Weil Image Sums
(and some related problems)

Joshua E. Hill

Department of Mathematics
University of California, Irvine

Advancement to Candidacy Exam
2011-Sep-26 (rev. 1.04)
http://bit.ly/AdvJEH

Talk Outline

Introduction Outline

General Exponential Sums: Weyl Sums

Definition

A Weyl Sum is any sum of the form

$$T = \sum_{j=1}^{N} e^{2\pi i P(j)}$$

where $P(x)$ is a polynomial over the real numbers.

First approximations for bounds:

- Trivially: $|T| \leq N$ (worst case)
- If P produced random outputs, then we would expect this to look like a 2-dimensional random walk: $|T| = O(\sqrt{N})$
- Generally, there is some structure and we are stuck with $|T| = o(N)$

Applications

Exponential sums are a reoccurring tool

- Number Theory
  - Sums of Squares
  - Class field theory
- Discrete Fourier Transform
  - Implemented by some style of FFT: "If you speed up any nontrivial algorithm by a factor of a million or so, the world will beat a path toward finding useful applications for it." – *Numerical Recipes* §13.0
- Paley graphs
- Computer Science
  - Graph theoretic applications
  - Random number generators

Characters

Definition

A character is a monoid homomorphism from a monoid G to the units of a field K^*.

- We will be principally working with finite fields, and our co-domain is $\mathbb{C}^*$.
- Fields have two obvious group structures we can use:
  - Additive
  - Multiplicative
- For this discussion, we are mainly concerned with additive characters.

Additive Characters

We can represent all additive characters of the form $\mathbb{F}_q \to \mathbb{C}^*$ nicely.

Definition

Let $\mathbb{F}_q$ be a finite field of $q = p^m$ elements (where p is prime). The (absolute) trace of $\alpha \in \mathbb{F}_q$ is $\mathrm{Tr}(\alpha) = \sum_{j=0}^{m-1} \alpha^{p^j}$.

Theorem (Weber 1882)

All additive characters of this type are of the form $\psi_\gamma(\alpha) = e^{\frac{2\pi i}{p}\mathrm{Tr}(\gamma\alpha)}$ *for some* $\gamma \in \mathbb{F}_q$.

Weil Sums

Definition

A Weil Sum is any sum of the form

$$W_{f,\gamma} = \sum_{c \in \mathbb{F}_q} \psi_\gamma(f(c))$$

where $f(x)$ is a polynomial over $\mathbb{F}_q$ and ψ_γ is an additive character.

Weil determined bounds:

Theorem (Weil 1948)

If $f(x) \in \mathbb{F}_q[x]$ is of degree $d > 1$ with $p \nmid d$ and ψ_γ is a non-trivial additive character of $\mathbb{F}_q$, then $\left|W_{f,\gamma}\right| \leq (d-1)\sqrt{q}$.

A Quick Aside

Hermann Weyl
(1885-1955)

$\neq$

André Weil
(1906-1998)

UNIVERSITY of CALIFORNIA · IRVINE

Weil Image Sums

- We adopt the notation $V_f = f(\mathbb{F}_q)$
- We examine incomplete Weil sums on the image set

$$S_{f,\gamma} = \sum_{\alpha \in V_f} \psi_\gamma(\alpha)$$

- To remove the dependence on the choice of character, we look at the maximal such sum (over non-trivial additive characters)

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

University *of* California · Irvine

Weil Image Sum Example

Example

- In $\mathbb{F}_4$, we'll represent field elements as polynomials over $\mathbb{F}_2[t]$ mod the irreducible $t^2 + t + 1$.
- Examine $f(x) = x^3 + x$:

α	$f(\alpha)$	$\mathrm{Tr}(f(\alpha))$	$\mathrm{Tr}(tf(\alpha))$	$\mathrm{Tr}((t+1)f(\alpha))$
0	0	0	0	0
1	0	0	0	0
t	$t+1$	1	0	1
$t+1$	t	1	1	0

