	2.9000
89	28.3000	3.0000
97	30.9000	3.1400

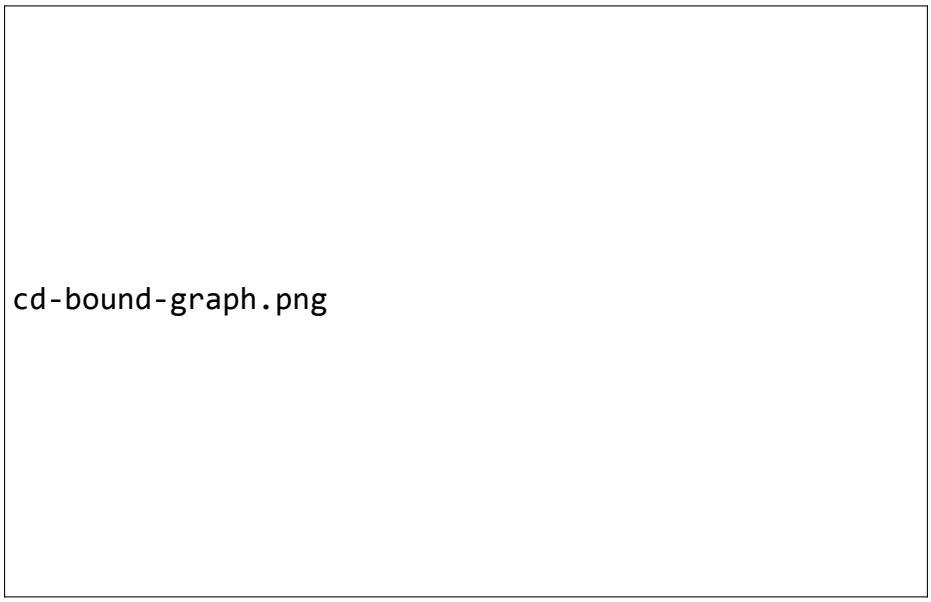

Figure 3.3: Minimum c_d Possible

NUMERICAL EVIDENCE

In attempting to seek evidence to support conjecture 1, we sought data. We estimated Φ_d by one of two techniques. When computationally feasible, Φ_d is calculated directly by calculating $|S_f|$ for every polynomial f of that degree of the form described above (we refer to this as "exhaustive search").

If exhaustive search is not computationally feasible, randomly chosen degree-d polynomials from $\mathbb{F}_q[x]$ of the form described above above are selected, and $|S_f|$ is calculated. We then use the maximal $|S_f|$ that was encountered as an estimate for Φ_d (we refer to this as "random sampling"). This random sampling process occurred in two rounds: in both rounds 1000 polynomials of each degree were sampled. All finite fields up to 100 elements were tested. This data is summarized in figure 3.3. The information is additionally normalized by dividing by a factor of $\sqrt{d}$ in figure 3.4.

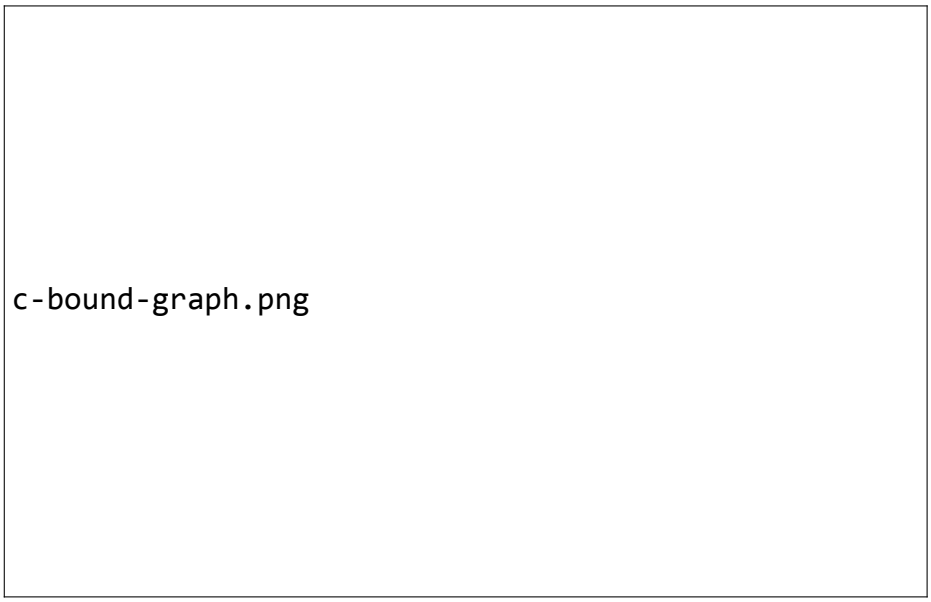

Figure 3.4: Minimum c Possible

3.2 IMAGE SET CARDINALITY

NAÏVE ALGORITHMS

To compare approaches, we'll use the "Big-Oh" and "Soft-Oh" notations. Let A and B be two eventually positive real valued functions $A, B : \mathbb{N}^k \to \mathbb{R}$ under $|\mathbf{x}|_{\min} = \min_i x_i$. $A(\mathbf{x}) = O(B(\mathbf{x}))$ if and only if there exists a positive real constant C and an integer N so that if $|\mathbf{x}|_{\min} > N$ then $A(\mathbf{x}) \leq CB(\mathbf{x})$. Similarly, $A(\mathbf{x}) = \tilde{O}(B(\mathbf{x}))$ if and only if there exists a positive real constant C' so that $A(\mathbf{x}) = O(B(\mathbf{x}) \log^{C'}(B(\mathbf{x}) + 3))$. "Soft-Oh" notation is used to dispense with log terms that might otherwise obscure the main thrust of "Big-Oh" notation.

There are several naïve methods of calculating $\#\left(V_f\right)$. Perhaps the most obvious method is to evaluate the polynomial at each point in $\mathbb{F}_q$ and count how many unique images result. This approach uses q evaluations, each of which can be evaluated using the Horner scheme [Knuth, 1998] in $2d - 1$ field multiplications, each in $O(m^{1+\lg 3} \lg^2 p)$ bit operations (here lg is the logarithm base 2), and d field additions, each in $O(m \lg p)$ bit operations.[1] The final counting can occur in $O(q)$, which is negligible in comparison to the other operations.

