The Fourier Transform and Equations over Finite Abelian Groups

An introduction to the method of trigonometric sums

LECTURE NOTES BY
László Babai
Department of Computer Science
University of Chicago

December 1989
Updated June 2002

VERSION 1.3

The aim of these notes is to present an easily accessible introduction to a powerful method of number theory.

The punchline will be the following finite counterpart of Fermat's Last Theorem:

Theorem 0.1 *If k is an integer, q a prime power, and $q \geq k^4 + 4$, then the Fermat equation*

$$x^k + y^k = z^k \tag{1}$$

has a nontrivial solution in the finite field $\mathbb{F}_q$ of order q.

This result seems to belong to algebraic geometry over finite fields: we have an algebraic variety and we assert that it has points over $\mathbb{F}_q$ other than certain "trivial" ones. In fact, we can asymptotically estimate the number of solutions if q/k^4 is large.

As we shall see, algebraic equations have little to do with the method. Indeed, a much more general result will follow easily from the basic theory. Let $\mathbb{F}_q^\times = \mathbb{F}_q \backslash \{0\}$.

Theorem 0.2 *Let k be an integer, $A_1, A_2 \subseteq \mathbb{F}_q, l_i = (q-1)/|A_i|$ (not necessarily integers), and assume that*

$$q \geq k^2 l_1 l_2 + 4. \tag{2}$$

Then the equation

$$x + y = z^k \quad (x \in A_1, y \in A_2, z \in \mathbb{F}_q^\times) \tag{3}$$

has at least one solution.

Theorem 0.1 follows from this result if we set $A_1 = A_2 = \{a^k : a \in \mathbb{F}_q^\times$. Clearly, $|A_i| = \frac{q-1}{\text{g.c.d.}(k,q-1)} \geq (q-1)/k$ and therefore $l_i \leq k$ in this case.

Note that in Theorem 0.2, the sets A_1 and A_2 are arbitrary (as long as they are not too small compared to q). This result has a flavor of combinatorics where the task often is to create order out of nothing (i.e., without prior structural assumptions). Results like this one have wide applicability in combinatorial terrain such as combinatorial number theory (to which they belong) and even in the theory of computing.

Notation

$\mathbb{C}$: field of complex numbers

$\mathbb{C}^\times = \mathbb{C}\backslash\{0\}$: multiplicative group of complex numbers

$\mathbb{Z}$: ring of integers

$\mathbb{Z}_n = \mathbb{Z}/n\mathbb{Z}$: ring of mod n residue classes

$\mathbb{F}_q$: field of q elements where q is a prime power

$(\mathbb{F}_q, +)$: the additive group of $\mathbb{F}_q$

$\mathbb{F}_q^\times = \mathbb{F}_q\backslash\{0\}$: the multiplicative group of $\mathbb{F}_q$.

1 Characters

Let G be a finite abelian group of order n, written additively.

A *character* of G is a homomorphism $\chi : G \to \mathbb{C}^\times$ of G to the multiplicative group of (nonzero) complex numbers:

$$\chi(a+b) = \chi(a)\chi(b) \quad (a, b \in G). \tag{4}$$

Clearly,

$$\chi(a)^n = \chi(na) = \chi(0) = 1 \quad (a \in G), \tag{5}$$

so the values of χ are n^{th} roots of unity. In particular,

$$\chi(-a) = \chi(a)^{-1} = \overline{\chi(a)} \tag{6}$$

where the bar indicates complex conjugation.

The *principal character* is defined by

$$\chi_0(a) = 1 \quad (a \in G). \tag{7}$$

Proposition 1.1 *For any nonprincipal character χ of G,*

$$\sum_{a \in G} \chi(a) = 0. \tag{8}$$

Proof: Let $b \in G$ be such that $\chi(b) \neq 1$, and let S denote the sum on the left hand side of equation (8). Then

$$\chi(b) \cdot S = \sum_{a \in G} \chi(b)\chi(a) = \sum_{a \in G} \chi(b+a) = S$$

hence

$$S(\chi(b) - 1) = 0,$$

proving the claim. □

Corollary 1.2 (First orthogonality relation for characters) *Let χ and ψ be two characters of G. Then*

$$\sum_{a\in G} \overline{\chi(a)}\psi(a) = \begin{cases} n & \text{if } \chi = \psi \\ 0 & \text{otherwise.} \end{cases} \tag{9}$$

Proof: The case $\chi = \psi$ follows from equation (6). If $\chi \neq \psi$, then $\overline{\chi}\psi$ is a nonprincipal character, hence Proposition 1.1 applies. □

As observed in the last proof, the pointwise product of the characters χ and ψ is a character again:

$$(\chi\psi)(a) := \chi(a)\psi(a) \tag{10}$$

Let $\widehat{G}$ denote the set of characters. It is easy to see that this set forms an abelian group under operation (10). $\widehat{G}$ is called the *dual group* of G.

Proposition 1.3 *Let ω be a primitive n^{th} root of unity. Then the map $\chi_j : \mathbb{Z}_n \to \mathbb{C}^\times$ defined by*

$$\chi_j(a) := \omega^{ja} \tag{11}$$

is a character of $\mathbb{Z}_n$ for every $j \in \mathbb{Z}$. Moreover,

(a) *$\chi_j = \chi_k$ if and only if $j \equiv k \bmod n$;*
(b) *$\chi_j = \chi_1^j$;*
(c) *$\widehat{\mathbb{Z}}_n = \{\chi_0, \ldots, \chi_{n-1}\}$.*
(d) *Consequently, $\widehat{\mathbb{Z}}_n \cong \mathbb{Z}_n$.*

Proof: (a) and (b) are straightforward. Let now χ be an arbitrary character; then $\chi(1) = \omega^j$ for some $j, 0 \leq j \leq n-1$ by eqn. (5). If follows that $\chi = \chi_j$. Now, (d) is immediate. □

Proposition 1.4 *If G is a direct sum: $G = H_1 \oplus H_2$, and $\varphi_i : H_i \to \mathbb{C}^\times$ is a character of H_i $(i = 1, 2)$, then $\chi = \varphi_1 \oplus \varphi_2$, defined by*

$$\chi(h_1, h_2) := \varphi_1(h_1) \cdot \varphi_2(h_2), \tag{12}$$

is a character of G. Moreover, all characters of G are of this form. Consequently,

$$\widehat{G} \cong \widehat{H}_1 \oplus \widehat{H}_2 \tag{13}$$

Proof: The first statement is clear, and it is easy to verify that the map $\widehat{H}_1 \oplus \widehat{H}_2 \to \widehat{G}$ defined by (12) is injective. Let now $\chi \in \widehat{G}$. The restriction $\varphi_i = \chi|_{H_i}$ is clearly a character of H_i, and it is easy to verify that $\chi = \varphi_1 \oplus \varphi_2$. □

Corollary 1.5 $G \cong \widehat{G}$.

