Since $|A \pm A| \geq |A|$ and $|A| \leq |F|^{1/2}$, we have $|G| \leq c\frac{|A|^{5/2}}{|A \pm A|} \leq c|F|^{3/4}$. But then, from Theorem 9.41, (2.8) and the fact that $|A \cap G| \geq |A|/3$ we see that if $|G| = O(|F|^{3/4})$, then

$$\begin{aligned} |A \pm A| &\geq |(A \cap G) \pm (A \cap G)| \\ &\geq c\frac{|A \cap G|^4}{E(A \cap G, A \cap G)} \\ &\geq c|A \cap G|^{5/2}/|G| \\ &\geq c|A|^{5/2}/|G|, \end{aligned}$$

a contradiction. □

Exercises

9.7.1 Prove Lemma 9.39. (Hint: in Section 9.4 it was demonstrated that $F^\times$ is a cyclic group of order $|F^\times| = |F| - 1$.)

9.7.2 Prove (9.16).

9.7.3 Prove Lemma 9.40. (Hint: use (9.16) and Lemma 2.9.)

9.7.4 [44] Let $F = F_p$ be a finite field of prime order, and let A be an additive set in $F^\times$ such that $|A| \leq |F|^{1/2}$. Using Theorem 9.45, prove that $|A \cdot (A - A) + A \cdot (A - A)| = \Omega(|A|^{5/4})$. Use this to derive another proof of Corollary 2.58.

9.8 Cyclotomic fields, and the uncertainty principle

We now recall some of the elementary theory of cyclotomic fields $\mathbf{Q}(\omega)$, and apply this to obtain an uncertainty principle for the Fourier transform on $\mathbf{Z}_p$.

Definition 9.47 (Cyclotomic field) Let $n \geq 1$ be any positive integer. An nth *root of unity* is any complex number $\omega \in \mathbf{C}$ such that $\omega^n = 1$. An nth root of unity ω is said to be *primitive* if ω is not an mth root of unity for any $1 \leq m < n$. We define the *cyclotomic field of order n* to be the field $\mathbf{Q}(\omega)$ obtained by adjoining a primitive nth root of unity to the rationals $\mathbf{Q}$. We define the nth *cyclotomic polynomial* $\Phi_n \in \mathbf{C}[z]$ to be the polynomial $\Phi_n(z) := \prod_\omega (z - \omega)$, where ω ranges over the primitive nth roots of unity.

It is easy to see that for each n, there are $\phi(n)$ primitive roots of unity, and they are all powers of each other. Thus there is only one cyclotomic field $\mathbf{Q}(\omega)$ for each order n. In particular we see that Φ_n is a monic polynomial of degree $\phi(n)$. Some further basic properties of Φ_n as follows.

Lemma 9.48 *Φ_n has integer coefficients (thus $\Phi_n \in \mathbf{Z}[z]$), and is irreducible in $\mathbf{Z}[z]$. Furthermore we have $\Phi_n(1) = p$ when n is a prime power $n = p^k$, $\Phi_1(1) = 0$, and $\Phi_n(1) = 1$ otherwise.*

Proof We first observe from the factor theorem that

$$z^n - 1 = \prod_{\omega:\omega^n=1} (z - \omega) \text{ for any } n \geq 1.$$

Since every nth root of unity is a primitive dth root of unity for some d, we obtain

$$z^n - 1 = \prod_{d|n} \Phi_d(z). \tag{9.21}$$

Thus one can obtain $\Phi_n(z)$ by factoring out $\prod_{d|n;d<n} \Phi_d(z)$ from $z^n - 1$. By an easy induction on n this implies that Φ_n is a monic polynomial with integer coefficients. Since $(z^n - 1)/\Phi_1(z) = (z^n - 1)/(z - 1)$ approaches n as $z \to 1$, we obtain the formula

$$n = \prod_{d|n;d>1} \Phi_d(1).$$

Taking logarithms and using Exercise 9.4.1 we conclude that

$$\sum_{d|n;d>1} \Lambda(d) = \sum_{d|n;d>1} \log \Phi_d(1)$$

for all $n \geq 1$, where $\Lambda(d) := \log p$ when d is a prime power $d = p^k$ for some $k \geq 1$, and $\Lambda(d) = 0$ otherwise (cf. Exercise 1.10.6). Another easy induction on n then shows that $\Phi_n(1) = e^{\Lambda(n)}$ for all $n > 1$, which gives the desired formula for $\Phi_n(1)$.