- $W_{f,1} = e^{\pi i 0} + e^{\pi i 0} + e^{\pi i 1} + e^{\pi i 1} = 0$
- $\#\left(V_f\right) = 3$
- $S_{f,1} = e^{\pi i 0} + e^{\pi i 1} + e^{\pi i 1} = -1$
- $\left|S_f\right| = 1$ (this is maximal)

Conjecture

Conjecture (Wan)

For all polynomials of degree d, with $p \nmid d$:

1. *There is a real number c_d such that $\left|S_f\right| \leq c_d \sqrt{q}$ for all q*
2. $c_d \leq c\sqrt{d}$
3. $c \leq 1$

Some notes about conjecture (1):

- (1) is true when $q \gg d$ as a consequence of Cohen / Chebotarev / Lenstra-Wan (unpublished).
- If $d = o(q)$, then (1) isn't very interesting.

What is Success?

Better information about $\left|S_f\right|$ or $\#\left(V_f\right)$:

- Better bounds
- An algorithm for computing or estimating
- Results that significantly refine the complexity class of these problems

Literature Survey Outline

Subsection 1

Exponential Sums

Gauss Sums

- Gauss sums were initially studied by...Gauss...
  Appeared in *Disquisitiones Arithmeticae*.
- If ψ is an additive character and χ is a multiplicative character, then a Gauss sum is as sum of the form

$$G(\psi, \chi) = \sum_{\alpha \in \mathbb{F}_q^*} \psi(\alpha)\chi(\alpha)$$

- This is a finite-field analog to the Γ function.
- This sum is used extensively in number theory
- Weil Image Sums are a variation of these sums (under the appropriate definitions of $\chi(0)$)

Weil Sums

Certain polynomial forms have special bounds for the associated Weil sum:

- $x^n + b$
- p-linear polynomials
- quadratics

Certain polynomials have explicit solutions for the associated Weil sum:

- $ax^{p^\alpha+1} + bx$ [Carlitz 1980 for $\alpha = 1$, Coulter 1998]

Some work for incomplete sums over alternate structures:

- Summed over quasi-projective varieties [Bombieri-Sperber, 1995]

Subsection 2

Cardinality of Image Sets

Cardinality of Image Sets

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

- These bounds are sharp!
- If $\#\left(V_f\right) = \left\lceil \frac{q}{d} \right\rceil$, then f is a polynomial with a minimal value set.
- If $\#\left(V_f\right) = q$, then f is a permutation polynomial.

University *of* California · Irvine

The Shape of the Problem (Average Results)

A vital companion function:

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

- If $f^*(u, v)$ is absolutely irreducible then on average $\#\left(V_f\right) \sim \mu_d q + O_d(1)$ with μ_d is the series $1 - e^{-1}$ truncated at d terms. [Uchiyama 1955]

Asymptotic Results I

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

First asymptotic results [Birch and Swinnerton-Dyer, 1959]

- μ is dependent on some Galois groups induced by f

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

  where $G^+(f)$ is viewed as a subgroup of $G(f)$.
- If $G^+(f) \cong S_d$ (f is a "general polynomial") then $\mu = \mu_d$.
- Otherwise μ depends only on $G(f)$, $G^+(f)$ and d.

Cohen gave a way to explicitly calculate μ [Cohen, 1970]

- Let K be the splitting field for $f(x) - t$ over $\mathbb{F}_q(t)$
- Denote $k' = K \cap \bar{\mathbb{F}}_q$
- $G^*(f) = \{\sigma \in G(f) \mid K_\sigma \cap k' = \mathbb{F}_q\}$
- $G_1(f) = \{\sigma \in G(f) \mid \sigma \text{ fixes at least one point}\}$
- $G_1^*(f) = G_1(f) \cap G^*(f)$
- We then have $\mu = \frac{\#(G_1^*)}{\#(G^*)}$.
- This provides a wonderful combinatorial explanation of μ_d (proportion of non-derangements!)