[1] Estimates of bit operations for arithmetic operations in $\mathbb{F}_q$ assume an iterated extension approach; see [Bach and Shallit, 1997, p. 348] for details.

Thus, our first naïve algorithm requires $O(qdm^{\lg 3-1} \lg^2 q)$ bit operations, or in "Soft-Oh" notation, $\tilde{O}(p^m d)$. This algorithm thus has polynomial complexity in d and p but exponential complexity in m. This algorithm is thus not polynomial in the input polynomial length, which in the dense polynomial model is assumed to be $O(d \lg q)$.

One can also approach this problem by operating on points in the co-domain. One has $f(x) = a$ for some $x \in \mathbb{F}_q$ if and only if $f_a(X) = f(X) - a$ has at least one linear factor. We can test for such factors by examining $\deg \gcd(f_a, X^q - X)$ (this is the first step of Rabin's irreducibility test from [Rabin, 1980]). This is computationally expensive for large q, so we instead examine $\deg \gcd(f_a, X^q - X \pmod{f_a})$, which is of the same degree.

Multiplication of polynomials of degree no greater than d can occur in $O(M(d))$ field operations, where $M(d) = d \log d \log\log d$. Modular reduction then requires $O(\lg q M(d))$ field operations, and the GCD calculation requires $O(\log d M(d))$ field operations. Repeating this process at most q times identifies the entire image set, requiring $O(q \lg q M(d))$ field multiplications. Combining, we get a complexity of $O(qdm^{\lg 3-1} \lg^3 q \log d \log\log d)$ bit operations, or in "Soft-Oh" notation, $\tilde{O}(p^m d)$.[2]

As a variant of this second algorithm, we could attempt to avoid doing repeated (expensive) modular reduction / GCD calculations, and instead work over $\mathbb{F}_q(t)$. To take this approach, we can replace the constant term with a transcendental term t (call this polynomial f_t) and then calculate the GCD only once.

The resulting polynomials are in $\mathbb{F}(t)[X]$ and have X-degree no greater than d. By applying the evaluation map $t \mapsto a$ to the GCD for each $a \in \mathbb{F}_q$, we would then know the cardinality of the image set. Each of these evaluation maps would require evaluation of at most d rational expressions (one for each positive power of X in the resulting polynomials). Applying this approach results in the same bounds with respect to field operations (but these field operations are more expensive). The resulting polynomial then has the evaluation map applied q times; each evaluation map results in the evaluation of at most d rational expressions, which makes this method considerably slower than either of the first two approaches. This method can lead to another possible approach; see chapter 4.

[2] If one was interested in estimating $\#(V_f)$, you could turn this algorithm into a probabilistic algorithm, as in [Ma and von zur Gathen, 1995b].

THE IMAGE SET AND POINT COUNTING

Theorem 3. *If $f \in \mathbb{F}_q[x]$ is a polynomial of degree $d > 0$, then the cardinality of its image set is*

$$\#\left(V_f\right) = \sum_{i=1}^{d} (-1)^{i-1} N_i \sigma_i \left(1, \frac{1}{2}, \ldots, \frac{1}{d}\right) \tag{3.1}$$

where $N_k = \#\left(\left\{(x_1, \ldots, x_k) \in \mathbb{F}_q^k \mid f(x_1) = \cdots = f(x_k)\right\}\right)$ and σ_i denotes the ith elementary symmetric function on d elements.

We present two proofs. The first proof suffices as a reasonable explanation of why the theorem is true, and the second is a more direct proof of correctness.

DERIVATION

Proof. Beginning as in [Uchiyama, 1954] and [Birch and Swinnerton-Dyer, 1959], we examine a family of subsets of V_f

$$V_{f,i} = \left\{x \in V_f \mid \#\left(f^{-1}(x)\right) = i\right\}, \quad 1 \leq i \leq d.$$

Each element $\beta \in V_f$ must have at least one pre-image (as if β had no points in its pre-image, it would not be in the image!) and can have at most d points in its pre-image (such pre-images are roots of the degree d polynomial $f(x) - \beta$, and there are at most d such roots.) Thus

$$V_f = \coprod_{1 \leq i \leq d} V_{f,i}$$

(where $\coprod$ denotes the disjoint union).

Proceeding as in both Uchiyama [Uchiyama, 1954] and Das [Das, 2003], denote the cardinality of each of these sets as $m_i = \#\left(V_{f,i}\right)$. Clearly,

$$m_1 + \ldots + m_d = \#\left(V_f\right). \tag{3.2}$$

We now count points in a particular subset of $\mathbb{A}^k_{\mathbb{F}_q}$; let

$$\tilde{N}_k = \left\{(x_1, \ldots, x_k) \in \mathbb{F}_q^k \mid f(x_1) = \cdots = f(x_k)\right\}.$$

We are generally going to be more interested in the number of elements of such sets; we have already denoted this as $N_k = \#\left(\tilde{N}_k\right)$. We'll categorize the points in $\tilde{N}_k$ by their (shared) image.

Starting with N_1, if $(\alpha_1) \in \tilde{N}_1$ with $f(\alpha_1) = \beta$, then β is in the image of f and so is in exactly one $V_{f,i}$. If $\beta \in V_{f,1}$, then there are m_1 distinct images, each of which must have a distinct pre-image, so there are m_1 choices for (α_1). If instead $\beta \in V_{f,2}$, then there are m_2 distinct images, each of which have exactly 2 distinct pre-images, so there would be $2m_2$ choices for (α_1). Similarly, if $\beta \in V_{f,\ell}$, then there are m_ℓ distinct images, each of which have exactly ℓ distinct pre-images, so there would be exactly ℓm_ℓ choices for (α_1). There can be no overlap between each of these cases, so we can then sum and find $N_1 = m_1 + 2m_2 + \ldots + d\,m_d$.

