TOPICS IN FOURIER ANALYSIS AND APPLICATIONS

MANJUNATH KRISHNAPUR

Contents

CHAPTER 1

Fourier analysis on finite abelian groups

1. Introduction

Let us start with perhaps the most important basic object in physics, the *simple harmonic motion*[1].

A single mass in simple harmonic motion: Consider a body of unit mass connected to a spring whose other end is fixed to a wall. When the spring is at its normal length, the location of the body is designated 0. If it is pulled or pushed and let go, the spring exerts a force proportional to the stretch, towards the mean position. Therefore, the position $x(t)$ undergoes a motion according to Newton's law that says $\frac{d^2}{dt^2}x(t) = -\kappa x(t)$. The general solution to this differential equation is

$$x(t) = a\cos(\sqrt{\kappa}t) + b\sin(\sqrt{\kappa}t).$$

If the initial position $x(0)$ and initial velocity $x'(0)$ are specified, we can solve for the coefficients as $a = x(0)$ and $b\sqrt{\kappa} = x'(0)$.

Finitely many masses in simple harmonic motion: Now consider N bodies of mass m each, connected in a line by springs, with the first and last connected to fixed walls by springs. We assume that all springs and masses are identical, and that when at rest position, the masses are at locations $k/(N+1)$, $1 \le k \le N$. If the bodies are pulled from their rest positions and let go (may be with certain initial velocities), they perform a complicated oscillatory motion influencing each other. To describe the equations, let $x_k(t)$ denote the *displacement* of the kth body from its mean position. We also set $x_0(t) = 0$ and $x_{N+1}(t) = 0$ (the walls are immovable). Then the force kth mass feels a force of $-\kappa(x_k(t) - x_{k-1}(t))$ from the spring to its left, and a force of $\kappa(x_{k+1}(t) - x_k(t))$ from the spring to its right. The total force is therefore $\kappa(x_{k+1}(t) - 2x_k(t) + x_{k-1}(t))$. Hence the equations of motion are

$$m\frac{d^2}{dt^2}x_k(t) = \kappa(x_{k+1}(t) - 2x_k(t) + x_{k-1}(t)), \text{ for } 1 \le k \le N. \tag{1}$$

Unlike the case of one mass, these appear difficult to solve. Let us look for simple solutions of the form $x_k(t) = v(t)w_k$. Then, the equations become (we also set the spring strength $\kappa = 1$ and the

[1] In a course on mathematical methods in physics that I attended in graduate school, Alberto Grunbaum quoted V I Arnold as saying Mathematics $\subseteq$ Physics, and extended it by adding Physics $\subseteq$ SHM!

mass $m = 1$)

$$w_k v''(t) = v(t)(w_{k+1} - 2w_k + w_{k-1})$$

which when rearranged become (what if we are dividing by zero? I leave you to worry about that case separately)

$$\frac{v''(t)}{v(t)} = \frac{w_{k+1} - 2w_k + w_{k-1}}{w_k}.$$

The left side depends on t alone while the right side depends on k alone. Since the two variables can be independently changed, this forces that both sides must be constant. Hence

$$v''(t) = -\lambda v(t) \qquad \text{and} \qquad w_{k+1} - 2w_k + w_{k-1} = -\lambda w_k.$$

We have written the constant as $-\lambda$ anticipating that it will turn out negative. Indeed, the equations for $\mathbf{w} = (w_1, \ldots, w_N)^t$ is the eigenvalue equation $L\mathbf{w} = \lambda w$, where

$$L_{N\times N} = \begin{bmatrix} 2 & -1 & & & & \\ -1 & 2 & -1 & & & \\ & -1 & \ddots & \ddots & & \\ & & \ddots & \ddots & 1 & \\ & & & -1 & 2 & -1 \\ & & & & -1 & 2 \end{bmatrix}.$$

All entries except on the three diagonals are zero. It can be checked (can you?) that L is a positive semi-definite matrix, hence $\lambda \geq 0$. Its eigenvalues and eigenvectors are explicitly found in the following exercise.

Exercise 1. Let $\mathbf{v}_\theta = (\sin\theta, \sin 2\theta, \ldots, \sin N\theta)^t$. Let $\theta_r = \frac{\pi r}{N+1}$ and $\lambda_r = 2 - 2\cos\theta_r = 4\sin^2\frac{\pi r}{2N+2}$ for $1 \leq r \leq N$. Show that $L\mathbf{v}_{\theta_r} = \lambda_r \mathbf{v}_{\theta_r}$. Argue that these are all the eigenvectors and eigenvalues of L.

Returning to the problem of springs, we see that the equations for w can be solved if and only if $\lambda = \lambda_r$ for some $1 \leq r \leq N$, and then $w_k = \sin\frac{\pi r k}{N+1}$ for all k (observe that plugging in $k = 0$ or $k = N+1$ automatically gives 0, ensuring the boundary conditions $x_0(t) = 0 = x_{N+1}(t)$). The general solutions for v is

$$v(t) = a\cos(\sqrt{\lambda_r}t) + b\sin(\sqrt{\lambda_r}t).$$

With this choice of v and w, we have arrived at the solution

$$x_k(t) = (a\cos(\sqrt{\lambda_r}t) + b\sin(\sqrt{\lambda_r}t))\sin\frac{\pi r k}{N+1}, \qquad 0 \leq k \leq N+1.$$

It satisfies the boundary conditions $x_0(t) = 0 = x_{N+1}(t)$ and the initial condition $x_k(0) = aw_k$ and $x_k'(0) = b\sqrt{\lambda_r}w_k$.