Proof: $G \cong \mathbb{Z}_{n_1} \oplus \cdots \oplus \mathbb{Z}_{n_k}$, hence $\widehat{G} \cong \widehat{\mathbb{Z}}_{n_1} \oplus \cdots \oplus \widehat{\mathbb{Z}}_{n_k} \cong G$ using the previous two propositions. $\square$

We remark that there is no *natural* isomorphism between G and $\widehat{G}$; even for cyclic groups, the isomorphism selected depends on the arbitrary choice of ω. The consequent isomorphism $G \cong \widehat{\widehat{G}}$ is, however, natural:

Corollary 1.6 *G can be identified with $\widehat{\widehat{G}}$ in the following natural way: for $a \in G$, define $\tilde{a} \in \widehat{\widehat{G}}$ by*

$$\tilde{a}(\chi) = \chi(a) \quad (\chi \in \widehat{G}). \tag{14}$$

The map $a \mapsto \tilde{a}$ is an isomorphism of G and $\widehat{\widehat{G}}$.

Proof: Left to the reader. $\square$

Let $\mathbb{C}^G$ denote the space of functions $f : G \to \mathbb{C}$. This is an n-dimensional linear space over $\mathbb{C}$. We introduce an inner product over this space:

$$(f, g) = \frac{1}{n} \sum_{a \in G} \overline{f(a)} g(a) \quad (f, g \in \mathbb{C}^G). \tag{15}$$

Theorem 1.7 *$\widehat{G}$ forms an orthonormal basis in $\mathbb{C}^G$.*

Proof: Orthonormality follows from Cor. 1.2. Completeness follows from Cor. 1.5 which implies that $|\widehat{G}| = n = \dim(\mathbb{C}^G)$. $\square$

Let $\chi_0, \ldots, \chi_{n-1}$ be the characters of $G = \{a_0, \ldots, a_{n-1}\}$. The $n \times n$ matrix

$$C = (\chi_i(a_j)) \tag{16}$$

is the *character table* of G.

Corollary 1.8 *The matrix $A = \frac{1}{\sqrt{n}} C$ is unitary, i. e., $AA^* = A^*A = I$. (A^* is the conjugate transpose of A; I is the $n \times n$ identity matrix.)*

Proof: $A^*A = I$ follows immediately from Theorem 1.7 in view of the formula (15). $\square$

Corollary 1.9 (Second orthogonality relation for characters) *Let $a, b \in G$. Then*

$$\sum_{\chi \in \widehat{G}} \overline{\chi(a)} \chi(b) = \begin{cases} n & \text{if } a = b \\ 0 & \text{otherwise.} \end{cases} \tag{17}$$

First proof: This is a restatement of the fact that $AA^* = I$ in Corollary 1.8. $\square$

Second proof: In view of the identification of G and $\widehat{\widehat{G}}$ (Cor. 1.6), Cor. 1.9 is a restatement of Cor. 1.2 for the abelian group $\widehat{G}$ in place of G. $\square$

We state a special case separately. The following is the dual of Proposition 1.1.

Corollary 1.10 *For any non-zero element $a \in G$,*

$$\sum_{\chi \in \widehat{G}} \chi(a) = 0. \qquad \square$$

2 Fourier Transform

Corollary 2.1 *Any function $f \in \mathbb{C}^G$ can be written as a linear combination of characters:*

$$f = \sum_{\chi \in \widehat{G}} c_\chi \chi. \tag{18}$$

Such a linear combination is also called a *trigonometric sum* since $f(a)$ is expressed as a combination of n^{th} roots of unity. The coefficients c_χ are called the *Fourier coefficients* and are given by the formula

$$c_\chi = (\chi, f). \tag{19}$$

Proof: Expansion (18) exists by Theorem 1.7. The inner product (χ, f) is equal to c_χ by orthonormality. $\square$

The function $\widehat{f} : \widehat{G} \to \mathbb{C}$, defined by

$$\widehat{f}(\chi) = n c_{\overline{\chi}} = \sum_{a \in G} \chi(a) f(a) \quad (\chi \in \widehat{G}), \tag{20}$$

is called the *Fourier Transform* of f. This transformation is easily inverted: using equations (18) and (20), we see that

$$f = \sum_{\chi \in \widehat{G}} c_\chi \chi = \sum_{\chi \in \widehat{G}} \frac{1}{n} \widehat{f}(\overline{\chi}) \chi,$$

hence the formula for the *Inverse Fourier Transform* is

$$f(a) = \frac{1}{n} \sum_{\chi \in \widehat{G}} \widehat{f}(\chi) \chi(-a) \quad (a \in G). \tag{21}$$

We derive a simple consequence.
Let $\delta \in \mathbb{C}^G$ be defined by

$$\delta(a) = \begin{cases} 1 & \text{if } a = 0 \\ 0 & \text{if } a \neq 0 \end{cases} \qquad (a \in G).$$

Corollary 2.2 (a)

$$\widehat{\delta}(\chi) = 1 \quad (\chi \in \widehat{G}). \tag{22}$$

(b)

$$\delta = \frac{1}{n} \sum_{\chi \in \widehat{G}} \chi. \tag{23}$$

Proof: (a) follows from eqn. (20). (b) follows from eqn. (21). (Note that (b) also follows from the second orthogonality relation (17) with $a = 0$.) $\square$

Applying formula (15) to $\widehat{G}$ we obtain the inner product

$$(f, g) = \frac{1}{n} \sum_{\chi \in \widehat{G}} \overline{f(\chi)} g(\chi) \quad (f, g \in \mathbb{C}^{\widehat{G}}) \tag{24}$$

over the space $\mathbb{C}^{\widehat{G}}$. Corollary 1.8 tells us that Fourier transformation is $\sqrt{n}$ times a unitary transformation between $\mathbb{C}^G$ and $\mathbb{C}^{\widehat{G}}$:

Theorem 2.3 (Plancherel formula) *For any $f, g \in \mathbb{C}^G$,*

$$(\widehat{f}, \widehat{g}) = n(f, g). \tag{25}$$

First proof: Using the notation introduced before Cor. 1.8, let

$$\underline{f} = (f(a_0), \ldots, f(a_{n-1})), \underline{g} = (g(a_0), \ldots, g(a_{n-1})),$$
$$\underline{\widehat{f}} = (\widehat{f}(\chi_0), \ldots, \widehat{f}(\chi_{n-1})), \underline{\widehat{g}} = (\widehat{g}(\chi_0), \ldots, \widehat{g}(\chi_{n-1}))$$

As in (16), let $C = (\chi_i(a_j))$ be the character table of G. Then $\underline{\widehat{f}} = \underline{f}C$, $\underline{\widehat{g}} = \underline{g}C$, and

$$\overline{(\widehat{f}, \widehat{g})} = \frac{1}{n} \cdot \underline{\widehat{f}} \cdot \underline{\widehat{g}}^* = \frac{1}{n} \underline{f} C C^* \underline{g}^* = \underline{f} \cdot \underline{g}^* = n \cdot \overline{(f, g)}. \tag{26}$$

(As before, * denotes conjugate transpose.) We made use of the fact that $C = \sqrt{n}A$, hence $CC^* = nAA^* = nI$. (Cor. 1.8). $\square$