Now we prove the irreducibility. When n is prime this can be easily verified from Eisenstein's criterion (Exercise 9.8.3), but the general case is trickier. We use an argument of Gauss. Suppose for contradiction that Φ_n is reducible in $\mathbf{Z}[z]$, then we can partition the primitive nth roots of unity into two disjoint non-empty classes A and B such that the monic polynomials $f(z) := \prod_{\omega\in A}(z - \omega)$ and $g(z) := \prod_{\omega\in B}(z - \omega)$ lie in $f, g \in \mathbf{Z}[z]$. Of course we have $\Phi_n = fg$. Since any two primitive nth roots are powers of each other, we can find an $\omega \in A$ such that $\omega^m \in B$ for some integer m. By decomposing m into primes and arguing by contradiction, we can in fact locate a prime p and an $\omega \in A$ such that $\omega^p \in B$. This implies that the polynomials $f(z)$ and $g(z^p)$ have a common root, and hence by the Euclidean algorithm we can find a non-trivial monic polynomial $h(z) \in \mathbf{Z}[z]$ which divides both $f(z)$ and $g(z^p)$. This implies that $\Phi_n(z^p) = f(z^p)g(z^p)$ contains a factor of $h(z^p)h(z)$; by (9.21) we see that $z^{np} - 1$ also contains a factor of $h(z^p)h(z)$.

Now we work in the finite field F_p. In that setting we have $h(z^p) = h(z)^p$ and $(z^n - 1)^p = z^{np} - 1$ (cf. Exercise 9.4.7) and hence $(z^n - 1)^p$ contains a factor of

$h(z)^{p+1}$; in particular $z^n - 1$ must contain a factor of $h(z)^2$ in F_p (cf. Exercise 9.4.1). Taking formal derivatives, this implies that $z^n - 1$ and nz^{n-1} have a common factor of $h(z)$; but from the Euclidean algorithm and the fact that $n \neq 0 \pmod{p}$ we see that these polynomials have a least common multiple of 1, contradiction. □

As a consequence of Lemma 9.48 we obtain a useful criterion for non-vanishing of polynomial expressions of roots of unity, which was already exploited in the proof of Theorem 9.20.

Lemma 9.49 *Let p be a prime and q be a power of p. Let $P \in \mathbf{Z}[t_1, \ldots, t_k]$ be a polynomial with integer coefficients such that $P(z_1, \ldots, z_k) = 0$ for some qth roots of unity $z_1, \ldots, z_k$. Then the integer $P(1, \ldots, 1)$ is divisible by p.*

Proof Let ω be a primitive qth root of unity, then $z_i = \omega^{n_i}$ for some integers n_i. If we let $Q(t) := P(t^{n_1}, \ldots, t^{n_k})$, then $Q(\omega) = 0$. Thus $Q(t)$ shares a root in common with the irreducible polynomial $\Phi_q(t)$, which must then be a factor of $Q(t)$. Thus $Q(1) = P(1, \ldots, 1)$ has $\Phi_q(1) = p$ as a factor. □

We apply this lemma to prove a non-vanishing result on generalized Vandermonde determinants. We first need a coefficient computation.

Proposition 9.50 *[355] Let $n_1, \ldots, n_k$ be non-negative integers, and let $P \in \mathbf{Z}[z_1, \ldots, z_k]$ be the polynomial*

$$P(z_1, \ldots, z_k) = \sum_{\pi \in S_k} \operatorname{sgn}(\pi) \prod_{i=1}^{k} z_i^{n_{\pi(i)}}$$

(cf. (9.3)). Then we can factor $P = \Delta_k Q$, where $Q \in \mathbf{Z}[z_1, \ldots, z_k]$ is such that

$$Q(1, \ldots, 1) = \Delta_k(n_1, \ldots, n_k)/\Delta_k(1, \ldots, k).$$

Proof The expression $P(z_1, \ldots, z_k)$ can also be interpreted as the determinant of the $k \times k$ matrix $(z_i^{n_j})_{1 \leq i,j \leq k}$. This shows in particular that P vanishes when any two of the z_i are equal. Dividing out the factors of $z_i - z_j$ using long division and applying Definition 9.6 we conclude the existence of a polynomial $Q \in \mathbf{Z}[z_1, \ldots, z_k]$ such that $P = \Delta_k Q$. It remains to compute $Q(1, \ldots, 1)$. To do this we introduce the normalized differentiation operators $D_i := z_i \frac{d}{dz_i}$, and consider the expression $D_1^0 D_2^1 \ldots D_k^{k-1} P(1, \ldots, 1)$. We split P into factors

$$P(z_1, \ldots, z_k) = \prod_{1 \leq i < j \leq k} (z_j - z_i) \times Q(z_1, \ldots, z_k)$$