Since the problem for x is linear, taking linear combinations such as

$$u_k(t) = \sum_{r=1}^{N} (a_r \cos(\sqrt{\lambda_r} t) + b_r \sin(\sqrt{\lambda_r} t)) \sin \frac{\pi r k}{N+1},$$

we get another solution to the system (1). The boundary conditions $u_0(t) = 0 = u_{N+1}(t)$ are also trivially satisfied. The initial conditions are

$$u(0) = \sum_{r=1}^{N} a_r \mathbf{v}_{\theta_r}, \qquad u'(0) = \sum_{r=1} b_r \sqrt{\lambda_r} \mathbf{v}_{\theta_r}.$$

As v_{θ_r}, $1 \le r \le N$, form a basis for $\mathbb{R}^N$ (being eigenvectors of a symmetric matrix), for any given $u(0)$ and $u'(0)$, we can find a unique set of coefficients a_r, b_r, so that the initial conditions are satisfied.

Exercise 2. Consider N masses placed on a circle (they cannot leave the circle, but they can move on it) and each mass connected by springs to its two adjacent neighbours (there is no wall). Can you work out the equations and the solutions to this system?

A continuum of masses in simple harmonic motion: In the previous setting, suppose that the mass of each body is $1/N$ and the spring strength is N. Then the equations (1) become

$$\frac{d^2}{dt^2} x_k(t) = \frac{x_{k+1}(t) - 2x_k(t) + x_{k-1}(t)}{1/N^2}, \text{ for } 1 \le k \le N.$$

We think of $N \to \infty$. Instead of indexing the bodies by $1, 2, \dots, N$, let us index them by $x = k/N$, $1 \le k \le N$ and write $u(t,x) = x_k(t)$ where $x = k/N$. Then the above equations become

$$\frac{d^2}{dt^2} u(t,x) = \frac{u(t, x + \frac{1}{N}) - 2u(t,x) + u(t, x - \frac{1}{N})}{1/N^2}, \text{ for } 1 \le k \le N \text{ with } x = \frac{k}{N}.$$

As $N \to \infty$, formally we arrive at

$$\frac{d^2}{dt^2} u(t,x) = \frac{d^2}{dx^2} u(t,x).$$

This is called the *wave equation* and describes a vibrating string. The boundary conditions $u(t,0) = 0 = u(t,1)$ describes a string with end-points fixed. The initial conditions are $u(0,x) = f(x)$ and $\frac{d}{dt} u(0,x) = g(x)$.

Can we solve this problem? Again it looks difficult, but taking a cue from the case of finitely many masses, we attempt a solution of the form $u(t,x) = v(t)w(x)$ and arrive at the equations

$$\frac{v''(t)}{v(t)} = \frac{w''(x)}{w(x)}$$

which forces that both sides must be some constant $-\lambda$. Thus we get

$$v''(t) = -\lambda v(t), \qquad w''(x) = -\lambda w(x).$$

General solutions to these ODEs are, assuming $\lambda > 0$,

$$v(t) = a \cos \sqrt{\lambda} t + b \sin \sqrt{\lambda} t, \qquad w(x) = c \cos \sqrt{\lambda} x + d \sin \sqrt{\lambda} x.$$

The boundary conditions require $w(0) = 0 = w(1)$, hence $c = 0$ and $\sqrt{\lambda} = n\pi$ for some $n \in \mathbb{Z}$. We take $d = 1$ and $n \geq 1$, because otherwise we just get a linear multiple of the solutions with $d = 1$ and $-n$. The initial conditions for the resulting solution are

$$u(0, x) = a \sin \pi n x, \qquad \frac{d}{dt} u(0, x) = b \pi n \sin \pi n x,$$

At all stages, observe the similarity to the case of finitely many masses.

As before, the wave equation is linear in the initial conditions f, g, meaning that if u_i solves the equation with f_i, g_i, for $i = 1, 2$, then $\alpha u_1 + \beta u_2$ satisfies it with $\alpha f_1 + \beta f_2, \alpha g_1 + \beta g_2$. Therefore, *formally* (i.e., without paying attention to what convergence means etc.)

$$u(t, x) := \sum_{n=1}^{\infty} (a_n \cos(\pi n t) + b_n \sin(\pi n t)) \sin(\pi n x)$$

satisfies the equation with initial conditions

$$u(0, x) = \sum_{n=1}^{\infty} a_n \sin(\pi n x), \qquad \frac{d}{dt} u(0, x) = \sum_{n=1}^{\infty} \pi n b_n \sin(\pi n x)$$

Therefore, to solve the problem for given f, g, the question becomes: Can we find a_n, b_n so that

$$f(x) = \sum_{n=1}^{\infty} a_n \sin(\pi n x), \qquad g(x) = \sum_{n=1}^{\infty} \pi n b_n \sin(\pi n x).$$

Apart from the convergence issues and what kind of convergence we need, here we are in an infinite dimensional setting. The functions $x \mapsto \sin(\pi n x)$ are elements in a function space, and we want to know what their span is. In particular, is it true that any smooth function f (satisfying $f(0) = 0 = f(1)$) is in the span? If the answer is yes, then we can presumably solve the vibrating string problem for smooth initial conditions f, g. Smoothness may be too restrictive - for example, when a string is plucked at the mid-point and let go, then we may model $f(x) = c(\frac{1}{2} - |\frac{1}{2} - x|)$ and $g(x) = 0$. Hence one may want to ask the question for continuous functions, or some other class of functions. This was the starting point of Fourier analysis - can a function be written as a linear combination of sines (and cosines)?