Second proof: The map $f \mapsto \widehat{f}$ is clearly linear. Therefore it suffices to prove (25) for elements f, g of a basis of $\mathbb{C}^G$. The functions of δ_a defined by

$$\delta_a(b) = \delta(b - a) \quad (b \in G) \tag{27}$$

(the characteristic vectors of the singletons) form a basis of $\mathbb{C}^G$. Clearly,

$$\widehat{\delta}_a(\chi) = \chi(a), \tag{28}$$

hence by the second orthogonality relation,

$$(\widehat{\delta}_a, \widehat{\delta}_b) = \frac{1}{n} \sum_{\chi \in \widehat{G}} \overline{\chi(a)} \chi(b) = \begin{cases} 1 & \text{if } a = b \\ 0 & \text{otherwise} \end{cases}$$

On the other hand, obviously,

$$(\delta_a, \delta_b) = \left\{ \begin{array}{ll} 1/n & \text{if } a = b \\ 0 & \text{otherwise.} \end{array} \right\} = \frac{1}{n} (\widehat{\delta}_a, \widehat{\delta}_b). \qquad \square$$

Let $\|f\| = \sqrt{(f, f)}$.

Corollary 2.4 $\|\widehat{f}\| = \sqrt{n}\|f\|$. $\square$

The *characteristic function* of a set $A \subseteq G$ is the function $f_A \in \mathbb{C}^G$ defined by

$$f_A(a) = \begin{cases} 1 & \text{if } a \in A \\ 0 & \text{otherwise.} \end{cases} \tag{29}$$

Proposition 2.5 *For $A, B \subseteq G$,*

$$(f_A, f_B) = \frac{1}{n}|A \cap B|. \tag{30}$$

In particular,

$$\sqrt{n}\ \|f_A\| = \sqrt{|A|}. \tag{31}$$

Proof: Evident. $\square$

We note that

$$\widehat{f}_A(\chi_0) = |A|. \tag{32}$$

This is n-times the principal Fourier coefficient of f_A. The remaining Fourier coefficients give important "randomness" information on the set A. Let

$$\Phi(A) = \max\{|\widehat{f}_A(\chi)| : \chi \in \widehat{G}, \chi \neq \chi_0\}.$$

The smaller $\Phi(A)$, the "smoother," more "random looking" the set A is. We shall estimate $\Phi(A)$ for specific "smooth" sets in Sections 5 and 6. Here we give a lower bound which holds for every set $A \subseteq G$.

Proposition 2.6 *For every $A \subseteq G$, if $|A| \leq n/2$, then*

$$\Phi(A) \geq \sqrt{|A|/2}. \tag{33}$$

Proof: By Cor. 2.4 and eqn. (31) we have:

$$\|\widehat{f}_A\|^2 = n\|f_A\|^2 = |A|.$$

On the other hand,

$$n\|\widehat{f}_A\|^2 = \sum_{\chi \in \widehat{G}} |\widehat{f}_A(\chi)|^2 \leq (\widehat{f}_A(\chi_0))^2 + (n-1)\Phi(A)^2 = |A|^2 + (n-1)\Phi(A)^2.$$

Consequently,

$$|A|^2 + (n-1)\Phi(A)^2 \geq n|A|; \qquad \Phi(A)^2 \geq \frac{(n-|A|)|A|}{n-1} \geq \frac{|A|}{2}.$$

$\square$

The "smooth" sets will be those which come close to this bound. The assumption $|A| \leq n/2$ is justified by the following exercise.

Exercise 2.7 Prove: $\Phi(A) = \Phi(G \backslash A)$ for every $A \subseteq G$.

For $A \subseteq G$ and $k \in \mathbb{Z}$, let $kA = \{ka : a \in A\}$.

Exercise 2.8 Prove: If g.c.d.$(k, n) = 1$ then $\Phi(kA) = \Phi(A)$ for every $A \subseteq G$.

In particular $\Phi(-A) = \Phi(A)$.

More generally, Φ is invariant under automorphism of G. Let Aut G denote the automorphism group of G.

Exercise 2.9 Prove: if $\alpha \in$ Aut G then $\Phi(A) = \Phi(\alpha A)$ for every $A \subseteq G$.

Exercise 2.10 Prove: $\Phi(A + a) = \Phi(A)$ for every $a \in G$.
(Here $A + a = \{u + a : u \in A\}$.)

3 Equations over finite abelian groups

We shall consider the following general problem: Let $A_1, \dots, A_k \subseteq G$ and let a be a fixed element of G. Estimate the number of solutions of the equation

$$x_1 + \cdots + x_k = a \quad (x_i \in A_i, i = 1, \dots, k). \tag{34}$$

In particular, decide whether or not a solution exists.

Let $|A_i| = m_i$. Assume for a moment that while the sets A_i are fixed, the element $a \in G$ is selected at random. This makes the *expected number of solutions* equal to

$$\frac{m_1 \cdots m_k}{n} \tag{35}$$

the numerator being the number of k-tuples from $A_1 \times \cdots \times A_k$, and $\frac{1}{n}$ being the chance that a random element $a \in G$ happens to be equal to $\sum_{i=1}^{k} x_i$ for fixed $x_i \in G$.

It is remarkable that under fairly general circumstances, the quantity $m_1 \cdots m_k / n$ will be close to the actual number of solutions for *every* $a \in G$.

We shall give a sufficient condition for this to happen. First of all we observe that the number of solutions will not change if we replace A_k by $A_k - a = \{u - a : u \in A_k\}$ and set the right hand side in eqn. (34) to zero. So it suffices to consider the homogeneous equation

$$x_1 + \cdots + x_k = 0 \quad (x_i \in A_i, i = 1, \dots, k). \tag{36}$$

Let N denote the number of solutions of eqn. (36).

We first describe an explicit formula for N.

Theorem 3.1. *The number of solutions of eqn.* (36) *is*

$$N = \frac{m_1 \cdots m_k}{n} + R, \tag{37}$$

where $m_i = |A_i|$ *and*

$$R = \frac{1}{n} \sum_{\substack{\chi \in \widehat{G} \\ \chi \neq \chi_0}} \prod_{i=1}^{k} \widehat{f}_{A_i}(\chi). \tag{38}$$

Proof: The number of solutions is clearly

$$N = \sum_{\substack{(x_1, \ldots, x_n) \\ x_i \in A_i}} \delta(x_1 + \cdots + x_n) = \frac{1}{n} \sum_{\chi \in \widehat{G}} \sum_{\substack{(x_1, \ldots, x_n) \\ x_i \in A_i}} \chi(x_1 + \cdots + x_k)$$

(we have used eqn. (23)). Since $\chi(x_1 + \cdots + x_n) = \chi(x_1) \cdots \chi(x_k)$, the rightmost sum factors as $\prod_{i=1}^{k} (\sum_{x_i \in A_i} \chi(x_i))$. We recognize the term in the parentheses as $\widehat{f}_{A_i}(\chi)$. In summary,

$$N = \frac{1}{n} \sum_{\chi \in \widehat{G}} \prod_{i=1}^{k} \widehat{f}_{A_i}(\chi).$$

We separate out the term corresponding to χ_0:

$$N = \frac{1}{n} \prod_{i=1}^{k} \widehat{f}_{A_i}(\chi_0) + R.$$

By eqn. (32), $\widehat{f}_{A_i}(\chi_0) = m_i$. This observation concludes the proof. □

The value of this formula depends on our ability to estimate the "error-term" R. In order to be able to conclude that equation (36) has a solution at all, we need to prove that $|R| < (m_1 \cdots m_k)/n$.