and apply the Leibniz rule $D_i(fg) = (D_i f)g + f(D_i g)$ repeatedly. Observe that there are $\binom{k}{2}$ linear factors in the expression to be differentiated, all of which vanish at $(1, \ldots, 1)$. There are also $\binom{k}{2}$ derivatives to be applied. Thus the only

terms in the Leibniz rule which do not vanish at $(1, \ldots, 1)$ are those in which all the derivatives land on the linear factors. Furthermore each derivative must land on a distinct linear factor to yield a non-zero term. But this means that each of the D_k derivatives must land on one of the $z_k - z_i$ factors with $i < k$ (and there are $(k-1)!$ ways this can happen); similarly the D_{k-1} derivatives must then land on one of the $z_{k-1} - z_i$ factors with $i < k-1$ (with $(k-2)!$ ways this can happen), and so forth. We conclude that

$$D_1^0 D_2^1 \ldots D_k^{k-1} P(1, \ldots, 1) = (k-1)! \cdots 1!0! Q(1, \ldots, 1) = \Delta_k(1, \ldots, k) Q(1, \ldots, 1).$$

On the other hand, since each monomial $z_1^{n_1} \cdots z_k^{n_k}$ is an eigenfunction of D_i with eigenvalue n_i, we see from definition of P that

$$D_1^0 D_2^1 \cdots D_k^{k-1} P(z_1, \ldots, z_k) = \sum_{\pi \in S_k} \mathrm{sgn}(\pi) \prod_{i=1}^{k} {n_{\pi(i)}}^{i-1} z_i^{n_{\pi(i)}}.$$

Substituting $z_1 = \cdots = z_k = 1$ and applying (9.3) we obtain

$$D_1^0 D_2^1 \cdots D_k^{k-1} P(z_1, \ldots, z_k) = \Delta_k(n_1, \ldots, n_k).$$

Combining this with the previous identity, the claim follows. □

Combining Proposition 9.50 with Lemma 9.49 we obtain

Lemma 9.51 (Chebotarev's lemma) *Let $q = p^\alpha$ be a prime power, let $1 \leq k < p$, and let $z_1, \ldots, z_k$ be distinct qth roots of unity. Let $n_1, \ldots, n_k$ be integers which are distinct modulo p. Then the $k \times k$ matrix $(z_i^{n_j})_{1 \leq i,j \leq k}$ has non-zero determinant.*

Indeed, Chebotarev's lemma follows since $\Delta_k(z_1, \ldots, z_k)$ is non-zero and $\Delta_k(n_1, \ldots, n_k)$ is not divisible by p. We note that while this result was proved by Chebotarev in 1926 (see [338]), it has been independently rediscovered and reproved a number of times [278], [71], [263], [102], [355], [120], [131]. As a consequence of this lemma, one easily establishes the following uncertainty principle for $\mathbf{Z}_p$:

Theorem 9.52 *[355] Let p be a prime number. Let $f : \mathbf{Z}_p \to \mathbf{C}$ be a random variable, and let $\hat{f} : \mathbf{Z}_p \to \mathbf{C}$ be its Fourier transform (using the standard bicharacter $e(x, \xi) = \exp(2\pi i x \xi / p)$). Then we have $|\mathrm{supp}(f)| + |\mathrm{supp}(\hat{f})| \geq p + 1$. Conversely, if A and B are two non-empty subsets of $\mathbf{Z}/p\mathbf{Z}$ such that $|A| + |B| \geq p + 1$, then there exists a function f such that $\mathrm{supp}(f) = A$ and $\mathrm{supp}(\hat{f}) = B$.*

We leave the deduction of Theorem 9.52 from Lemma 9.51 to Exercise 9.8.9. This result should be compared with (4.21). As an application of this theorem we give yet another proof of the Cauchy–Davenport inequality, this proof being Fourier-analytic (or more precisely Fourier-algebraic) in nature.

Theorem 9.53 (Cauchy–Davenport inequality, yet again) *Let $F = F_p$ be a finite field of prime order. If A, B are two additive sets in F, then*

$$|A + B| \geq \min(|A| + |B| - 1, p).$$

Proof ([355] and Robin Chapman, private communication) Since A and B are non-empty, we may find two subsets X and Y of $\mathbf{Z}/p\mathbf{Z}$ such that $|X| = p + 1 - |A|$, $|Y| = p + 1 - |B|$, and $|X \cap Y| = \max(|X| + |Y| - p, 1)$. By Theorem 9.52 we may find a function f such that $\mathrm{supp}(f) = A$ and $\mathrm{supp}(\hat{f}) = X$, and a function g such that $\mathrm{supp}(g) = B$ and $\mathrm{supp}(\hat{g}) = Y$. Then $f * g$ has support contained in $A + B$ and has Fourier support equal to $X \cap Y$ (in particular, $f * g$ is non-zero), and hence by Theorem 9.52 again we have $|A + B| + |X \cap Y| \geq p + 1$, which gives $|A + B| \geq \max(|A| + |B| - 1, p)$ as desired. □

One can iterate Theorem 9.52 to also apply to the group $\mathbf{Z}_p^n$ for any $n \geq 1$, which we endow with the standard bilinear form, as in Example 4.2.