Exact Results

Exact values for $\#\left(V_f\right)$ are known for very few classes of polynomials:

- Permutation polynomials (and exceptional polynomials)
- Polynomials with a minimal (or very small) value set
- Other

Permutation Polynomials

The class of polynomials where $\#\left(V_f\right) = q$

1. These polynomials are uncommon ($\sim e^{-q}$ for large q)
2. Dickson found all of the permutation polynomials $d \leq 6$ [Dickson 1896]
3. There is a ZPP algorithm to test to see if f is a permutation polynomial. [Ma and von zur Gathen, 1995]
4. There is a deterministic algorithm to see if f is a permutation polynomial that runs slightly sub-linear in q. [Shparlinski, 1992]

Exceptional Polynomials

Hayes harmonized these apparently disparate results by casting this into an Algo-Geometric setting [Hayes 1967]

Definition

$f(X) \in \mathbb{F}_q[X]$ is an exceptional polynomial if when $f^*(X, Y)$ is factored into irreducibles over $\mathbb{F}_q[X, Y]$ and all of these irreducible factors are not absolutely irreducible (that is, each irreducible factor cannot be irreducible over $\bar{\mathbb{F}}_q[X, Y]$.)

- All exceptional polynomials are permutation polynomials [Cohen 1970], [Wan, 1993]
- If $d > 1$, $p \nmid d$ and $q > d^4$, then all permutation polynomials are exceptional polynomials. (by Lang-Weil Bound)
- f is an exceptional polynomial if and only if $\mu = 1$.

Small Image Set Polynomials

- All polynomials with minimal value sets with $d \leq \sqrt{q}$ were characterized in [Carlitz, Lewis, Mills, Straus 1961/1964]
- All polynomials with $d^4 < q$ with $\#\left(V_f\right) < 2q/d$ were characterized in [Gomez-Calderon, 1986]

UNIVERSITY of CALIFORNIA · IRVINE

$\#\left(V_f\right)$ is known in a few other cases:

- Degree 0 and 1 cases are clear
- Degree 2,3 cases are due to [Kantor 1915] and [Uchiyama 1955]
- p-linear polynomials are known due to linearity
- Dickson Polynomials [Chou Gomez-Calderon, Mullen 1988]
- $f(x) = x^k(1+x)^{2^m-1}$ in $\mathbb{F}_{2^m}$ (for $k = \pm 1, \pm 2, 4$) and $f(x) = (x+1)^d + x^d + 1$ for particular values of d [Cusick 2005]

An Important Note

- These results may seem to suggest that V_f can only be of certain forms. This is completely false.
- Lagrange interpolation can be used to build a polynomial with any image set.
- The restrictions discussed tell us that some of these image sets cannot be associated with polynomials of certain degrees.
- Note that we can always reduce mod $X^q - X$ and get the same image set.

Subsection 3

p-adic Point Counting

The Zeta Function on Algebraic Sets

Consider the simultaneous zeros of a set of polynomials $f_1, \ldots, f_s \in \mathbb{F}_q[x_1, \ldots, x_n]$ over $\bar{\mathbb{F}}_q$; call this variety X.

- Let $X(\mathbb{F}_{q^k}) = X \cap \mathbb{F}_{q^k}$.

Definition

The zeta function of the algebraic set X is defined to be

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

Curiouser and Curiouser

- Weil conjectured that the zeta function is rational.
- This conjecture was first proven by Dwork in 1960 using p-adic methods.
- This conjecture was again proven by Grothendieck in 1964 using ℓ-adic cohomological methods.
- If it's rational, then intuitively there is only a fixed amount of information necessary to fully establish $Z(X)$. This is fundamentally what enables the p-adic approach to calculating $Z(X)$.
- Approaches to building up $Z(X)$ generally start by calculating $X(\mathbb{F}_{q^k})$ up to a suitably large k.
- We only care about the number of points in $\mathbb{F}_q$, so we only need to look at $X(\mathbb{F}_q)$.