Remark 3. It is a running theme of this course that the discrete situation is as interesting and useful as the continuous situation. Historically the continuous objects were defined first, but in recent years the discrete objects are found to be very useful in a variety of fields. That is the reason why we elaborated on the case of finitely many masses. Those who prefer an even more discrete setting, where time is also discretized, may go over the next example.

Random walk on a discrete cycle: Consider a particle that is moving on the discrete cycle $\{0, 1, \dots, N-1\}$. At each discrete time point $t = 0, 1, 2, \dots$, a coin is tossed, and if it falls head, the particles moves one step up (modulo N), and if the coin falls tails, it moves one step down (modulo N). If the starting position at time 0 is 0, what is the probability distribution after t steps?

Let $p_t(k)$ be this probability, for $t \geq 0$ and $k \in \{0, 1, \dots, N-1\}$. Then,

$$p_{t+1}(k) = \frac{1}{2}p_t(k-1) + \frac{1}{2}p_t(k+1). \tag{2}$$

Can we solve for this? There are different approaches, but we take a route that is close to the situation considered earlier. Subtract $p_t(k)$ from the above equation to get

$$p_{t+1}(k) - p_t(k) = \frac{1}{2}[p_t(k+1) - 2p_t(k) + p_t(k-1)].$$

On the left side we have the first difference in t, while on the right, we have the second difference in k. The analogous continuous equation is $\frac{d}{dt}p_t(x) = \frac{1}{2}\frac{d^2}{dx^2}p_t(x)$, which is not the wave equation, but the *heat equation*. Nevertheless, the same ideas may be repeated.

First we attempt a solution of the form $v(t)w(k)$. The equations become (everywhere $k \pm 1$ to be interpreted modulo N)

$$\frac{v(k+1) - v(k)}{v(k)} = \frac{1}{2}\frac{w(k+1) - 2w(k) + w(k-1)}{w(k)}.$$

Both sides must be constant, say $-\lambda$. The equations for w are $-w(k+1) + 2w(k) - w(k-1) = 2\lambda w(k)$, which is the eigenvalue equation $Lw = 2\lambda w$, with (all entries not shown are zero)

$$L_{N \times N} = \begin{bmatrix} 2 & -1 & & & & -1 \\ -1 & 2 & -1 & & & \\ & -1 & \ddots & \ddots & & \\ & & \ddots & \ddots & 1 & \\ & & & -1 & 2 & -1 \\ -1 & & & & -1 & 2 \end{bmatrix}.$$

Observe the subtle difference from the earlier matrix. Let $\mathbf{v}_\theta = (1, e^{i\theta}, \dots, e^{i(N-1)\theta})^t$. Then it is easy to see that if $e^{iN\theta} = 1$, then $L\mathbf{v}_\theta = 2(1 - \cos\theta)\mathbf{v}_\theta$. Hence the eigenvectors are $\mathbf{v}_{\theta_r}$ with eigenvalues $2(1 - \cos\theta_r)$, where $0 \leq r \leq N-1$ and $\theta_r = \frac{2\pi r}{N}$. Therefore, choices for w are $\mathbf{v}_{\theta_r}$, $0 \leq r \leq N-1$, and then $\lambda = 1 - \cos\theta_r$. The equation for v becomes $v_{k+1} = v_k(1 - \lambda) = v_k \cos\theta_r$, which means that $v_t = v_1(\cos\theta_r)^t$. Thus, we have arrived at the solutions

$$u_r(t, k) := (\cos\theta_r)^t e^{ik\theta_r}, \quad 0 \leq k \leq N-1, \;\; 0 \leq r \leq N-1.$$

The initial condition is $u_r(0, k) = e^{ik\theta_r}$ (i.e., $u_r(0, \cdot) = \mathbf{v}_{\theta_r}$). Taking linear combinations of u_r, we can get solutions with general initial conditions.

In particular, for the original problem of $p_t(k)$, observe that $p_0(k) = \delta_0(k)$, since the random walk starts at 0. It is easy to see that $\delta_0 = \frac{1}{N}\sum_{r=0}^{N-1} \mathbf{v}_{\theta_r}$, hence

$$\begin{aligned}
p_t(k) &= \frac{1}{N}\sum_{r=0}^{N-1} u_r(t,k) \\
&= \frac{1}{N2^t}\sum_{r=0}^{N-1}(e^{i\theta_r} + e^{-i\theta_r})^t e^{ik\theta_r} \\
&= \frac{1}{N2^t}\sum_{r=0}^{N-1}\sum_{j=0}^{t}\binom{t}{j} e^{i\theta_r(t-2j+k)} \\
&= \frac{1}{2^t}\sum_{0\le j\le t:\ t-2j+k=0 \pmod N}\binom{t}{j}.
\end{aligned}$$

If it was not clear earlier, it should be clear now that this answer could also have been arrived at by combinatorial methods, but that is not the point of our discussion here.

Exercise 4. The continuum analogue of (2) is $\frac{\partial}{\partial t}p(t,x) = \frac{\partial^2}{\partial x^2}p(t,x)$, with $(t,x) \in [0,\infty)\times\mathbb{T}$, where $\mathbb{T} = [0,1]/$ under the equivalence $0 \sim 1$ (so $\mathbb{T}$ is the circle). Use separation of variables to find many solutions and formulate the question on Fourier series regarding how one could find solutions for general initial conditions.