For N_k, if $(\alpha_1, \ldots, \alpha_k) \in \tilde{N}_k$ with $f(\alpha_1) = \beta$ and $\beta \in V_{f,\ell}$, then there are m_ℓ distinct images, each of which have exactly ℓ distinct pre-images, so there would be exactly ℓm_ℓ choices for α_1, and ℓ choices for each of $\alpha_2, \ldots, \alpha_k$, yielding a total of $\ell^k m_\ell$ choices for $(\alpha_1, \ldots, \alpha_k)$. Thus we see that in general

$$N_k = m_1 + 2^k m_2 + \cdots + d^k m_d. \tag{3.3}$$

Now, let us now introduce a new variable, say $\xi = -\#\left(V_f\right)$. We can then rewrite (3.2) to be $m_1 + \ldots + m_d + \xi = 0$, and (3.3) to $m_1 + 2^k m_2 + \cdots + d^k m_d + 0\xi = N_k$ with $1 \leq k \leq d$; this system of equations yields

$$\begin{pmatrix} 1 & 1 & \cdots & 1 & 1 \\ 1 & 2 & \cdots & d & 0 \\ 1 & 2^2 & \cdots & d^2 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & 2^d & \cdots & d^d & 0 \end{pmatrix} \begin{pmatrix} m_1 \\ m_2 \\ m_3 \\ \vdots \\ \xi \end{pmatrix} = \begin{pmatrix} 0 \\ N_1 \\ N_2 \\ \vdots \\ N_d \end{pmatrix}. \tag{3.4}$$

We then solve for ξ using Cramer's rule.

For Cramer's rule, we need two different determinants. First, we need the determinant of the $(d+1) \times (d+1)$ square matrix above, which we'll call A

Lemma 4.

$$\begin{aligned} \det A &= \det \begin{pmatrix} 1 & 1 & \cdots & 1 & 1 \\ 1 & 2 & \cdots & d & 0 \\ 1 & 2^2 & \cdots & d^2 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & 2^d & \cdots & d^d & 0 \end{pmatrix} \\ &= (-1)^d\, d!(d-1)!(d-2)! \cdots 2!1! \end{aligned}$$

Proof. For the determinant of A, we can expand along the last column and then factor out the common terms from each column:

$$\begin{aligned}\det A &= \det \begin{pmatrix} 1 & 1 & \cdots & 1 & 1 \\ 1 & 2 & \cdots & d & 0 \\ 1 & 2^2 & \cdots & d^2 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & 2^d & \cdots & d^d & 0 \end{pmatrix} \\ &= (-1)^{d+2} \det \begin{pmatrix} 1 & 2 & \cdots & d \\ 1 & 2^2 & \cdots & d^2 \\ \vdots & \vdots & \vdots & \vdots \\ 1 & 2^d & \cdots & d^d \end{pmatrix} \\ &= (-1)^d d! \det \begin{pmatrix} 1 & 1 & \cdots & 1 \\ 1 & 2 & \cdots & d \\ \vdots & \vdots & \vdots & \vdots \\ 1 & 2^{d-1} & \cdots & d^{d-1} \end{pmatrix}.\end{aligned}$$

This sub-matrix is the transpose of a Vandermonde matrix, so the determinant of the original matrix is:

$$\begin{aligned}\det A &= (-1)^d d! \prod_{1 \le i < j \le d} (j - i) \\ &= (-1)^d d!(d-1)!(d-2)! \cdots 2!1!\end{aligned}$$

□

We'll need the determinant of a new matrix for Cramer's rule. This new matrix, B, will be based on A, but with the last column replaced by the column vector on the right hand side of equation (3.4). Calculating this determinant will require a modest effort.

Lemma 5.

$$\begin{aligned}\det B &= \det \begin{pmatrix} 1 & 1 & \cdots & 1 & 0 \\ 1 & 2 & \cdots & d & N_1 \\ 1 & 2^2 & \cdots & d^2 & N_2 \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ 1 & 2^d & \cdots & d^d & N_d \end{pmatrix} \\ &= (d-1)!(d-2)! \cdots 2!1! \sum_{i=1}^{d} (-1)^{d+i} N_i \sigma_{d-i}(1, 2, \ldots, d)\end{aligned}$$

Proof. For the determinant of B, we have a somewhat similar looking determinant; again expanding along the last column:

$$\det B = \sum_{i=1}^{d} (-1)^{d+i} N_i M_{i+1,d+1}.$$

where $M_{i+1,d+1}$ is the corresponding minor for B.

Each of these $M_{i+1,d+1}$ are "simple alternants" (see [Muir and Metzler, 1960, Chapter XI] or alternately [Aitken, 1959, Chapter VI] for a slightly more readable account). If we generalize B by replacing the bases $(1, \ldots, d)$ with a corresponding unique variable $(X_1, \ldots, X_d)$ and write the (non-eliminated) powers as $\alpha = (\alpha_1, \ldots, \alpha_d)$ then denote:

$$\begin{aligned} a_\alpha &= \hat{M}_{i+1,d+1} \\ &= \det \left(x_n^{\alpha_m}\right)_{m,n=1}^{d} \text{ where } \alpha_m = \begin{cases} m-1 & 1 \le m < i+1 \\ m & i+1 \le m \le d \end{cases} \end{aligned}$$

Proceeding (roughly) as in [Stanley, 2001, p. 335], we define $\delta = (0, 1, \ldots, d-1)$, then

$$a_\alpha = a_{\lambda+\delta} = \det \left(x_n^{\lambda_m + (m-1)} \right)_{n,m=1}^{d}$$

which forces $\lambda = (\underbrace{0, \ldots, 0}_{i \text{ terms}}, \underbrace{1, \ldots, 1}_{d-i \text{ terms}})$.

It is evident that if $x_m = x_n$ for any $m < n$ then a_α is 0. This implies that $(x_n - x_m)$ divides a_α for all $1 \le m < n \le d$, thus a_α is divisible by a_δ (the Vandermonde determinant). The quotient $a_{\delta+\lambda}/a_\delta$ (called a "bialternant") is the historical definition of the Schur polynomial of shape λ:

$$s_\lambda (X_1, \ldots, X_d) = \frac{a_{\lambda+\delta} (X_1, \ldots, X_d)}{a_\delta (X_1, \ldots, X_d)}.$$

Comparing this to the more standard combinatorial definition of the Schur polynomials:

$$s_\lambda = \sum_\beta K_{\lambda\beta} x^\beta$$

where β runs over all weak compositions. (*i.e.*, start with an integer partition of $\ell = \sum_m \lambda_m$ padded with 0s to bring the partition length

to the same length as λ. The set of weak compositions are every possible ordering of every such partition.) Here, $K_{\lambda\beta}$ is the Kostka number, the number of semi-standard Young tableaux (SSYT) of shape λ and type β.