UNIVERSITY *of* CALIFORNIA · IRVINE

The point counting algorithm of Lauder and Wan [Lauder-Wan 2008]:

Lemma

If f has total degree d in n variables and $p = O((d \log q)^C)$ for some constant C, then $\#\left(X(\mathbb{F}_{q^k})\right)$ can be calculated in polynomial time (polynomial in p, m, k, and d; exponential in n).

Preliminary Results Outline

Attribution

All of these results are taken from joint work with Daqing Wan.

Subsection 1

Weil Image Sum Bounds

Too Many Polynomials on the Dance Floor I

- Start with an arbitrary degree d polynomial $f(x) = a_d x^d + \cdots + a_0$, $a_i \in \mathbb{F}_q$.
- $f(x)$ and $f(x - \lambda)$ have the same image set.
  - Setting $\lambda = \frac{a_{d-1}}{d a_d}$ removes x^{d-1} term.
  - Thus, WLOG we can examine $f(x) = a_d x^d + a_{d-2} x^{d-2} + \cdots + a_0$.
- We can do better: $f(x) = x^d + a_{d-2} x^{d-2} + \cdots + a_1 x$.

Too Many Polynomials on the Dance Floor II

Let I_f be some minimal preimage set that produces V_f.

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

Bounding $\left|S_f\right|$

We introduce two expressions to help us discuss bounds:

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

- Examining Φ_d gives us insight into the value c_d: For all q, $c_d \geq \Phi_d$.
- A related question: for a given q, what is the maximum $\left|S_f\right|$ possible?

$$\left|S_{A_q}\right| = \max_{A \subset \mathbb{F}_q} \left| \sum_{\alpha \in A} \psi_1(\alpha) \right|$$

University *of* California · Irvine

A Word of Warning

- At least one polynomial produces A_q as an image set.
- This polynomial does not necessarily have degree relatively prime to p.
- Not every image set can be obtained as the image of a polynomial whose degree is relatively prime to p.

An Example of Warning

Example

- In $\mathbb{F}_4$ again.
- Examine $f(x) = x^2 + x$ (p-linear!):

α	$f(\alpha)$
0	0
1	0
t	1
$t+1$	1

- Clearly no polynomial with degree 0 or 1 will have this image.
- Idea: We don't expect that degree 3 polynomials would be linear.
- Actual Proof: Just evaluate all degree 3 polynomials in $\mathbb{F}_4[x]$ and note that none of them have this image.

Bounding Theorem Proof Outline I

Theorem

If $q = p^m$ then

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

The "interesting part" of the proof:

- Trace is an $\mathbb{F}_p$-linear transform, and surjects onto $\mathbb{F}_p$.
- $\#\left(\ker \mathrm{Tr}\right) = p^{m-1}$
- Thus each element is hit p^{m-1} times.
- To find A_q, find A_p and then choose all the elements in the same equivalence classes.

This reduces the question to the case where $q = p$. The rest is "proof by calculus".

University *of* California · Irvine

Bounding Theorem Proof Outline II

- We are now summing distinct pth roots of unity, seeking the largest modulus possible.
- A proposed maximal sum must include all the roots of unity with angle $\leq \pi/2$ to the sum.
- $p = 2$ case is trivial. Assume p is odd.
- First stab: All of the pth roots of unity in quadrants I and IV?

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

- This is maximal, but obviously not unique.

UNIVERSITY *of* CALIFORNIA · IRVINE

Corollary

As $p \to \infty$ along the primes, $\left|S_{A_q}\right| \searrow \frac{q}{\pi}$

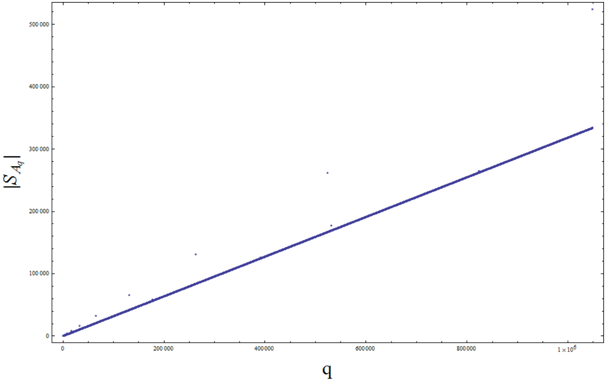

Estimation of c_d

- Estimate Φ_d by looking at all polynomials of that degree for various q or random sampling.
- We did both with fields up to 100 elements.
- The maximal value for each degree is plotted below.