The art of estimating $|R|$ and variations of it constitute the *method of trigonometric sums*.

4 The Cauchy-Schwarz trick

In the case $k = 3$, one can give a strong *explicit upper bound* on R under surprisingly general circumstances. It will follow from the estimate that if at least *one of the sets* A_i *is smooth* (all non-principal Fourier coefficients of A_i are small) and the sets are *not too small*, then equation (36) has approximately $(m_1 \cdots m_k)/n$ solutions. Along the way, we shall experience a little Cauchy-Schwarz magic.

Theorem 4.1. *Let $A_1, A_2, A_3 \subseteq G$, $a \in G$, and let N denote the number of solutions of the equation*

$$x_1 + x_2 + x_3 = a \quad (x_i \in A_i, i = 1, 2, 3). \tag{39}$$

Then

$$\left| N - \frac{|A_1||A_2||A_3|}{n} \right| < \Phi(A_3)\sqrt{|A_1||A_2|}. \tag{40}$$

Proof: Applying a translation by $-a$ to A_3 as mentioned in Section 3, transform eqn. (40) into its homogeneous version ($a = 0$). By Exercise 2.10, inequality (40) is invariant under this transformation, so it suffices to consider the homogeneous case.

We thus have to estimate $|R|$ where R is defined by eqn. (38), $k = 3$:

$$R = \frac{1}{n} \sum_{\substack{\chi \in \widehat{G} \\ \chi \neq \chi_0}} \widehat{f}_{A_1}(\chi)\widehat{f}_{A_2}(\chi)\widehat{f}_{A_3}(\chi) \tag{41}$$

It follows that

$$\begin{aligned} |R| &\leq \tfrac{1}{n} \sum_{\substack{\chi \in \widehat{G} \\ \chi \neq \chi_0}} |\widehat{f}_{A_1}(\chi)| \cdot |\widehat{f}_{A_2}(\chi)| \cdot |\widehat{f}_{A_3}(\chi)| \\ &\leq \tfrac{1}{n}\Phi(A_3) \sum_{\chi \in \widehat{G}} |\widehat{f}_{A_1}(\chi)| \cdot |\widehat{f}_{A_2}(\chi)| \end{aligned} \tag{42}$$

By the Cauchy-Schwarz inequality, the right hand side can now be bounded as

$$\sum_{\chi \in \widehat{G}} |\widehat{f}_{A_1}(\chi)| \cdot |\widehat{f}_{A_2}(\chi)| \leq ((\sum_{\chi \in \widehat{G}} |\widehat{f}_{A_1}(\chi)|^2)(\sum_{\chi \in \widehat{G}} |\widehat{f}_{A_2}(\chi)|^2))^{1/2}. \tag{43}$$

By definition (24) and Cor. 2.4, the right hand side here is

$$(n\|\widehat{f}_{A_1}\|^2 n\|\widehat{f}_{A_2}\|^2)^{1/2} = n^2\|f_{A_1}\| \cdot \|f_{A_2}\| = n\sqrt{|A_1||A_2|},$$

using Corollary 2.4 and eqn. (31). Substituting back into (42), we obtain the desired bound. □

Corollary 4.2. *If*

$$\frac{\Phi(A_3)}{|A_3|} < \frac{\sqrt{|A_1||A_2|}}{n} \tag{44}$$

then equation (39) *has at least one solution.*

Proof: Inequality (44) is a restatement of the condition that $|R| < \frac{m_1 m_2 m_3}{n}$. □

The value of this result depends on our ability to bound $\Phi(A)$ for various sets $A \subseteq G$.

Our next objective is to show that "smooth" sets abound.

5 Almost every set is smooth[1]

Recall that we are especially interested in sets for which all the non-principal Fourier coefficients are small. We referred to such sets as "smooth." In this section we show that almost all t-subsets of a finite abelian group G are smooth.

We have already seen (Proposition 2.6) that, for all sets $A \subseteq G$, if $|A| \leq n/2$, then

$$\sqrt{|A|/2} \leq \Phi(A) \leq |A|.$$

This is quite a large range! However, we will now show that, when A *is randomly selected*, $\Phi(A)$ almost always lies near the bottom of this interval; for almost all sets A of size t, we have $\Phi(A) = O(\sqrt{t \ln n})$ (see Theorem 5.14).

Some familiarity with finite probability spaces will required for this section (see e.g. L.B., "Finite Probability Spaces" Lecture Notes). In particular, a key tool is Chernoff's bound, which we state here without proof (see e.g. Alon, Spencer, Erdős pp. 83–85, or L.B. Finite Probability Spaces, Theorem 5.4).

First, we review some basic terminology. A *finite probability space* is a pair (Ω, P) where the nonempty finite set Ω is the *sample space*, thought of as the set of possible outcomes of an experiment. The *probability distribution* $P : \Omega \to \mathbb{R}$ must satisfy $(\forall x \in \Omega)\,(P(x) > 0)$ and $\sum_{x \in \Omega} P(x) = 1$.

An *event* is a subset $A \subseteq \Omega$. We define $P(A) = \sum_{x \in \Omega} P(x)$. Singletons $\{x\}$ are called *elementary events*.

A *random variable* is a function $\xi : \Omega \to \mathbb{R}^n$. For a "real-valued" random variable, we have $n = 1$. We shall treat "complex-valued" random variables $\xi : \Omega \to \mathbb{C}$ as 2-dimensional random variables $\xi : \Omega \to \mathbb{R}^2$. The *expected value* of ξ is defined as $\mathrm{E}(\xi) = \sum_{x \in \Omega} \xi(x) P(x)$.

For concepts of independence, we refer to the lecture notes cited above, or any text on probability. Later in this section, we give a self-contained introduction to a class of discrete stochastic processes called *martingales*.

Theorem 5.1 (Chernoff) *Let ξ_i be real-valued independent random variables satisfying $|\xi_i| \leq 1$ and $E(\xi_i) = 0$. Let $\eta = \sum_{i=1}^{n} \xi_i$. Then for any $a > 0$,*

$$\Pr(\eta \geq a) < e^{-a^2/2n}$$

and

$$\Pr(|\eta| \geq a) < 2e^{-a^2/2n}.$$

Theorem 5.2 *Let $\epsilon > 0$. For all but a $O(n^{-\epsilon})$ fraction of subsets $A \subseteq G$,*

$$\Phi(A) < \sqrt{(1+\epsilon)\, n \ln(n)}.$$

[1] This chapter was contributed by T. Hayes in June 2002.