In this case, the form of λ causes all such tableaux to be a single column of length $d-i$; a tableau forms a valid SSYT only if the integers that fill the tableau strictly increase down the column. Each weak composition, β, establishes values that must be used to fill the tableau; there must be β_m total m's present in the tableau. As we are required to *strictly* increase down the column, this tells us that $K_{\lambda\beta} = 0$ for any β that contains any values other than 0 and 1, and there is exactly one way to arrange these numbers into our tableau: in increasing order. Thus

$$K_{\lambda\beta} = \begin{cases} 0 & \beta_m > 1 \text{ for any } i \\ 1 & \text{otherwise} \end{cases}$$

which suggests that each term in the sum s_λ has exactly $d-i$ distinct terms, and includes all possible arrangements. For this λ, we see that:

$$\begin{aligned} s_\lambda &= \sum_\beta K_{\lambda\beta} x^\beta \\ &= \sum_{1 \le j_1 < j_2 < \cdots < j_{d-i} \le d} X_{j_1} X_{j_2} \cdots X_{j_{d-i}} \\ &= \sigma_{d-i}\left(X_1, \ldots, X_d\right). \end{aligned}$$

That is, for this type of λ, s_λ is just the $(d-i)$th elementary symmetric polynomial on d variables, and thus

$$a_{\lambda+\delta}\left(X_1, \ldots, X_d\right) = \sigma_{d-i}\left(X_1, \ldots, X_d\right) a_\delta\left(X_1, \ldots, X_d\right).$$

We thus have:

$$\begin{aligned} M_{i+1,d+1} &= \sigma_{d-i}\left(1, \ldots, d\right) a_\delta\left(1, \ldots, d\right) \\ &= \sigma_{d-i}\left(1, \ldots, d\right) \prod_{1 \le m < n \le d} (n-m) \\ &= \sigma_{d-i}\left(1, \ldots, d\right) (d-1)!(d-2)! \cdots 2!1! \end{aligned}$$

Finally combining these results,

$$\det B = (d-1)!(d-2)! \cdots 2!1! \sum_{i=1}^{d} (-1)^{d+i} N_i \sigma_{d-i}\left(1, 2, \ldots, d\right).$$

□

Combining our results and applying Cramer's rule:

$$\begin{aligned}\xi &= \frac{\det B}{\det A} \\ &= \frac{((d-1)!(d-2)!\cdots 2!1!)\sum_{i=1}^{d}(-1)^{d+i}N_i\sigma_{d-i}(1,2,\ldots,d)}{(-1)^d d!(d-1)!(d-2)!\cdots 2!1!} \\ &= \frac{1}{d!}\sum_{i=1}^{d}(-1)^i N_i\sigma_{d-i}(1,2,\ldots,d) \\ &= \sum_{i=1}^{d}(-1)^i N_i\sigma_i\left(1,\frac{1}{2},\ldots,\frac{1}{d}\right).\end{aligned}$$

Consequently, we have the desired result. □

One could similarly solve for any particular m_j in this fashion.

Proposition 1.

$$m_j = \binom{d}{j}\frac{1}{j}\sum_{i=1}^{d}(-1)^{j+i}N_i\sigma_{i-1}\left(1,\frac{1}{2},\ldots,\frac{1}{j-1},\frac{1}{j+1},\ldots,\frac{1}{d}\right).$$

Proof. Start with equation (3.4) and apply Cramer's Rule.

Lemma 6. *If we let*

$$B_j = \det\begin{pmatrix} 1^0 & 2^0 & \cdots & (j-1)^0 & 0 & (j+1)^0 & \cdots & d^0 & 1 \\ 1^1 & 2^1 & \cdots & (j-1)^1 & N_1 & (j+1)^1 & \cdots & d^1 & 0 \\ \vdots & \vdots & \cdots & \vdots & \vdots & \vdots & \cdots & \vdots & \vdots \\ 1^d & 2^d & \cdots & (j-1)^d & N_d & (j+1)^d & \cdots & d^d & 0 \end{pmatrix}$$

then $\det B_j$ *is of the form*

$$(d-1)!\cdots 1!\binom{d}{j}\sum_{i=1}^{d}(-1)^{i+j+d}N_i\sigma_{d-i}(1,\ldots,j-1,j+1,\ldots,d).$$

Proof. We start by expanding the determinant along the $(d+1)$th column, arriving at the moderately nicer

$$\begin{aligned}&\det B_j \\ &= (-1)^{d+2}\det\begin{pmatrix} 1^1 & \cdots & (j-1)^1 & N_1 & (j+1)^1 & \cdots & d^1 \\ \vdots & & \vdots & \vdots & \vdots & & \vdots \\ 1^d & \cdots & (j-1)^d & N_d & (j+1)^d & \cdots & d^d \end{pmatrix} \\ &= (-1)^d\frac{d!}{j}\det\begin{pmatrix} 1^0 & \cdots & N_1 & \cdots & d^0 \\ \vdots & & \vdots & & \vdots \\ 1^{d-1} & \cdots & N_d & \cdots & d^{d-1} \end{pmatrix}.\end{aligned}$$

Expanding this new matrix along the jth column results in minors of the form

$$C_{i,j} = \begin{pmatrix} 1^0 & 2^0 & \cdots & (j-1)^0 & (j+1)^0 & \cdots & d^0 \\ \vdots & \vdots & & \vdots & \vdots & & \vdots \\ 1^{i-2} & 2^{i-2} & \cdots & (j-1)^{i-2} & (j+1)^{i-2} & \cdots & d^{i-2} \\ 1^i & 2^i & \cdots & (j-1)^i & (j+1)^i & \cdots & d^i \\ \vdots & \vdots & & \vdots & \vdots & & \vdots \\ 1^{d-1} & 2^{d-1} & \cdots & (j-1)^{d-1} & (j+1)^{d-1} & \cdots & d^{d-1} \end{pmatrix}$$

whence we find $\det B_j$ is

$$\det B_j = (-1)^d \frac{d!}{j} \sum_{i=1}^{d} (-1)^{i+j} N_i \det C_{i,j}.$$

It simplifies our notation if we take $\Gamma = (1, \ldots, j-1, j+1, \ldots, d)$. As before, we have a bialternant of a very similar form. Here we have $\delta = (0, 1, \ldots, d-2)$ and $\lambda = (\underbrace{0, \ldots, 0}_{i-1 \text{ terms}}, \underbrace{1, \ldots, 1}_{d-i \text{ terms}})$, yielding

$$\det C_{i,j} = \sigma_{d-i}(\Gamma)\, a_\delta(\Gamma).$$

This is slightly more complex, as the Vandermonde determinant is no longer a simple product of factorials: in particular,

$$\begin{aligned} a_\delta(\Gamma) &= \prod_{\substack{1 \le u < v \le d \\ u,v \neq j}} (v-u) \\ &= \frac{(d-1)!}{(d-j)} \frac{(d-2)!}{(d-j-1)} \cdots \frac{j!}{1} \frac{(j-1)!}{(j-1)!} (j-2)! \cdots 2!1! \\ &= \frac{(d-1)! \cdots 1!}{(d-j)!(j-1)!}. \end{aligned}$$