2. The groups of interest

As we saw in the context of solving the vibrating string problem, the basic question of Fourier series is about writing 2π-periodic functions on $\mathbb{R}$ (equivalently thought of as functions on the unit circle $\mathbb{T} = \{e^{it} : 0 \le t < 2\pi\}$) as linear combinations of sines and cosines with integer frequencies, i.e., as $a_0 + \sum_{n\ge1} a_n\cos(nx) + b_n\sin(nx)$. Fourier transform, which you may also have encountered as characteristic functions in probability class, involves writing functions on $\mathbb{R}$ as superpositions of sines and cosines with arbitrary frequencies, i.e., as $\int_{-\infty}^{\infty}[g(\lambda)\cos(\lambda x) + h(\lambda)\sin(\lambda x)]dx$.

In analysis it is almost always better to work over complex numbers, hence, in place of sines and cosines one may use complex exponentials $e_\lambda(x) := e^{i\lambda x}$ where $\lambda \in \mathbb{Z}$ for Fourier series and $\lambda \in \mathbb{R}$ for Fourier transform. One can recover $\sin(\lambda x)$ and $\cos(\lambda x)$ from $e_\lambda(x)$ and $e_{-\lambda}(x)$, and vice versa. The spaces $\mathbb{T}^1$ and $\mathbb{R}$, are groups under multiplication and addition. The sets of continuous homomorphisms from these groups into the circle group $\mathbb{T}$ are precisely $\{e_n : n \in \mathbb{Z}\}$ and $\{e_\lambda : \lambda \in \mathbb{R}\}$.

This suggests the generalization to other groups, and asking if continuous homomorphisms (the group must have a topology to talk about continuity) from the group into S^1 give a good collection of functions whose linear superpositions give a large class of functions from the group into $\mathbb{C}$. The matters get more subtle in non-commutative groups, but if we restrict to Abelian groups (with a locally compact topology), then the whole story is clean and complete. We shall indicate this general situation later, but for the purposes of this course, we are only concerned with the following four examples. Henceforth we shall use the standard terms "character" for continuous homomorphisms into $\mathbb{C}$ and $\hat{G}$ for the set of all characters of a group G.

(1) The group $\mathbb{R}$ under addition. If $\chi : \mathbb{R} \mapsto \mathbb{T}$, then $\chi(nx) = \chi(x)^n$ for $x \in \mathbb{R}$ and $n \in \mathbb{N}$, from which it follows that $\chi(x) = \chi(1)^x$ for $x \in \mathbb{Q}$. By continuity, the same holds for all $x \in \mathbb{R}^d$, showing that $\chi(x) = e^{i\lambda x}$ where $\chi(1) = e^{i\lambda}$. Thus, $\hat{\mathbb{R}} = \{e_\lambda : \lambda \in \mathbb{R}\}$.

(2) The 1-dimensional torus or circle grup $\mathbb{T} = \{e^{it} : 0 \leq t < 2\pi\}$ with with multiplication and the usual topology derived from embedding in $\mathbb{C}$ (we can take the metric $d(e^{it}, e^{is}) = |e^{it} - e^{is}|$). Since $\varphi : \mathbb{R} \mapsto \mathbb{T}$ define by $\varphi(x) = e^{ix}$ is a continuous homomorphism, if $\chi : \mathbb{T} \mapsto \mathbb{T}$ is a continuous homomorphism then $\chi \circ \varphi$ is a homomorphism from $\mathbb{R}$ to $\mathbb{T}$, showing that $\chi \circ \varphi = e_\lambda$ for some $\lambda \in \mathbb{R}$. As $\varphi(x + 2\pi) = \varphi(x)$, the same must hold for e_λ and that forces $\lambda \in \mathbb{Z}$. Thus, $\hat{\mathbb{T}} = \{e_m : m \in \mathbb{Z}\}$.

(3) The finite cyclic group $\mathbb{Z}_n = \mathbb{Z}_n = \{0, 1, \ldots n-1\}$ with addition modulo n. On finite or countable groups, we always take the discrete topology (in other words, the metric is $d(k, \ell) = 1$ if $k \neq \ell$), hence continuity is not a restriction. A homomorphism χ is determined by $\chi(0)$. Also $\chi(0)^n = 1$, hence $\chi(0) = e^{2\pi i k/n}$ for some $0 \leq k \leq n-1$. Therefore, $\widehat{\mathbb{Z}_n} = \{e_{2\pi k/n} : 0 \leq k \leq n-1\}$.

(4) The group $\mathbb{Z}_2^n$ which we represent as $\{0, 1\}^n$ with coordinatewise addition modulo 2 or alternately as $\{-1, +1\}^n$ with coordinatewise multiplication. For each $S \subseteq [n]$, the function $\chi_S : \{-1, +1\}^n \mapsto \mathbb{T}$ defined by $\chi_S(x) = \prod_{i \in S} x_i$, is a homomorphism. By completing the exercise below or directly, show that these are all the homomorphisms from $\mathbb{Z}_2^n$ into $\mathbb{T}$. Thus, $\widehat{\mathbb{Z}_2^n} = \{\chi_S : S \subseteq [n]\}$.

Exercise 1. If G is a finite group and χ is a character, show that it is in fact a homomorphism into the unit circle $\mathbb{T}$ (which is a group under multiplication). Going further, it is a homomorphism into the subgroup of nth roots of unity for $n = |G|$.