Corollary 9.54 *[249] Let p be a prime, $n \geq 1$ be an integer, and $f : \mathbf{Z}_p^n \to \mathbf{C}$ be a non-zero random variable. Then we have*

$$p^k |\mathrm{supp}(f)| + p^{n-k-1} |\mathrm{supp}(\hat{f})| \geq p^n + p^{n-1}$$

for all $0 \leq k \leq n - 1$.

Remark 9.55 These bounds can be seen to be sharp in a large number of situations, by taking the Cartesian product of the examples in Theorem 9.52 with subgroups of $\mathbf{Z}_p$. It has a nice geometric interpretation: if one plots the point $(|\mathrm{supp}(f)|, |\mathrm{supp}(\hat{f})|)$ in $\mathbf{Z} \times \mathbf{Z}$, then this point lies on or above the convex hull of the points (p^j, p^{n-j}) for $0 \leq j \leq n$, which correspond to the cases where f is the indicator function of a subgroup of $\mathbf{Z}_p^n$; this convex hull should be contrasted with the hyperbola corresponding to (4.21). In [249], this result was generalized further to arbitrary finite additive groups Z, see Exercise 9.8.11.

Proof We prove this by induction on n. For $n = 1$ this is just Theorem 9.52. Now suppose that $n > 1$, and the Corollary has already been proven for all smaller values of n. Fix f. We parameterize $\mathbf{Z}_p^n$ as $x = (\underline{x}, x_n)$, where $\underline{x} \in \mathbf{Z}_p^{n-1}$ and $x_n \in \mathbf{Z}_p$. If $g(\underline{\xi}, x_n)$ is the Fourier transform of $f(\underline{x}, x_n)$ in the $\underline{x}$ variable (with x_n fixed), then $\hat{f}(\underline{\xi}, \xi_n)$ is the Fourier transform of $g(\underline{\xi}, x_n)$ in the x_n variable (keeping $\underline{x}$ fixed).

Let $A \subset \mathbf{Z}_p$ be the set of all x_n such that $f(\cdot, x_n)$ (and hence $g(\cdot, x_n)$) is not identically zero. Observe that $1 \leq |A| \leq p$ and

$$|\mathrm{supp}(f)| = \sum_{x_n \in A} |\mathrm{supp}(f(\cdot, x_n))|.$$

Thus by the pigeonhole principle there exists an x_n such that

$$|A||\mathrm{supp}(f(\cdot, x_n))| \leq |\mathrm{supp}(f)|. \tag{9.22}$$

Fix this x_n. By induction we have

$$p^{k'}|\mathrm{supp}(f(\cdot, x_n))| + p^{n-k'-1}|\mathrm{supp}(g(\cdot, x_n))| \geq p^{n-1} + p^{n-2} \tag{9.23}$$

for all $0 \leq k' \leq n-2$. Also, for any $\underline{\xi}$ in the support of $g(\cdot, x_n)$, we see that $g(\underline{\xi}, \cdot)$ is supported in A, so by Theorem 9.52

$$|\mathrm{supp}(\hat{f}(\underline{\xi}, \cdot))| \geq p + 1 - |A|.$$

Summing this over all $\underline{\xi}$ in the support of $g(\cdot, x_n)$ we obtain

$$|\mathrm{supp}(\hat{f})| \geq (p + 1 - |A|)|\mathrm{supp}(g(\cdot, x_n))|.$$

Combining this with (9.22) we obtain

$$p^k|\mathrm{supp}(f)| + p^{n-k-1}|\mathrm{supp}(\hat{f})| \geq p^k|A||\mathrm{supp}(f(\cdot, x_n))| + (p + 1 - |A|)p^{n-k-1}|\mathrm{supp}(g(\cdot, x_n))|.$$

When $|A|$ is equal to 1 or p then the right-hand side here is at least $p^n + p^{n-1}$ thanks to (9.23). Since the right-hand side is linear in $|A|$, the same is true for the intermediate cases $1 < |A| < p$. This completes the induction. □

Exercises

9.8.1 Let p be a prime and $k \geq 1$. Prove that $\Phi_p(z) = 1 + z + z^2 + \cdots + z^{p-1}$ and $\Phi_{p^k}(z) = \Phi_p(z^{p^{k-1}})$.