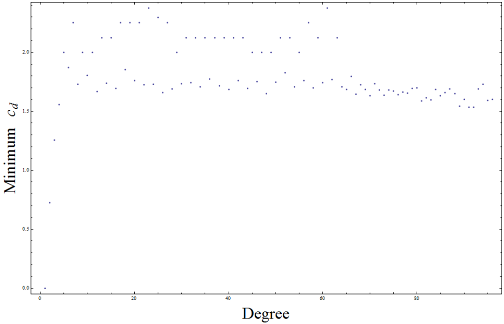

Estimation of c

- The same data, but additionally normalized by a factor of $\sqrt{d}$.

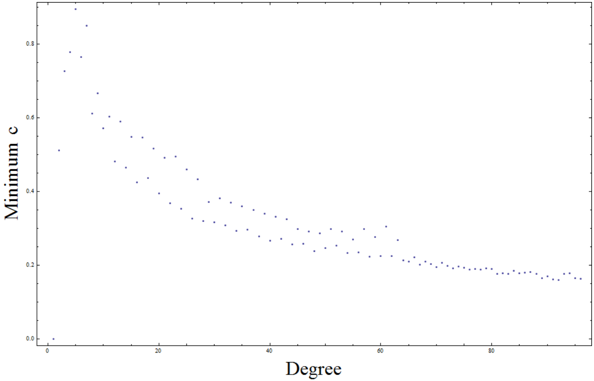

Subsection 2

Image Set Cardinality

Big-O and Soft-O Notation

- We have two eventually positive real valued functions $A, B : \mathbb{N}^k \to \mathbb{R}^+$. Take $\mathbf{x}$ as an n-tuple, with $\mathbf{x} = (x_1, \ldots, x_n)$
- We'll write $|\mathbf{x}|_{\min} = \min_i x_i$.

Definition

1. $A(\mathbf{x}) = O(B(\mathbf{x}))$ if there exists a positive real constant C and an integer N so that if $|\mathbf{x}|_{\min} > N$ then $A(\mathbf{x}) \leq CB(\mathbf{x})$.
2. $A(\mathbf{x}) = \tilde{O}(B(\mathbf{x}))$ if there exists a positive real constant C' so that $A(\mathbf{x}) = O(B(\mathbf{x}) \log^{C'}(B(\mathbf{x}) + 3))$

University *of* California · Irvine

How to calculate $\#\left(V_f\right)$?

- Evaluate f at each point in $\mathbb{F}_q$. Cost: $\tilde{O}(qd)$ bit operations.
- For each $a \in \mathbb{F}_q$, $a \in V_f \Leftrightarrow \deg\gcd(f(x) - a, X^q - X) > 0$. Cost: $\tilde{O}(qd)$ bit operations.

Another connection between $\#\left(V_f\right)$ and an algo-geometric structure:

Theorem

If $f \in \mathbb{F}_q[x]$ of positive degree d, then

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

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 is the ith elementary symmetric function on d elements.

Proof Outline I

- $V_{f,i} = \{x \in V_f \mid \#\left(f^{-1}(x)\right) = i\}$ with $1 \leq i \leq d$ forms a partition of V_f.
- Let $m_i = \#\left(V_{f,i}\right)$. Thus $m_1 + \cdots + m_d = \#\left(V_f\right)$. Introduce a new value $\xi = -\#\left(V_f\right)$. We then have:

$$m_1 + \cdots + m_d + \xi = 0 \tag{1}$$

- Define the space $\tilde{N}_k = \left\{(x_1, \ldots, x_k) \in \mathbb{F}_q^k \mid f(x_1) = \cdots = f(x_k)\right\}$. Then $N_k = \#\left(\tilde{N}_k\right)$.
- By a counting argument,