Exercise 2. If G_1, G_2 are groups with topology, then $G_1 \times G_2$ is a group with the product topology and co-ordinatewise multiplication. Show that $\chi : G_1 \times G_2 \mapsto \mathbb{T}$ is a character if and only if it is of the form $\chi_1 \otimes \chi_2$, where χ_1 and χ_2 are characters of G_1 and G_2 respectively. (Recall the tensor product notation: $(f \otimes g)(x, y) := f(x)g(y)$)

Exercise 3. Show that the characters of

(1) $\mathbb{R}^n$ are precisely e_λ, $\lambda \in \mathbb{R}^n$, where $e_\lambda(x) := e^{2\pi i \langle \lambda, x \rangle}$,

(2) $\mathbb{T}^n$ are precisely e_λ, $\lambda \in \mathbb{Z}^n$,

(3) $\mathbb{Z}_2^n$ are precisely χ_S, $S \subseteq [n]$, that were introduced above.

3. Fourier analysis on $\mathbb{Z}_n$

For any finite group G with $|G| = n$, we denote by $L^2(G)$ the n-dimensional complex vector space of functions $f : G \mapsto \mathbb{C}$. We can and do identify f with the column vector $(f(0), \dots, f(n-1))^t \in \mathbb{C}^n$. The inner product on $\mathbb{C}^n$ gives the inner product on $L^2(G)$:

$$\langle f, g\rangle = \sum_{x\in G} f(x)\overline{g(x)}$$

Now let $G = \mathbb{Z}_n$. The characters are $\chi_k(j) = e^{2\pi ijk/n}$, $0 \le k, j \le n-1$ of $\mathbb{Z}_n$. Of these $\chi_0 = 1$ is the trivial character. The single most important point about these characters is the orthogonality relationship

$$\langle \chi_k, \chi_\ell\rangle = \sum_{j=0}^{n-1} e^{2\pi ij(k-\ell)/n} = \begin{cases} n & \text{if } k = \ell \\ 0 & \text{if } k \neq \ell. \end{cases}$$

Indeed, as j runs over $\mathbb{Z}_n$, if $k \neq \ell$, then the summand runs over all mth roots of unity for some $m \ge 2$. When $k = \ell$, all the summands are identically equal to 1.

Thus, $\{\frac{1}{\sqrt{n}}\chi_k\}_{0\le k\le n-1}$ is orthonormal, and since there are n of them, they form an orthonormal basis for $L^2(\mathbb{Z}_n)$. Hence, any $f \in L^2(\mathbb{Z}_n)$ can be written as

$$f = \frac{1}{\sqrt{n}}\sum_{k=0}^{n-1} \hat{f}(k)\chi_k, \qquad \hat{f}(k) := \frac{1}{\sqrt{n}}\langle f, \chi_k\rangle = \frac{1}{\sqrt{n}}\sum_{j=0}^{n-1} f(j)e^{-2\pi ijk/n}.$$

Functions on $\mathbb{Z}_n$ can be identified in a natural way with n-periodic functions (satisfying $f(x+n) = f(x)$) on $\mathbb{Z}$ via the composition map $\mathbb{Z} \mapsto \mathbb{Z}_n$ given by $k \mapsto k \pmod{n}$. Thus, χ_k can also be thought of as functions on $\mathbb{Z}$, the formula is exactly the same $\chi_k(j) = e^{2\pi ijk/n}$, since this is already n-periodic. The function $\hat{f} : \{0, 1, \dots, n-1\} \mapsto \mathbb{C}$ is called the discrete Fourier transform of f.

From the orthogonality of characters, we get the *Plancherel relation*

$$\langle f, g\rangle = \sum_{k=0}^{n-1} \hat{f}(k)\overline{\hat{g}(k)}, \qquad \|f\|^2 = \sum_{k=0}^{n-1} |\hat{f}(k)|^2.$$

Observe that the relationship between f and $\hat{f}$ may also be written as

$$f(j) = \frac{1}{\sqrt{n}}\sum_{j=0}^{n-1} \hat{f}(k)e^{2\pi ijk/n}, \qquad \hat{f}(k) = \frac{1}{\sqrt{n}}\sum_{j=0}^{n-1} f(j)e^{-2\pi ijk/n}.$$

The two relationships are almost identical, except for the negative sign in the exponent. In other symbols, $\hat{\hat{f}}(j) = f(-j)$. This is called *Fourier inversion*. That is, if we define a new transformation (pronounced "g check")

$$\check{g}(k) = \frac{1}{\sqrt{n}}\sum_{j=0}^{n-1} g(j)e^{-2\pi ijk/n}$$

then it is the inverse of the Fourier transform.

Some examples.

Example 4. On $\mathbb{Z}_n$, let $f = \mathbf{1}$ and $g = \mathbf{1}_0$ (two extremes in terms of support size). Then

$$\hat{f}(k) = \begin{cases} \sqrt{n} & \text{if } k = 0, \\ 0 & \text{otherwise.} \end{cases} \qquad \hat{g}(k) = \frac{1}{\sqrt{n}} \text{ for any } k.$$

Observe that $\hat{f} \propto g$ and $\hat{g} \propto f$.