9.8.2 (Eisenstein's criterion) Let p be a prime, and let $P(t) = a_n t^n + \cdots + a_0 \in \mathbf{Z}[t]$ be such that a_n is not divisible by p, that $a_{n-1}, \ldots, a_0$ are divisible by p, and a_0 is not divisible by p^2. Show that P is irreducible in $\mathbf{Z}[t]$.

9.8.3 Let p be a prime. Compute the polynomial $\Phi_p(t-1)$ explicitly, and then use Eisenstein's criterion to give a proof that $\Phi_p(t-1)$, and hence Φ_p itself, is irreducible in $\mathbf{Z}[t]$, without using Lemma 9.48.

9.8.4 Let $n \geq 1$ be an integer, and suppose that $x \in F_p^\times$ is such that $\Phi_n(x) = 0$. Show that $\mathrm{ord}_\times(x) = n$, and in particular n divides $p-1$.

9.8.5 Let n, m be integers. Using Exercise 9.8.4, show that all the prime factors of $\Phi_n(m)$ are equal to 1 mod n and are coprime to m. Using this (and modifying Euclid's proof of the infinitude of primes) show that there are infinitely many primes equal to 1 mod n; this is a special case of *Dirichlet's theorem*.

9.8.6 Let $n \geq 1$, and let ω be a primitive nth root of unity. Show that the cyclotomic field $\mathbf{Q}(\omega)$ is a $\phi(n)$-dimensional vector space over $\mathbf{Q}$, and

that the complex numbers $1, \omega, \omega^2, \ldots, \omega^{\phi(n)-1}$ form a linear basis for $\mathbf{Q}(\omega)$.

9.8.7 Let p be a prime, and let ω be a primitive pth root of unity. Let $\mathbf{Z}[\omega]$ be the ring generated by ω. Show that the quotient ring $\mathbf{Z}[\omega]/((1-\omega)\cdot\mathbf{Z}[\omega])$ is isomorphic to the field F_p. (Hint: exploit the fact that $\Phi_p(1) = p$, and hence $\Phi_p(\omega) - p$ contains a factor of $(1-\omega)$.)

9.8.8 [120] Let p be a prime, let ω be a primitive pth root of unity, let $z_1, \ldots, z_k$ be distinct pth roots of unity, and let $n_1, \ldots, n_k$ be distinct integers in $[0, p)$. Suppose there exists a polynomial $P \in \mathbf{Z}[\omega][z]$ of degree at most $p-1$ which vanishes at $z_1, \ldots, z_k$ and has at most k non-zero coefficients. Using Exercise 9.8.7 and Lemma 9.26, show that P is a multiple of $(1-\omega)$. Using this and an infinite descent argument, obtain another proof of Lemma 9.51 (at least in the case $q = p$, which is all one needs for Theorem 9.52).

9.8.9 [355] Deduce Theorem 9.52 from Lemma 9.51. (Hint: Lemma 9.51 implies that all the minors of the Fourier matrix $(e^{2\pi ijk/p})_{1\le j,k\le p}$ are invertible.) Conversely, show that Theorem 9.52 implies the $q = p$ case of Lemma 9.51.

9.8.10 Let p be a prime, let $G := \{z \in \mathbf{C} : z^p = 1\}$ be the pth roots of unity, and let $P \in \mathbf{C}[z]$ be a non-zero polynomial with $\deg(P) < p$. Show that that the number of zeroes of P in G cannot exceed the number of non-zero coefficients in P.

9.8.11 [249] Given any finite additive group Z and any real number k, let $\theta(Z;k)$ denote the quantity

$$\theta(Z;k) := \inf\{|\mathrm{supp}(\hat{f})| : f \in L^2(Z); f \not\equiv 0; |\mathrm{supp}(f)| \le k\}.$$

Show that for every subgroup G of Z and any $1 \le k \le |Z|$, we have the inequality

$$\theta(Z;k) \ge \sup_{st=k} \theta(G;s)\theta(Z/G;t)$$

by adapting the proof of Corollary 9.54. Conclude via an inductive argument that for any non-zero function f in $L^2(Z)$, the lattice point $(|\mathrm{supp}(f)|, |\mathrm{supp}(\hat{f})|)$ lies on or above the convex hull of the points $(|G|, |Z|/|G|)$ as G ranges over all subgroups of Z.