Proof: As before, let $G = \{a_0, \ldots, a_{n-1}\}$. Choose a random subset $A \subseteq G$ by independently flipping a fair coin to decide whether to include each element a_i.

Let ζ_i be the ± 1-valued indicator variable for the event $a_i \in A$, i. e.,

$$\zeta_i = \begin{cases} 1 & \text{if } a_i \in A \\ -1 & \text{otherwise.} \end{cases}$$

Let $\chi \in \widehat{G}$ be any non-principal character of G. Let $\eta_\chi = \sum_{a_i \in A} \chi(a_i) = \sum_{i=0}^{n-1} \frac{\zeta_i + 1}{2} \chi(a_i)$. By Proposition 1.1, this can be rewritten as

$$2\eta_\chi = \sum_{i=0}^{n-1} \zeta_i \chi(a_i).$$

Now, the summands $\zeta_i \chi(a_i)$ are independent random variables satisfying $|\zeta_i \chi(a_i)| \leq 1$ and $E(\zeta_i \chi(a_i)) = 0$, but unfortunately, since they are complex-valued, Chernoff's bound cannot be applied directly. Instead, we look at the real and imaginary parts of $\zeta_i \chi(a_i)$.

Let $\xi_i = \zeta_i \mathrm{Re}(\chi(a_i))$. Let $\nu = \sum_{i=0}^{n-1} \xi_i$. Then, by Chernoff's bound,

$$\Pr(|\nu| \geq a) < 2e^{-a^2/2n}.$$

Similarly, let $\mu = \sum_{i=1}^{n} \zeta_i \mathrm{Im}(\chi(a_i))$.

$$\Pr(|\mu| \geq a) < 2e^{-a^2/2n}.$$

Since $2\eta_\chi = \nu + \mu i$, the event $\{|\eta_\chi| \geq a\}$ is contained in the union $\{|\nu| \geq a\sqrt{2}\} \cup \{|\mu| \geq a\sqrt{2}\}$. By the union bound, this event has probability at most $4e^{-a^2/n}$.

The event $\{\Phi(A) \geq a\}$ is contained in the union $\cup_{i=1}^{n-1} \{\eta_{\chi_i} > a\}$. Thus, its probability is at most $4(n-1)e^{-a^2/n}$. When $a \geq \sqrt{(1+\epsilon) n \ln n}$, this probability is $O(n^{-\epsilon})$. $\square$

In Theorem 5.2, we assumed that A was a randomly chosen subset of G, and we know that randomly chosen subsets almost always have size approximately $n/2$. We would also like to show that most subsets of a given size are also smooth. To do this, we will need a strengthened version of Chernoff's bounds due to Azuma (cf. Alon, Spencer, Erdős [pp. 83–85, 233–240], or Azuma's original paper). Azuma's inequality applies to random variables which are not independent, but instead satisfy a weaker condition, that of being a "martingale."

Definition 5.3 Let $\xi : \Omega \to \mathbb{R}^n, \zeta : \Omega \to \mathbb{R}^m$ be two random variables defined over the same finite probability space. The *conditional expectation* of ζ with respect to ξ, denoted $\mathrm{E}(\zeta \mid \xi)$, is the unique random variable which is constant on atoms of ξ, and whose value on an atom A of ξ is $E(\zeta \vartheta_A)/\Pr(A)$, where ϑ_A denotes the indicator variable of A, i. e.,

$$\vartheta_A(x) = \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{otherwise.} \end{cases}$$

The *atoms* of ξ are the non-empty preimage sets $\xi^{-1}(y)$, $y \in \mathbb{R}^m$. (Note that these form a partition of Ω.)

Note 5.4 Observe that $\eta = \mathrm{E}(\zeta \mid \xi)$ is a *random variable*. The value $\eta(x)$ depends only on $\xi(x)$, and not on the specific elementary event x.

Exercise 5.5 Prove that $\mathrm{E}(\eta) = \mathrm{E}(\zeta)$, where as above, $\eta = \mathrm{E}(\zeta \mid \xi)$.

Exercise 5.6 Suppose that an elementary event $x \in \Omega$ is drawn at random. You are told $\xi(x)$, and asked to predict $\zeta(x)$. Show that, to minimize your expected error, you should answer $\mathrm{E}(\zeta \mid \xi)(x)$. In this sense, $\mathrm{E}(\zeta \mid \xi)$ is the best guess for the value of ζ given the value of ξ.

Exercise 5.7 Let $\zeta, \xi_1, \xi_2 : \Omega \to \mathbb{R}$ be random variables. The conditional expectation $\mathrm{E}(\zeta \mid \xi_1, \xi_2)$ is defined to equal $\mathrm{E}(\zeta \mid (\xi_1, \xi_2))$. Check that this also equals $\mathrm{E}(\mathrm{E}(\zeta \mid \xi_1) \mid \xi_2)$. $\mathrm{E}(\zeta \mid \xi_1, \xi_2)$ is a "best" guess for ζ given the values for both ξ_1 and ξ_2. Extend this to $\zeta, \xi_1, \ldots, \xi_k$.

After these introductory definitions, we come to the central concept.

Definition 5.8 A sequence of random variables $\sigma_0, \sigma_1, \ldots, \sigma_n, \ldots$ is called a *martingale* if, for every $i \geq 1$, $\mathrm{E}(\sigma_i \mid \sigma_0, \ldots, \sigma_{i-1}) = \sigma_{i-1}$.

Example 5.9 If $\xi_1, \ldots, \xi_n$ are independent random variables satisfying $E(\xi_i) = 0$, then the partial sums

$$\sigma_i = \sum_{j=1}^{i} \xi_j$$

form a martingale. (By convention, the empty sum $\sigma_0 = 0$.)

Example 5.10 Consider the following gambling proposition. A player, P, is allowed to flip a fair coin up to 10 times, but may stop earlier. If heads comes up at least 2 more times than tails, the player wins \$100, otherwise, he pays \$30. Let μ be the amount which P wins (the random variable μ is thus defined in terms of P's strategy!)

Obviously, P's best strategy is to stop flipping iff heads has already come up two more times than tails. (What is $E(\mu)$ in this case?)

As the game progresses, P's expected winnings fluctuate based on the outcomes of the coin flips. Let σ_i be P's expected winnings after the coin has been flipped i times. Obviously, σ_i is a random variable, determined by the first i coin flips. Formally, σ_i is defined as

$$\sigma_i = \mathrm{E}(\mu \mid \xi_1, \ldots, \xi_i),$$

where ξ_i are indicator variables for the coin flips.

Exercise 5.11 Show that $\sigma_0 = E(\mu)$ and $\sigma_{10} = \mu$. Show that $\sigma_0, \sigma_1, \ldots, \sigma_{10}$ is a martingale (regardless of P's strategy!) *Hint:* For a given initial sequence $\xi_1, \ldots, \xi_i$ (for which P would not already stop) one way to compute σ_i would be to compute σ_{i+1} for each of the two possible values of ξ_{i+1}, then average.