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

Proof Outline II

Arrange this into a system of equations:

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

Solve for ξ using Cramer's rule. There are some unfortunate details. See the paper. :-)

You can just as reasonably solve for m_j through the same process:

Proposition

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

Application of Lauder-Wan

- This equation is in terms of N_k, which we must establish.
- $\tilde{N}_k$ isn't of any particularly desirable form: in particular, we can't assume that it is non-singular projective or an abelian variety (if it were, faster algorithms would apply!)
- We'll proceed through trickery.

Algorithm for finding $\#\left(V_f\right)$

Theorem

There is a an explicit polynomial R and a deterministic algorithm which, for any $f \in \mathbb{F}_q[x]$ (with $q = p^m$, p a prime, f degree d), calculates $\#\left(V_f\right)$. This algorithm requires a number of bit operations bounded by $R(m^d d^d p^d)$.

More explicit performance: $\tilde{O}\left(2^{8d+1} m^{6d+4} d^{12d-1} p^{4d+2}\right)$ bit operations.

Proof Outline

Define:

$$F_k(\mathbf{x}, \mathbf{z}) = z_1 \left(f(x_1) - f(x_2)\right) + \cdots + z_{k-1} \left(f(x_1) - f(x_k)\right)$$

- If $\gamma \in \tilde{N}_k$ then $F_k(\gamma, \mathbf{z}) = 0$.
- If $\gamma \in \mathbb{F}_q^k \setminus \tilde{N}_k$ then the solutions to $F_k(\gamma, \mathbf{z})$ form a $(k-2)$-dimensional subspace of $\mathbb{F}_q^{k-1}$.
- If we denote the number of solutions to $F_k(\mathbf{x}, \mathbf{z})$ as $\#(F_k)$, then we have
$$\#(F_k) = q^{k-1} N_k + q^{k-2}(q^k - N_k)$$
- So, we can solve:
$$N_k = \frac{\#(F_k) - q^{2k-2}}{q^{k-2}(q-1)}$$
- And that's it!

Proposal Outline

Proposal: Goals

- A first step at understanding this style of sum is understanding V_f.
  - Calculating V_f.
  - Estimating V_f.
  - Refining bounds for or estimating μ.
  - Refining the constant associated with the $O_d(\sqrt{q})$ term; current term is highly exponential in d; $d^{O(1)}$ may be possible.
- We seek to investigate incomplete exponential sums evaluated on image sets.
  - Work thus far has been with additive characters and Weil sums.
  - Many of the same approaches would work with Weil sums of multiplicative characters.
  - Other sum styles can also be investigated: incomplete Gauss and Jacobi sums may also yield results.

Proposal: Style of Results

We look for results of the following styles:

- Improved explicit bounds.
- Algorithms for explicitly calculating values.
- Algorithms for producing estimates.
- Refinements to the complexity class of these problems.

Conclusion Outline

Conclusion

- We outlined problems in finite fields concerning:
  - incomplete Weil exponentials sums (Weil Image Sums)
  - the image set of a polynomial
- We surveyed literature relevant to these problems.
- We discussed new findings related to these problems.
- Proposed a course of investigation for further work.

Colophon

- The principal font is Evert Bloemsma's 2004 humanist san-serif font Legato. This font is designed to be exquisitely readable, and is a significant departure from the highly geometric forms that dominate most san-serif fonts. Legato was Evert Bloemsma's final font prior to his untimely death at the age of 46.
- Equations are typeset using the MathTime Professional II (MTPro2) fonts, a font package released in 2006 by the great mathematical expositor Michael Spivak.
- The serif text font (which appears mainly as text within mathematical expressions) is Jean-François Porchez's wonderful 2002 Sabon Next typeface.
- The URLs are typeset in Luc(as) de Groot's 2005 Consolas, a monospace font with excellent readability.
- Diagrams were produced in Mathematica.