Putting it together, we get the desired result. □

Thus, we can solve for $m_j = \frac{\det B_j}{\det A}$, and get:

$$m_j = \binom{d}{j} \frac{1}{d!} \sum_{i=1}^{d} (-1)^{i+j} N_i \sigma_{d-i}(\Gamma).$$

Distributing in the $\frac{1}{d!}$ term into the symmetric polynomial, we get products of i terms, each of the form $\frac{1}{k}$ $(1 \le k \le d)$, each with a $\frac{1}{j}$ term. Removing this common term, we are left with the desired result. □

A MORE DIRECT PROOF

The above demonstrates how one would arrive at this theorem, but more direct proofs are possible. One such proof follows.

An alternate (opaque) proof. Let's begin by introducing a variation of a prior notation. For any $y \in V_f$, define

$$\tilde{N}_{k,y} = \left\{(x_1, \ldots, x_k) \in \mathbb{F}_q^k \mid f(x_1) = \cdots = f(x_k) = y\right\}$$

and denote the corresponding cardinality of these sets as

$$N_{k,y} = \#\left(\tilde{N}_{k,y}\right)$$

and finally, note that

$$N_k = \sum_{y \in V_f} N_{k,y}. \tag{3.5}$$

Let's refer to the right hand side of (3.1) as η; plugging (3.5) into this expression, we get

$$\begin{aligned}
\eta &= \\
&= \sum_{i=1}^{d} (-1)^{i-1} \sum_{y \in V_f} N_{i,y}\, \sigma_i \left(1, \frac{1}{2}, \ldots, \frac{1}{d}\right) \\
&= \sum_{y \in V_f} \sum_{i=1}^{d} (-1)^{i-1} N_{i,y}\, \sigma_i \left(1, \frac{1}{2}, \ldots, \frac{1}{d}\right).
\end{aligned}$$

Let's call the inner sum ω_y, that is:

$$\omega_y = \sum_{i=1}^{d} (-1)^{i-1} N_{i,y}\, \sigma_i \left(1, \frac{1}{2}, \ldots, \frac{1}{d}\right).$$

If we can show that for all $y \in V_f$ we have $\omega_y = 1$, then we clearly have $\eta = \#\left(V_f\right)$.

Let $y \in V_f$ be fixed. Using the same reasoning as above, we know that $y \in V_{f,k}$ for some $1 \leq k \leq d$, and thus $N_{i,y} = k^i$. Substituting this in, our expression mercifully becomes somewhat nicer:

$$
\begin{aligned}
\omega_y &= \sum_{i=1}^{d} (-1)^{i-1} k^i \sigma_i \left(1, \frac{1}{2}, \ldots, \frac{1}{d}\right) \\
&= 1 + \sum_{i=0}^{d} (-1)^{i-1} k^i \sigma_i \left(1, \frac{1}{2}, \ldots, \frac{1}{d}\right) \\
&= 1 - \sum_{i=0}^{d} (-1)^{i} k^i \sigma_i \left(1, \frac{1}{2}, \ldots, \frac{1}{d}\right) \\
&= 1 - \sum_{i=0}^{d} (-1)^{i} \sigma_i \left(k1, k\frac{1}{2}, \ldots, k\frac{1}{d}\right) \qquad (3.6) \\
&= 1 - \left[(1-k1)\left(1-k\frac{1}{2}\right)\cdots\left(1-k\frac{1}{d}\right)\right] \qquad (3.7) \\
&= 1.
\end{aligned}
$$

From step (3.6) to step (3.7), we are using the identity

$$\prod_{j=1}^{n} (\lambda - X_j) = \sum_{j=0}^{n} (-1)^j \lambda^{n-j} \sigma_j (X_1, \ldots, X_n).$$

Note that the bracketed term of (3.7) is 0, as k must be an integer such that $1 \leq k \leq d$, so one term in the product will be 0.

Thus, we have $\eta = \#\left(V_f\right)$, as desired. □

Theorem 3 gives us a way to express $\#\left(V_f\right)$ in terms of the points on a family of curves on $\mathbb{F}_q^k$. If we had a way of getting N_k for $1 \leq k \leq d$, then it would be easy to calculate $\#\left(V_f\right)$.

APPLICATION OF p-ADIC POINT COUNTING

We now proceed by using p-adic point counting of the points in the described spaces. For an introduction to this area, the expository papers [Wan, 2008] and [Lauder, 2005] nicely outline an approach which is fundamentally enabled by [Dwork, 1960].

The spaces $\tilde{N}_k$ aren't of any nice form (we cannot assume they are non-singular projective, abelian varieties, etc.), so we proceed by using the p-adic point counting method described in [Lauder and Wan, 2008], which runs in polynomial time for small characteristic (characters on the order of $p = O((d \log q)^C)$ for some positive constant C, see [Wan, 2008] for details).