Example 5. Let $n = k\ell$ and let $f(x) = \mathbf{1}_{x=0 \pmod k}$. Then

$$\hat{f}(r) = \frac{1}{\sqrt{n}} \sum_{j=0}^{\ell-1} e^{2\pi i jkr/n} = \frac{1}{\sqrt{n}} \sum_{j=0}^{\ell-1} e^{2\pi i jr/\ell} = \begin{cases} \frac{\ell}{\sqrt{n}} & \text{if } r = 0 \pmod \ell) \\ 0 & \text{otherwise.} \end{cases}$$

Thus the support of f has size ℓ and the support of $\hat{f}$ has size k. For fixed n and different choices of (k, ℓ), the support sizes have an inverse relationship. This is an illustration of *uncertainty principle*, more of which we shall see later.

Exercise 6. Write the functions $n \mapsto \mathbf{1}_{n=1 \pmod 3}$ and $n \mapsto \mathbf{1}_{n=3 \pmod 7}$ on $\mathbb{Z}$ in terms of the characters of $\mathbb{Z}_3$ and $\mathbb{Z}_7$ respectively (as remarked above, we extend these functions to all of $\mathbb{Z}$).

The indicator functions capturing modular properties can thus be expressed in terms of the much nicer functions, namely the characters. This simple observation will be crucial to applications later.

The discrete Fourier matrix: Fourier analysis on $\mathbb{Z}_n$ can be numerically executed quite efficiently using matrices. For this, identify $L^2(G)$ with $\mathbb{C}^n$ as before. Define the discrete Fourier matrix $\mathcal{F}_n := \frac{1}{\sqrt{n}}(e^{2\pi i jk/n})_{0 \le j,k \le n-1}$, whose columns represent the characters. The orthogonality of characters is the same as saying $\mathcal{F}_n^* \mathcal{F}_n = I_n$, i.e., that $\mathcal{F}_n$ is a unitary matrix. This also implies that $\mathcal{F}_n \mathcal{F}_n^* = I_n$, which gives a different set of orthogonality relations (observe that the character is fixed here)

$$\frac{1}{n} \sum_{k=0}^{n-1} \chi_k(j) \overline{\chi_k(j')} = \begin{cases} 1 & \text{if } j = j' \\ 0 & \text{if } j \neq j'. \end{cases}$$

The relationship between a function $f = (f(0), \ldots, f(n-1))^t$ and its Fourier transform $\hat{f} = (\hat{f}(0), \ldots, \hat{f}(n-1))^t$ is given by

$$\hat{f} = \mathcal{F}_n^* f \qquad \text{and} \qquad f = \mathcal{F}_n \hat{f}.$$

For $g \in \mathbb{C}^n$, it is customary to denote $\check{g} = \mathcal{F}_n g$. Then, $f \mapsto \hat{f}$ and $g \mapsto \check{g}$ are inverses of each other. But they also look very similar, a fact made more precise in the following exercise.

Exercise 7. Show that $\hat{\hat{f}}(x) = f(-x)$. Conclude that $\mathcal{F}_n^4 = I_n$. What are the possible eigenvalues of $\mathcal{F}_n$? Find the actual eigenvalues (with multiplicities) of $\mathcal{F}_n$ (use a computer to compute for small n to guess the answer).

Exercise 8. Let $n = 2m + 1$ and let $p \leq m$. Let $f(j) = 1$ if $j \wedge (n - j) \leq p$ and 0 otherwise. Find $\hat{f}$.

4. Fourier analysis on $\mathbb{Z}_2^n$

Since we already understand the character theory of $\mathbb{Z}_2$, and most considerations easily carry over to products of groups (see Exercise 2), we can deduce the statements below from those of the previous section. But we give a direct presentation anyway.

Again $L^2(G)$, with $G = \mathbb{Z}_2^n$ is the vector space of $\mathbb{C}$-valued functions on G. The dimension of this space is 2^n. It has the natural inner product

$$\langle f, g \rangle = \sum_{x \in \{-1,+1\}^n} f(x)\overline{g(x)}.$$

The scaled characters $2^{-n/2}\chi_S$, $S \subseteq [n]$, are orthonormal. Indeed,

$$\begin{aligned}\langle \chi_S, \chi_T \rangle &= \sum_{x \in \{-1,+1\}^n} \prod_{i \in S} x_i \prod_{j \in T} x_j \\ &= \sum_{x \in \{-1,+1\}^n} \prod_{i \in S \Delta T} x_i.\end{aligned}$$

This summation factors over $i \in [n]$. The factor corresponding to i gives 2 if $i \notin S\Delta T$ and gives $-1 + 1 = 0$ if $i \in S\Delta T$. Therefore, $\langle \chi_S, \chi_T \rangle = 2^n$ if $S = T$ and $\langle \chi_S, \chi_T \rangle = 0$ otherwise.

The number of subsets of $[n]$ is exactly 2^n, hence the number of elements in the orthonormal set matches the dimension of $L^2(G)$, showing that it is in fact an orthonormal basis. Thus, any $f \in L^2(G)$ may be written as (as the characters are real-valued, we drop the conjugates)

$$f(x) = \frac{1}{2^{n/2}} \sum_{S \subseteq [n]} \hat{f}(S)\chi_S(x), \quad \text{where}$$

$$\hat{f}(S) = \frac{1}{2^{n/2}} \langle f, \chi_S \rangle = \frac{1}{2^{n/2}} \sum_{x \in \{-1,+1\}^n} f(x) \prod_{i \in S} x_i.$$

The Plancherel relations $\langle f, g \rangle = \langle \hat{f}, \hat{g} \rangle$ take the form

$$\sum_{x \in \{-1,1\}^n} f(x)\overline{g(x)} = \sum_{S \subseteq [n]} \hat{f}(S)\overline{\hat{g}(S)}, \qquad \sum_{x \in \{-1,1\}^n} |f(x)|^2 = \sum_{S \subseteq [n]} |\hat{f}(S)|^2$$

It may be useful to see some examples. Some of the picturesque names are motivated by thinking of n voters voting for one of two candidates ± 1, and x_i as the vote of the ith person. A function $f : \{-1, 1\}^n \mapsto \mathbb{R}$ (particularly $f : \{-1, 1\}^n \mapsto \{-1, 1\}$) is the method by which these votes are combined to make a decision.