Exercise 5.12 Suppose that $\sigma_0, \ldots, \sigma_n, \ldots : \Omega \to \mathbb{C}$ satisfies the martingale condition, for every $i \geq 1$ $\mathrm{E}(\sigma_i \mid \sigma_0, \ldots, \sigma_{i-1}) = \sigma_{i-1}$ (a *complex-valued martingale*). Show that the sequences $\mathrm{Re}(\sigma_0), \ldots, \mathrm{Re}(\sigma_n), \ldots : \Omega \to \mathbb{R}$ and $\mathrm{Im}(\sigma_0), \ldots, \mathrm{Im}(\sigma_n), \ldots : \Omega \to \mathbb{R}$ are real-valued martingales.

Theorem 5.13 (K. Azuma) *Let $\sigma_0, \ldots, \sigma_n, \ldots$ be a real-valued martingale such that for every i, $|\sigma_i - \sigma_{i-1}| \leq 1$. Then*

$$\Pr(|\sigma_n - \sigma_0| \geq a) \ < \ 2e^{-a^2/2n}. \tag{45}$$

Theorem 5.14 (T. Hayes) *Let $\epsilon > 0$. Let $k \leq n/2$. For all but an $O(n^{-\epsilon})$ fraction of subsets $A \subseteq G$ such that $|A| = t$,*

$$\Phi(A) < 4\sqrt{(1+\epsilon)\,|A| \ln(n)}.$$

Proof: Let A be selected uniformly at random from all t-subsets of G. As in the proof of Theorem 5.2, let ζ_i be the indicator variable for the event $a_i \in A$. In this case, $\zeta_0, \ldots, \zeta_{n-1}$ are far from independent; for instance, they always satisfy $\sum_{i=0}^{n-1} \zeta_i = t$.

Let χ be a non-principal character, and define $\eta_\chi = \sum_{a_i \in A} \chi(a_i) = \sum_{i=0}^{n} \zeta_i \chi(a_i)$. Suppose we look at the ζ_i one at a time, and, after each one, make the best prediction we can about the value of η_χ. In other words, let $\sigma_0 = \mathrm{E}(\eta_\chi) = 0$, and, for $1 \leq i \leq n$, let $\sigma_i = \mathrm{E}(\eta_\chi \mid \zeta_0, \ldots, \zeta_{i-1})$. Note that, since $\zeta_0, \ldots, \zeta_{i-1}$ determine η_χ, $\sigma_n = \eta_\chi$.

Claim: $\sigma_0, \ldots, \sigma_n$ is a martingale, and for $1 \leq i \leq n$, $|\sigma_i - \sigma_{i-1}| \leq 2$.

The formal proof of this claim is an exercise in conditional probabilities, and is left to the reader. The intuition is: by definition, σ_i is the unique best prediction of η_χ given $\zeta_0, \ldots, \zeta_{i-1}$. But σ_{i+1} would be the best possible prediction if we knew ζ_i as well. So another best prediction would be $\mathrm{E}(\sigma_{i+1} \mid \zeta_0, \ldots, \zeta_{i-1})$, which equals $\mathrm{E}(\sigma_{i+1} \mid \sigma_0, \ldots, \sigma_i)$. This proves the martingale condition.

To see that $|\sigma_i - \sigma_{i-1}| \leq 2$, observe that substituting one element of A for another can change η_χ by at most two. Therefore, learning whether any particular element is included or excluded cannot change the expectation of η_χ by more than two.

By Exercise 5.12, the real and imaginary parts of $\sigma_0, \ldots, \sigma_n$ are martingales, which clearly satisfy the same distance bound $|\mathrm{Re}(\sigma_i) - \mathrm{Re}(\sigma_{i-1})| \leq 2$. By Azuma's Inequality applied to $\mathrm{Re}(\sigma_n)/2$ (and since $\sigma_0 = 0$), we have

$$\Pr\left(|\mathrm{Re}(\sigma_n)| \geq a/\sqrt{2}\right) < 2e^{-a^2/4n},$$

and similarly,

$$\Pr\left(|\mathrm{Im}(\sigma_n)| \geq a/\sqrt{2}\right) < 2e^{-a^2/4n}.$$

From this it follows that

$$\Pr(|\eta_\chi| \geq a) = \Pr(|\sigma_n| \geq a) < 4e^{-a^2/4}.$$

By the usual union bound, this gives us

$$\Pr(|\Phi(A)| \geq a) < 4(n-1)e^{-a^2/4n},$$

which is $O(n^{-\epsilon})$ when $a \geq 2\sqrt{(1+\epsilon)n \ln n}$. $\square$

Remark: The constants appearing in Theorems 5.2 and 5.14 can be improved by a factor of $\sqrt{2}$, using an extension of Azuma's Inequality to higher dimensions, due to T. Hayes. He showed that for $\mathbb{R}^n$-valued martingales $\sigma_0, \ldots, \sigma_n$ satisfying $|\sigma_i - \sigma_{i-1}| \leq 1$, the inequality

$$\Pr(|\sigma_n - \sigma_0| \geq a) < 2e^2 e^{-a^2/2n}$$

holds. Using the isometry between $\mathbb{C}$ and $\mathbb{R}^2$ as Euclidean spaces, this can also be applied to complex-valued martingales.

6 Gauss sums and the Fourier coefficients of cyclotomic classes

Fourier transforms in this section will be taken over $G = (\mathbb{F}_q, +)$, the additive group of the field of the order q. Our aim is to estimate the Fourier coefficients of sets of the form

$$H(q,k) = \{a^k : a \in \mathbb{F}_q^\times\} \tag{46}$$

and related sets. We shall see that all these sets are "smooth" (all Fourier coefficients are small) (see Theorem 6.8).

Exercise 6.1. Prove: If $q = p^s$, p prime, then $(\mathbb{F}_q, +) \cong \mathbb{Z}_p \times \cdots \times \mathbb{Z}_p$ (s times).

Exercise 6.2. Let χ_1 be a nonprincipal character of $(\mathbb{F}_q, +)$. For $a, b \in \mathbb{F}_q$, set

$$\chi_a(b) = \chi_1(ab). \tag{47}$$

Prove:

(i) For every $a \in \mathbb{F}_q$, χ_a is a character of $(\mathbb{F}_q, +)$.

(ii) For $a \neq b$, $\chi_a \neq \chi_b$.

(iii) Every character of $(\mathbb{F}_q, +)$ is of the form χ_a for some $a \in \mathbb{F}_q$.

(Observe that this notation is compatible with our previous notation χ_0 for the principal character.)

Exercise 6.3. Prove: if $A \subseteq \mathbb{F}_q$ and $a, b \in \mathbb{F}_q$, $a \neq 0$, then $\Phi(aA + b) = \Phi(A)$.

Exercise 6.4. Prove:

(i) $H(q,k)$ (defined by formula (46)) is a subgroup of the multiplicative group $\mathbb{F}_q^\times$.

(ii) If
$$d = \text{g.c.d.}(q-1,k) \quad \text{then} \quad H(q,k) = H(q,d) \tag{48}$$

(iii) Every subgroup of $\mathbb{F}_q^\times$ is of the form $H(q,k)$ for some $k|q-1$.