Theorem 4. *There exists an explicit deterministic algorithm and an explicit polynomial R such that for any $f \in \mathbb{F}_q[x]$ of degree d, where $q = p^m$ (p prime), the algorithm computes the cardinality of the image set, $\#\left(V_f\right)$, in a number of bit operations bounded by $R\left(m^d d^d p^d\right)$.*

Proof. Recall that $N_k = \#\left(\tilde{N}_k\right)$ with

$$\begin{aligned}\tilde{N}_k &= \left\{(x_1, \ldots, x_k) \in \mathbb{F}_q^k \mid f(x_1) = \cdots = f(x_k)\right\} \\ &= \left\{(x_1, \ldots, x_k) \in \mathbb{F}_q^k \;\middle|\; \begin{array}{ll} f(x_1) - f(x_2) & = 0 \\ f(x_1) - f(x_3) & = 0 \\ & \vdots \\ f(x_1) - f(x_k) & = 0 \end{array}\right\}.\end{aligned}$$

For reasons soon to become clear, we need to represent this as a single polynomial. Let us introduce additional variables z_1 to z_{k-1}, and denote $x = (x_1, \ldots, x_k)$ and $z = (z_1, \ldots, z_{k-1})$. Now examine the auxiliary function

$$F_k(x, z) = z_1 (f(x_1) - f(x_2)) + \cdots + z_{k-1} (f(x_1) - f(x_k)). \quad (3.8)$$

Clearly, if $\gamma \in \tilde{N}_k$, then $F_k(\gamma, z)$ is the zero function. If $\gamma \in \mathbb{F}_q^k \setminus \tilde{N}_k$, then the solutions of $F_k(\gamma, z) = 0$ specify a $(k-2)$-dimensional linear subspace of $\mathbb{F}_q^{k-1}$. Thus, if we denote the cardinality of the solution set to $F_k(x, z) = 0$ as $\#(F_k)$, then we see that

$$\begin{aligned}\#(F_k) &= q^{k-1} N_k + q^{k-2}\left(q^k - N_k\right) \\ &= N_k q^{k-2}(q-1) + q^{2k-2}.\end{aligned}$$

Solving for N_k, we find that

$$N_k = \frac{\#(F_k) - q^{2k-2}}{q^{k-2}(q-1)}. \tag{3.9}$$

thus we have an easy way to determine what N_k is depending on the number of $\mathbb{F}_q$-rational points on this single variety.

The main theorem in [Lauder and Wan, 2008, Theorem 28] yields an algorithm for toric point counting in $\mathbb{F}_{q^\ell}$ which runs in polynomial time for small characteristic (on the order of $p = O((d \log q)^C)$ for some positive constant C; see [Wan, 2008] for details) that works for general varieties. In [Lauder and Wan, 2008, section 6.4], this theorem is adapted to be a generic point counting algorithm.

Adapting this result to our problem, we see that F_k has a total degree of $d+1$, is in $2k-1$ variables, and that we only care about the case where $\ell = 1$. Thus, we have an expected complexity for this algorithm of $\tilde{O}(2^{8k+1} m^{6k+4} k^{6k+2} d^{6k-3} p^{4k+2})$ bit operations. In order to calculate $\#\left(V_f\right)$ using equation (3.1), we calculate N_k for $1 \leq k \leq d$, scaled by an elementary symmetric polynomial. All of the necessary elementary symmetric polynomials can be evaluated using Newton's identity (see [Mead, 1992]) in less than $O(d^2 log d)$ multiplications.

As such, the calculation requires $\tilde{O}(2^{8d+1} m^{6d+4} d^{12d-1} p^{4d+2})$ bit operations. For consistency with [Lauder and Wan, 2008], we can then note that as $d > 1$, we can write $2^{8d+1} = d^{(\log_d 2)(8d+1)}$. Thus, there is a polynomial, R, in one variable such that the runtime of this algorithm is bounded by $R(m^d d^d p^d)$ bit operations. In the dense polynomial model, the polynomial f has input size $O\left(d \log q\right)$, so this algorithm does not have polynomial runtime with respect to the input length in an unrestricted setting. This algorithm has runtime that is exponential in the degree of the polynomial, d, and polynomial in m and p. Thus in the case where d is fixed and where p sufficiently small (in the same sense as in Lauder-Wan), we have a runtime that is polynomial in input length. □

As noted, the initial approach for Theorem 3 included an approach similar to [Birch and Swinnerton-Dyer, 1959]. The difference is that they required that $x_i \neq x_j$ for $i \neq j$. The standard approach to representing such inequalities is the "Rabinovich trick". Using this trick, we introduce an additional variable, say y, and the additional equation

$$y \prod_{i<j} (x_j - x_i) = 1.$$

This is a degree $\binom{k}{2} + 1$ polynomial, which would lead to an equation corresponding to (3.8) of at least degree $\binom{k}{2} + 2$ with $2k + 1$ variables, which would increase the work factor of the algorithm significantly.

4

Proposal

We seek to develop a corresponding set of bounds for incomplete exponential sums based the algebraic structure of the sets we are summing over. This would include Weil sums based on additive and multiplicative characters, but could also include the analogous incomplete Gauss and Jacobi sums. We are also interested in investigating explicit calculation of complete and incomplete exponential sums algorithmically, including analysis of these algorithms given various input models.

Estimating Weil image sums can also be advanced through better understanding of the cardinality of the image set of polynomials. To this end, several additional approaches for estimating this value may yield productive refinements to the theory: investigation of the structure of certain Galois extensions, the algo-geometric structure of certain related varieties, and their underlying combinatorial structure. We seek to develop explicit algorithms that can be used to estimate these sums, and establish the computational complexity of these algorithms under a variety of input models. These results may further refine the computational complexity class of these problems. In particular, Cohen showed that

$$\left|V_f\right| = \mu q + O_d(\sqrt{q}).$$

We seek better methods for estimating and explicitly calculating μ, and a tighter estimate for the constant in the $O_d(\sqrt{q})$ term (the best bound known for this is exponential in d; a bound of the form $d^{O(1)}$ may be possible), together with computational complexity analysis for the resulting algorithms.

In advancing our current approaches, the third naïve approach developed for determining the image set cardinality of polynomials can be modified by clearing the denominators of the rational expressions, resulting in an alternate curve that we can use for point counting. It remains to be seen if this alternate approach produces a computational complexity advantage.

Further, our Theorem 3 can be adopted to provide an estimate for $\#\left(V_f\right)$ based on estimates of the number of $\mathbb{F}_q$-rational points on the constructed varieties. Two immediate candidates for these estimates are [Huang and Wong, 1998] for prime ordered fields, and [Grigoriev and Karpinski, 1991] for arbitrary finite fields.

5 Conclusion

Continued investigation into incomplete Weil sums could have significant impact on several applications; development of analogs to classical results for this style of incomplete exponential sum could result in significant improvements in the understanding of computer science applications, algo-geometric codes, and various areas of cryptography, and may have surprising applications to more general settings.