Example 9. Let $f(x) = x_1$ (the dictator function: whatever the first voter says goes). Then $\hat{f}(S) = 2^{n/2}$ if $S = \{1\}$ and 0 otherwise. Observe that Plancherel relation holds with $\|f\|^2 = \|\hat{f}\|^2 = 2^n$, but the left side sum is "spread out", with $|f(x)|^2 = 1$ for all x while the right side sum is "concentrated" with only one term being non-zero. This is again an example of an uncertainty principle.

Another feature that generalizes from the above example: If f does not depend on a variable j, then $\hat{f}(S) = 0$ for any set that contains j. Taking $j = n$ for convenience and $S = T \sqcup \{n\}$,

$$\hat{f}(S) = \sum_{x} f(x) \prod_{i \in S} x_i = \sum_{x_1, \ldots, x_{n-1}} f(x) \prod_{i \in T} x_i \sum_{x_n \in \{-1,1\}} x_n$$

and the inner sum is zero.

Example 10. Let $f(x) = \mathbf{1}_{x=-1} = 2^{-n} \prod_i (1 + x_i)$. Expanding, we see that $\hat{f}(S) = (-1)^{|S|} 2^{-n/2}$ for all S. In this case, f is concentrated while $\hat{f}$ is spread out.

The exercise below gives an alternate argument that any function on the hypercube can be written as a linear combination of characters. The group structure and inner product are not used, what is used is that $x_i^k = 1$ if k is even and $x_i^k = x_i$ if k is odd (hence any polynomial of x_is may be written as a multilinear polynomial).

Exercise 11. For $a \in \{-1, 1\}^n$, write $\mathbf{1}_a$ as a linear combination of χ_S. Hence argue that any $f : \{-1, 1\}^n \mapsto \mathbb{R}$ can be written as a linear combination of χ_S, $S \subseteq [n]$.

A probabilistic interpretation: Observe that

$$\hat{f}(\emptyset) = 2^{-n/2} \sum_{x \in \mathbb{Z}_2^n} f(x), \qquad \sum_{S} |\hat{f}(S)|^2 = \sum_{x} |f(x)|^2.$$

Hence, if we endow $\mathbb{Z}_2^n$ with uniform probability measure ($\mathbf{P}\{x\} = 2^{-n}$ for all x) and view f as a random variable on it, then $2^{-n/2}\hat{f}(\emptyset)$ is the mean value of f while $2^{-n} \sum_{S \neq \emptyset} |\hat{f}(S)|^2$ is the variance.

5. Fourier analysis on finite abelian groups

Consider a finite abelian group G. Will the characters of G form a basis for $L^2(G)$? Are they orthogonal?

The answer to the second question is yes. Let $\chi : G \mapsto \mathbb{T}$ be any character. Then for any $a \in G$,

$$\langle \chi, \mathbf{1} \rangle = \sum_{x \in G} \chi(x) = \sum_{x \in G} \chi(ax) = \chi(a) \sum_{x \in G} \chi(x) = \chi(a) \langle \chi, \mathbf{1} \rangle.$$

Thus, either χ is trivial (i.e., $\chi(a) = 1$ for all $a \in G$) or $\chi \perp \mathbf{1}$. Now, if χ_1, χ_2 are two characters, then so is $\chi = \chi_1 \overline{\chi_2}$, and χ is trivial if and only if $\chi_1 = \chi_2$. Hence

$$\langle \chi_1, \chi_2 \rangle = \langle \chi, \mathbf{1} \rangle = \begin{cases} |G| & \text{if } \chi_1 = \chi_2, \\ 0 & \text{if } \chi_1 \neq \chi_2. \end{cases}$$

Thus any two distinct characters are orthogonal. In particular there can be at most $N = |G|$ of them (as $L^2(G)$ has dimension N). Observe that we did not use the fact that G is abelian: orthogonality of characters is true for any finite group.

To return to the first question of whether they provide a basis for $L^2(G)$, it is now clear that for that to happen, we need exactly N distinct characters. But if χ is a character, then $\chi(xyx^{-1}) = \chi(y)$ for any $x, y \in G$, implying that χ is constant on each conjugacy class. But G has N conjugacy classes if and only if it is abelian. Hence, whenever G is non-abelian, characters do not span $L^(G)$.

We do not even necessarily have as many characters as conjugacy classes, as the example of symmetric groups shows.

Exercise 12. Show that S_n has only two characters, the trivial one and the sign.

Exercise 13. If G is a finite simple group that has a non-trivial character, then it must be $\mathbb{Z}_p$ for some prime p.