(iv) If $k|q-1$ then $H(q,k)$ is the unique subgroup of $\mathbb{F}_q^\times$ of order $(q-1)/k$.

(Hint: Use the fact that $\mathbb{F}_q^\times$ is a cyclic group.)

The cosets of the subgroups of the multiplicative group $\mathbb{F}_q^\times$ are called *cyclotomic classes*. By the previous exercise these sets are of the form
$$b \cdot H(q,k) = \{ba^k : a \in \mathbb{F}_q^\times\} \tag{49}$$
form some $b \in \mathbb{F}_q^\times$ and $k|q-1$. By Ex.6.3, the parameter of Φ of such a set does not depend on b.

One of the most fascinating features of the theory of characters is the interplay between the additive and multiplicative structures of finite fields captured by the concept of Gauss sums.

Definition. An *additive character* of $\mathbb{F}_q$ is a character of the additive group $(\mathbb{F}_q, +)$. For such a character χ we have for all $a, b \in \mathbb{F}_q$:
$$\chi(a+b) = \chi(a)\chi(b); \tag{50}$$
$$\chi(0) = 1; \tag{51}$$
$$\chi(-a) = \overline{\chi(a)}. \tag{52}$$

A *multiplicative character* of $\mathbb{F}_q$ is a character ψ of the multiplicative group $\mathbb{F}_q^\times$, extended to all of $\mathbb{F}_q$ by setting $\psi(0) = 0$. For such a character we have for all $a, b \in \mathbb{F}_q$:
$$\psi(ab) = \psi(a)\psi(b); \tag{53}$$
$$\psi(1) = 1; \tag{54}$$
$$\psi(a^{-1}) = \overline{\psi(a)}. \tag{55}$$

Example. For p a prime, the Legendre symbol
$$\psi(a) = \left(\frac{a}{p}\right)$$
is a multiplicative character of $\mathbb{F}_p$.

Definition. Let χ be an additive character and ψ a multiplicative character of $\mathbb{F}_q$. The sum

$$S(\chi, \psi) = \sum_{a \in \mathbb{F}_q} \chi(a)\psi(a) \tag{56}$$

is called a *Gauss sum* over $\mathbb{F}_q$.

Exercise 6.5. Prove:

(i) $S(\chi_0, \psi_0) = q - 1$, where χ_0 and ψ_0 are the principal additive and multiplicative characters, resp., of $\mathbb{F}_q$.

(ii) $S(\chi_0, \psi) = 0$ if $\psi \neq \psi_0$.

(iii) $S(\chi, \psi_0) = -1$ if $\chi \neq \chi_0$.

In the cases not covered by this exercise, a result of startling simplicity and uniformity holds.

Theorem 6.6. *If neither χ nor ψ is principal, then*

$$|S(\chi, \psi)| = \sqrt{q}. \tag{57}$$

Proof:

$$\begin{aligned} |S(\chi,\psi)|^2 &= \overline{S(\chi,\psi)}S(\chi,\psi) \\ &= \textstyle\sum_{a\in\mathbb{F}_q}\sum_{b\in\mathbb{F}_q} \overline{\chi(a)}\chi(b)\overline{\psi(a)}\psi(b) \\ &= \textstyle\sum_{a\in\mathbb{F}_q^\times}\sum_{b\in\mathbb{F}_q^\times} \chi(b-a)\psi(ba^{-1}) \end{aligned}$$

Let $c = ba^{-1}$. Then our expression turns into

$$\sum_{c\in\mathbb{F}_q^\times}\sum_{a\in\mathbb{F}_q^\times} \chi(ac-a)\psi(c) = \sum_{c\in\mathbb{F}_q^\times} \psi(c) \sum_{a\in\mathbb{F}_q^\times} \chi(a(c-1)). \tag{58}$$

For $c \neq 1$, we have (by Prop. 1.1)

$$\sum_{a\in\mathbb{F}_q} \chi(a(c-1)) = 0. \tag{59}$$

Consequently,

$$\sum_{a\in\mathbb{F}_q^\times} \chi(a(c-1)) = -1. \tag{60}$$

For $c = 1$, this sum is $q - 1$. Substituting back into eqn. (58), we obtain

$$|S(\chi,\psi)|^2 = \psi(1)\cdot(q-1) - \sum_{\substack{c\,\in\,\mathbb{F}_q^\times \\ c\,\neq\,1}} \psi(c) = \psi(1)\cdot q - \sum_{c\in\mathbb{F}_q^\times} \psi(c).$$

The last term is zero again by Prop. 1.1, hence

$$|S(\chi,\psi)|^2 = \psi(1)\cdot q = q. \qquad \square$$

The following simple sieve idea links the Fourier coefficients of the cyclotomic classes to Gaussian sums.

Lemma 6.7. *Let $k|q-1$ and let A be the (unique) subgroup of index k in $\mathbb{F}_q^\times$. Let $\psi_0,\dots,\psi_{k-1}$ be the characters of $\mathbb{F}_q^\times/A$. For $a\in\mathbb{F}_q^\times$, define $\psi_i(a) := \psi_i(aA)$ to make ψ_i a multiplicative character of $\mathbb{F}_q$. With this notation, for every additive character χ, we have*

$$\widehat{f}_A(\chi) = \frac{1}{k}\sum_{i=0}^{k-1} S(\chi,\psi_i). \tag{61}$$

Proof:

$$\sum_{i=0}^{k-1} S(\chi,\psi_i) = \sum_{a\in\mathbb{F}_q^\times}\chi(a)\sum_{i=0}^{k-1}\psi_i(a). \tag{62}$$

By Corollary 1.10, $\sum_{i=0}^{k-1}\psi_i(a) = 0$ unless $a\in A$, in which case $\sum_{i=0}^{k-1}\psi_i(a) = k$. From (62) we thus obtain

$$\sum_{i=0}^{k-1} S(\chi,\psi_i) = k\sum_{a\in A}\chi(a) = k\widehat{f}_A(\chi). \quad \square$$

The estimate promised at the beginning of this section now follows.

Theorem 6.8. *Let A be a cyclotomic class in $\mathbb{F}_q$. Then*

$$\Phi(A) < \sqrt{q}.$$

.

Proof: By Exercises 6.3 and 6.4, we may assume that $A = H(q,k)$ for some $k|q-1$. Then A is a subgroup of index k in $\mathbb{F}_q^\times$. Let $\psi_0,\dots,\psi_{k-1}$ be the multiplicative characters of the factor group $\mathbb{F}_q^\times/A$. We shall regard ψ_i as multiplicative characters of $\mathbb{F}_q$ through the map $\mathbb{F}_q^\times \to \mathbb{F}_q^\times/A \to \mathbb{C}^\times$. For every nonprincipal additive character χ we have by Lemma 6.7:

$$|\widehat{f}_A(\chi)| \le \frac{1}{k}\sum_{i=0}^{k-1}|S(\chi,\psi_i)| \le \frac{1}{k}(1+(k-1)\sqrt{q}) < \sqrt{q}.$$

(We used Exercise 6.5 (iii) and Theorem 6.6.) $\square$

Remark 6.1 As $k' := \gcd(k, q-1)$ increases, A becomes smaller ($|A| = (q-1)/k'$), so one might expect $\Phi(A)$ to become smaller. This is not reflected in Theorem 6.8; indeed, for $k' > \sqrt{q}$ we have $|A| < \sqrt{q}$ and so the bound $\Phi(A) < \sqrt{q}$ becomes weaker than even the trivial upper bound $\Phi(A) \leq |A|$. In fact, I believe, but have not checked, that for cyclotomic sets A of size $|A| = o(\sqrt{q})$, the trivial bound $\Phi(A) \leq |A|$ is rather tight.