We have already made some progress in areas by supplying a Weil maximal bound for polynomials of any degree and by finding additional connections between the image set of a polynomial and its algo-geometric structure. Further work would include investigations of additional bounds, refinement of the connection between these problems and other areas of mathematics, and development and analysis of algorithms for accomplishing these tasks.

Bibliography

A.C. Aitken. *Determinants and Matrices*. Oilver and Boyd, ninth edition, 1959.

László Babai, Noam Nisan, and Márió Szegedy. Multiparty protocols, pseudorandom generators for logspace, and time-space trade-offs. *Journal of Computer and System Sciences*, 45(2):204–232, 1992. DOI: 10.1016/0022-0000(92)90047-M .

Eric Bach and Jeffrey Shallit. *Algorithmic Number Theory*, volume 1. The MIT Press, 1997.

P. T. Bateman, S. Chowla, and P. Erdös. Remarks on the size of $L(1, \chi)$. *Publicationes Mathematicae Debrecen*, 1:165–182, 1950.

B. J. Birch and H. P. F. Swinnerton-Dyer. Note on a problem of Chowla. *Polska Akademia Nauk. Instytut Matematyczny. Acta Arithmetica*, 5:417–423, 1959.

E. Bombieri and S. Sperber. On the estimation of certain exponential sums. *Acta Arithmetica*, 69(4):329–358, 1995.

L. Carlitz. Evaluation of some exponential sums over a finite field. *Mathematische Nachrichten*, 96:319–339, 1980. DOI: 10.1002/mana.19800960125 .

L. Carlitz, D. J. Lewis, W. H. Mills, and E. G. Straus. Polynomials over finite fields with minimal value sets. *Mathematika. A Journal of Pure and Applied Mathematics*, 8:121–130, 1961.

Leonard Carlitz. On the number of distinct values of a polynomial with coefficients in a finite field. *Proceedings of the Japan Academy*, 31:119–120, 1955.

Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, 1988. DOI: 10.1137/0217015 .

Wun Seng Chou, Javier Gomez-Calderon, and Gary L. Mullen. Value sets of Dickson polynomials over finite fields. *Journal of Number Theory*, 30(3): 334–344, 1988. DOI: 10.1016/0022-314X(88)90006-6 .

F. R. K. Chung. Diameters and eigenvalues. *Journal of the American Mathematical Society*, 2(2):187–196, 1989. DOI: 10.2307/1990973 .

F. R. K. Chung, R. L. Graham, and R. M. Wilson. Quasi-random graphs. *Combinatorica. An International Journal on Combinatorics and the Theory of Computing*, 9(4):345–362, 1989. DOI: 10.1007/BF02125347 .

Stephen D. Cohen. The distribution of polynomials over finite fields. *Polska Akademia Nauk. Instytut Matematyczny. Acta Arithmetica*, 17:255–271, 1970.

Robert S. Coulter. Explicit evaluations of some Weil sums. *Acta Arithmetica*, 83(3):241–251, 1998a.

Robert S. Coulter. Further evaluations of Weil sums. *Acta Arithmetica*, 86(3): 217–226, 1998b.

Thomas W. Cusick. Value sets of some polynomials over finite fields GF(2^{2m}). *SIAM Journal on Computing*, 27(1):120–131 (electronic), 1998. DOI: 10.1137/S0097539794270352 .

Thomas W. Cusick. Polynomials over base 2 finite fields with evenly distributed values. *Finite Fields and their Applications*, 11(2):278–291, 2005. DOI: 10.1016/j.ffa.2004.10.001 .

Thomas W. Cusick and Peter Müller. Wan's bound for value sets of polynomials. In *Finite fields and applications (Glasgow, 1995)*, volume 233 of *London Math. Soc. Lecture Note Ser.*, pages 69–72. Cambridge Univ. Press, Cambridge, 1996. DOI: 10.1017/CBO9780511525988.008 .

Pinaki Das. The number of polynomials of a given degree over a finite field with value sets of a given cardinality. *Finite Fields and their Applications*, 9 (2):168–174, 2003. DOI: 10.1016/S1071-5797(02)00020-5 .

H. Davenport and P. Erdös. The distribution of quadratic and higher residues. *Publicationes Mathematicae Debrecen*, 2:252–265, 1952.

Harold Davenport. *Multiplicative Number Theory*. Springer–Verlag, 3 edition, 2000.

Leonard Eugene Dickson. The analytic representation of substitutions on a power of a prime number of letters with a discussion of the linear group. *Annals of Mathematics*, 11(1-6):65–120, 1896/97. DOI: 10.2307/1967217 .

Bernard Dwork. On the rationality of the zeta function of an algebraic variety. *American Journal of Mathematics*, 82:631–648, 1960.

Carl Friedrich Gauss. *Disquisitiones Arithmeticae*. Gerhard Fleischer, Jr., 1801.

Javier Gomez-Calderon. *Polynomials with small value set over finite fields*. PhD thesis, The University of Arizona, 1986.

Javier Gomez-Calderon. A note on polynomials with minimal value set over finite fields. *Mathematika. A Journal of Pure and Applied Mathematics*, 35(1): 144–148, 1988. DOI: 10.1112/S0025579300006355 .

Javier Gomez-Calderon and Daniel J. Madden. Polynomials with small value set over finite fields. *Journal of Number Theory*, 28(2):167–188, 1988. DOI: 10.1016/0022-314X(88)90064-9 .

Dima Grigoriev and Marek Karpinski. An approximation algorithm for the number of zeros of arbitrary polynomials over GF[q]. In *32nd Annual Symposium on Foundations of Computer Science (San Juan, PR, 1991)*, pages 662–669. IEEE Comput. Soc. Press, Los Alamitos, CA, 1991. DOI: 10.1109/SFCS.1991.185433 .

Alexander Grothendieck. Formule de Lefschetz et rationalité des fonctions L. In *Séminaire Bourbaki, Vol. 9*, pages Exp. No. 279, 41–55. Soc. Math. France, Paris, 1964.

G.H. Hardy and E.M. Wright. *An Introduction to the Theory of Numbers*. Oxford University Press, 5 edition, 2000.

David R. Hayes. A geometric approach to permutation polynomials over a finite field. *Duke Mathematical Journal*, 34:293–305, 1967.

Ming-Deh Huang and Yiu-Chung Wong. An algorithm for approximate counting of points on algebraic sets over finite fields. In *Algorithmic number theory (Portland, OR, 1998)*, volume 1423 of *Lecture Notes in Comput. Sci.*, pages 514–527. Springer, Berlin, 1998. DOI: 10.1007/BFb0054889 .

Richard Kantor. Über die Anzahl inkongruenter Werte ganzer, rationaler Funktionen. *Monatshefte für Mathematik und Physik*, 26(1):24–39, 1915. DOI: 10.1007/BF01999438 .

Donald Knuth. *The Art of Computer Programming*, volume 2. Addison-Wesley, third edition, 1998.

N. M. Korobov. New number-theoretic estimates. *Doklady Akademii Nauk SSSR*, 119:433–434, 1958.

Alan G. B. Lauder. Rigid cohomology and p-adic point counting. *Journal de Théorie des Nombres de Bordeaux*, 17(1):169–180, 2005.

Alan G. B. Lauder and Daqing Wan. Counting points on varieties over finite fields of small characteristic. In J.P. Buhler and P. Stevenhagen, editors, *Algorithmic Number Theory*, pages 579 – 612. Cambridge University Press, 2008.

Rudolf Lidl and Harald Niederreiter. *Finite Fields*. Cambridge University Press, 1997.

Keju Ma and Joachim von zur Gathen. The computational complexity of recognizing permutation functions. *Computational Complexity*, 5(1):76–97, 1995a. DOI: 10.1007/BF01277957 .

Keju Ma and Joachim von zur Gathen. Tests for permutation functions. *Finite Fields and their Applications*, 1(1):31–56, 1995b. DOI: 10.1006/ffta.1995.1003 .

D. G. Mead. Newton's identities. *The American Mathematical Monthly*, 99(8): pp. 749–751, 1992.

W. H. Mills. Polynomials with minimal value sets. *Pacific Journal of Mathematics*, 14:225–241, 1964.

Thomas Muir and William H. Metzler. *A Treatise on the Theory of Determinants*. Dover Publications, Inc., 1960.

Gary L. Mullen. Permutation polynomials over finite fields. In *Finite fields, coding theory, and advances in communications and computing (Las Vegas, NV, 1991)*, volume 141 of *Lecture Notes in Pure and Appl. Math.*, pages 131–151. Dekker, New York, 1993.

Michael O. Rabin. Probabilistic algorithms in finite fields. *SIAM Journal on Computing*, 9(2):273–280, 1980. DOI: 10.1137/0209024 .

I. E. Shparlinski. A deterministic test for permutation polynomials. *Computational Complexity*, 2(2):129–132, 1992. DOI: 10.1007/BF01202000 .

Richard P. Stanley. *Enumerative Combinatorics*, volume 2. Cambridge University Press, 2001.

Saburô Uchiyama. Sur le nombre des valeurs distinctes d'un polynôme à coefficients dans un corps fini. *Proceedings of the Japan Academy*, 30:930–933, 1954.

Saburô Uchiyama. Note on the mean value of $V(f)$. *Proceedings of the Japan Academy*, 31:199–201, 1955.

J. F. Voloch. On the number of values taken by a polynomial over a finite field. *Polska Akademia Nauk. Instytut Matematyczny. Acta Arithmetica*, 52 (2):197–201, 1989.

Joachim von zur Gathen. Values of polynomials over finite fields. *Bulletin of the Australian Mathematical Society*, 43(1):141–146, 1991a. DOI: 10.1017/S0004972700028860 .

Joachim von zur Gathen. Tests for permutation polynomials. *SIAM Journal on Computing*, 20(3):591–602, 1991b. DOI: 10.1137/0220037 .

Daqing Wan. A p-adic lifting lemma and its applications to permutation polynomials. In *Finite fields, coding theory, and advances in communications and computing (Las Vegas, NV, 1991)*, volume 141 of *Lecture Notes in Pure and Appl. Math.*, pages 209–216. Dekker, New York, 1993.

Daqing Wan. Computing zeta functions over finite fields. In *Finite fields: theory, applications, and algorithms (Waterloo, ON, 1997)*, volume 225 of *Contemp. Math.*, pages 131–141. Amer. Math. Soc., Providence, RI, 1999.

Daqing Wan. Algorithmic theory of zeta functions over finite fields. In *Algorithmic number theory: lattices, number fields, curves and cryptography*, volume 44 of *Math. Sci. Res. Inst. Publ.*, pages 551–578. Cambridge Univ. Press, Cambridge, 2008.

André Weil. On some exponential sums. *Proceedings of the National Academy of Sciences of the United States of America*, 34:204–207, 1948.

Hermann Weyl. Über ein Problem aus dem Gebiet der Diophantischen Approximationen. *Nachrichten von der Königlichen Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-physikalische Klasse*, pages 234–245, 1914.

Hermann Weyl. Über die Gleichverteilung von Zahlen mod. Eins. *Mathematische Annalen*, pages 313–352, 1916.

David Zuckerman. General weak random sources. In *31st Annual Symposium on Foundations of Computer Science, Vol. I, II (St. Louis, MO, 1990)*, pages 534–543. IEEE Computer Society Press, Los Alamitos, CA, 1990. DOI: 10.1109/FSCS.1990.89574 .

Colophon

The text of this document is typeset in Jean-François Porchez's wonderful Sabon Next typeface (in the regular, display, and bold weights). Sabon Next is a modern (2002) revival of Jan Tschichold's 1967 Sabon typeface, which is in turn an adaptation of Claude Garamond's mid-16th century typeface.

Equations are typeset using the MathTime Professional II (MTPro2) fonts, a font package released in 2006 by the great mathematical expositor Michael Spivak. These fonts are designed to work with the Times typeface, but they blend well with many classical fonts.

Portions of the bibliography are set in Evert Bloemsma's 2004 humanist san-serif font Legato. Chapter numbers are set in Hermann Zapf's delightful Zapfino.

The very configurable Memoir class was used with XeTeX to typeset the document. XeTeX is a scion of Donald Knuth's profoundly important TeX.

Diagrams were produced in Mathematica.

The copies of this paper that I distributed were printed on $8.5'' \times 11''$ 28lb ($105 g/m^2$) smooth white Mohawk Superfine paper. The cover page and divider were printed on $8.5'' \times 11''$ 80lb ($216 g/m^2$) smooth white Mohawk Superfine cover stock.