Compare this with Theorem 5.14, which says that for a *random* set A of size t, we have almost surely $\Phi(A) = O(\sqrt{t \log n})$. So while arguments involving random sets of size greater than $c\sqrt{n}$ can be "derandomized" using cyclotomic sets, this method does not seem to work for $|A| = o(\sqrt{n})$.

7 Fermat's Last Theorem over finite fields

In order to prove Theorem 0.2, all we have to do is to combine Theorem 4.1 (the explicit estimate of the error term for the number of solutions) with Theorem 6.8 (the estimate of the Φ parameter for cyclotomic classes). We summarize the conclusion.

Theorem 7.1. *Let $k|q-1$ be an integer, A_1 $A_2 \subseteq \mathbb{F}_q$, and let N denote the number of solutions of the equation*

$$x + y = z^k \quad (x \in A_1, y \in A_2, z \in \mathbb{F}_q^\times).$$

Then

$$\left| N - \frac{|A_1||A_2|(q-1)}{q} \right| < k\sqrt{|A_1||A_2|q}. \tag{63}$$

Proof: Let $A_3 = \{a^k : a \in \mathbb{F}_q^\times\} = H(q,k)$, and let N' denote the number of solutions of the equation

$$x + y = u \quad (x \in A_1, y \in A_2, u \in A_3).$$

Note that $|A_3| = (q-1)/k$. Since $\mathbb{F}_q$ contains k k^{th} roots of unity, the equation $z^k = u$ has precisely k solutions for every $u \in A_3$. Consequently, $N = kN'$.

By Theorem 4.1,

$$\left| N' - \frac{|A_1||A_2|(q-1)}{kq} \right| < \Phi(A_3)\sqrt{|A_1||A_2|} < \sqrt{|A_1||A_2|q},$$

where the last inequality follows from Theorem 6.7. □

Now, let $l_i = (q-1)/|A_i|$ (not necessarily integers) and assume inequality (2) holds:

$$q \geq k^2 l_1 l_2 + 4.$$

Under this assumption it is easy to verify that the error bound in (63) is less than the main term. This implies $N \neq 0$, proving Theorem 0.2.

Here is what we have to verify:

$$k\frac{(q-1)\sqrt{q}}{\sqrt{l_1 l_2}} < \frac{(q-1)^3}{l_1 l_2 q},$$

or equivalently,

$$k^2 l_1 l_2 < q \cdot (1 - \frac{1}{q})^4.$$

But $q \cdot (1 - \frac{1}{q})^4 > q - 4 \geq k^2 l_1 l_2$ indeed. □

Strictly speaking, this proof verifies Theorem 0.2 for the case $k|q-1$ only. The general case reduces to this one via Exercise 6.4(ii):

$$\{a^k : a \in \mathbb{F}_q^\times\} = \{a^d : a \in \mathbb{F}_q^\times\},$$

where $d = \text{g.c.d.}(q-1, k)$. □

8 Some literature

The methods discussed have a long history, going back to the work of Gauss, Jacobi, and Eisenstein in the first half of the 19th century. An accessible source of substantial material on the subject is

K. Ireland and M. Rosen: *A Classical Introduction to Modern Number Theory*, Springer 1982,

esp. Chapters 8, 10, 11. Each chapter is followed by an illuminating discussion of the literature. This includes a brief outlook on A. Weil's conjectures regarding the zeta function of certain algebraic hypersurfaces over finite fields (stated in 1949 and proved by B. Dwork (1959) and P. Deligne (1973)), the key to estimating the number of solutions of polynomial equations over finite fields. No elementary treatment of this subject is known; for an exposition, see

N. Katz: An overview of Deligne's proof of the Riemann hypothesis for varieties over finite fields, *Proc. Symp. in Pure Math.* **28** (1976), pp. 275-305.

With the use of algebraic geometry, A. Weil solved the case of a nonsingular algebraic curve in a paper, published in 1948. An elementary treatment of this result originated with the work of S. A. Stepanov (1972). A complete account of the method is given by

W. M. Schmidt: *Equations over Finite Fields: An Elementary Approach*, Lect. Notes in Math. Vol. 536, Springer 1976.

The development of the method of trigonometric sums has been motivated by its applications to age-old problems of additive number theory; most notably to the Waring problem in the work of Hardy, Ramanujan, Littlewood, Vinogradov, Davenport, Hasse, Hua L-K. during the first half of this century. An exposition of their methods along with more recent results is given in

R. C. Vaughan: *The Hardy–Littlewood Method*, Cambridge University Press, 1981.

This monograph lists over 400 items of literature, and comments on a substantial portion of them.

The Cauchy-Schwarz trick is apparently folklore and it is difficult to pinpoint an original reference. Variants of it appear for instance in the following papers:

G. A. Freiman: What is the structure of K if $K + K$ is small? Lecture Notes in Math. Vol. 1240 (1987), pp. 109-134.

I. Z. Ruzsa: Essential components, *Proc. London Math. Soc.* **54** (1987), 38-56. See "Statement 7.1."

The result, in the compact form stated in Theorem 4.1, was described to me by *Endre Szemerédi.* I owe him gratitude for this classic gem, which provided the inspiration for these notes.

An application of Theorem 4.1 in the theory of computing (to the analysis of the computational model called "branching program") appears in the paper

M. Ajtai, L. Babai, P. Hajnal, J. Komlós, P. Pudlák, V. Rödl, E. Szemerédi, G. Turán: Two lower bounds for branching programs, Proc. 18th ACM Symp. on Theory of Computing, Berkeley CA 1986, pp. 30-38.

A very elegant proof of Azuma's Inequality, which we needed for the proof of Theorem 5.6, can be found in Appendix A of

N. Alon, J. Spencer, P. Erdős: *The Probabilistic Method.* John Wiley & Sons, Inc. New York.

This is a must-read book for any combinatorist.

The original source for Azuma's inequality is

K. Azuma: *Weighted Sums of Certain Dependent Random Variables.* Tôhoku Math. Journ. **19** (1967) 357–367.

The extension of Azuma's Inequality to martingales in higher dimensional spaces, referred to in Section 5, can be found in

T. Hayes: *A Large-Deviation Inequality for Vector-valued Martingales,* to appear in Combinatorics, Probability and Computing. (see http://www.cs.uchicago.edu/research/publications/techreports/TR-2001-24)