Returning to our question of whether characters provide a basis for $L^2(G)$, it is now clear that we must restrict ourselves to abelian groups. We have seen that for cyclic groups $\mathbb{Z}_n$, there is indeed a basis of characters. Now suppose G_1, G_2 are two groups whose characters form a basis of the corresponding L^2 spaces. From Exercise 2, it follows that $\chi_1 \otimes \chi_2$ is a character of G if χ_i is a character of G_i, for $i = 1, 2$. Since $|G| = |G_1| \times |G_2|$, it seems that this gives $|G|$ characters, and hence there is a full basis of character of $L^2(G)$. But one must check that these character are *distinct*. They are distinct, because χ_1 and χ_2 can be recovered from χ by $\chi_1(x) = \chi(x, 1)$ and $\chi_2(y) = \chi(1, y)$. By a somewhat tedious check one can also see that they are linearly independent, but it is easier to check that they are orthogonal.

Exercise 14. Let $G = G_1 \times G_2$ be a product of finite groups. Show that

$$\langle f_1 \otimes f_2, g_1 \otimes g_2 \rangle_{L^2(G)} = \langle f_1, g_1 \rangle_{L^2(G_1)} \langle f_2, g_2 \rangle_{L^2(G_2)}.$$

Generalize to a product of k groups.

As a consequence, we see that all finite products of cyclic groups have the desired number of characters, and that they form an orthogonal set in L^2. What groups can be realized as products of cyclic groups? As products of abelian groups are abelian, it is clear that we can only get abelian groups this way. In fact we can get all of them!

Structure theorem for finite abelian groups: Let G be a finite abelian group. The G is isomorphic to $\mathbb{Z}_{n_1} \times \ldots \times \mathbb{Z}_{n_k}$ for some $n_i = p_i^{m_i}$, where p_i is prime and $m_i \geq 1$. Of course the order of G must be $N = p_1^{m_1} \ldots p_k^{m_k}$ (but this is not necessarily the prime factorization of $|G|$, note that p_i need not be distinct).

In conclusion (recall the $\hat{G}$ is the set of characters of G),

Theorem 15. *If G is a finite abelian group of order N, then $|\hat{G}| = N$ and the collection $\{\frac{1}{\sqrt{N}}\chi : \chi \in \hat{G}\}$ forms an orthonormal basis for $L^2(G)$.*

Fourier transform: Let G be a finite abelian group of order N. For $f \in L^2(G)$, its *Fourier transform* is defined as $\hat{f} : \hat{G} \mapsto \mathbb{C}$ given by $\hat{f}(\chi) = \frac{1}{\sqrt{N}}\langle f, \chi\rangle$. Thus, for any $x \in G$,

$$f(x) = \frac{1}{\sqrt{N}} \sum_{\chi \in \hat{G}} \hat{f}(\chi)\chi(x).$$

As $\hat{G}$ is a set of the same cardinality, we can talk of $L^2(\hat{G}) = \{f : \hat{G} \mapsto \mathbb{C}\}$, a Hilbert space with inner product $\langle f, g\rangle_{L^2(\hat{G})} = \sum_{\chi \in \hat{G}} f(\chi)\overline{g(\chi)}$.

The Fourier transform is a mapping from $L^2(G)$ to $L^2(\hat{G})$. When the hat notation is not convenient, we denote this mapping by $\mathcal{F}$ or $\mathcal{F}_G$. Often it is a better to view properties of the Fourier transform $\mathcal{F}$ instead of writing it elaborately in terms of f and $\hat{f}$. Immediately from the definition, we get the all important *Plancherel relation*.

Theorem 16 (Plancherel relation)**.** *$\mathcal{F} : L^2(G) \mapsto L^2(\hat{G})$ is a unitary transformation. That is, $\langle f, g\rangle_{L^2(G)} = \langle \hat{f}, \hat{g}\rangle_{L^2(\hat{G})}$. In particular, $\|f\|^2_{L^2(G)} = \|\hat{f}\|^2_{L^2(\hat{G})}$.*

Proof. Write $f \in L^2(G)$ as $\frac{1}{\sqrt{N}} \sum_{\chi \in \hat{G}} \hat{f}(\chi)\chi$ and recall that $\chi/\sqrt{N}$ form an orthonormal basis to see that $\|f\|^2_{L^2(G)} = \sum_{\chi \in \hat{G}} |\hat{f}(\chi)|^2$. Similarly, expand g and take inner product between f and g to get $\langle f, g\rangle_{L^2(G)} = \langle \hat{f}, \hat{g}\rangle_{L^2(\hat{G})}$. This is the definition of unitarity. ■

Remark 17. Observe that in case of $\mathbb{Z}_n$, the characters were naturally indexed by $\mathbb{Z}_n$ too, hence we wrote $\hat{f}(k)$ rather than $\hat{f}(e_{2\pi k/n})$. That also made things like the inversion formula seem natural. But to discuss inversion formula for a general group, we need to first see that $\hat{G}$ itself has a group structure. We do this in a later section on duality.

Probabilistic interpretation: If $f \in L^2(G)$, and $\mathbf{1}$ is the trivial character, then $\hat{f}(\mathbf{1}) = \sum_{x \in G} f(x)$ and by the Plancherel relationship we know that $\sum_{\chi \in \hat{G}} |\hat{f}(\chi)|^2 = \sum_{x \in G} |f(x)|^2$. Hence, if we consider the uniform probability measure on G: $\mathbf{P}\{x\} = \frac{1}{N}$ for each $x \in G$, then the random variable f has mean $\frac{1}{N}\hat{f}(\mathbf{1})$ and variance $\frac{1}{N}\sum_{\chi \neq 1} |\hat{f}(\chi)|^2$.

6. Characters as eigenvectors

Let G be a finite abelian group. For $x \in G$, define $T_x : L^2(G) \mapsto L^2(G)$ by $(T_x f)(y) = f(y + x)$. Then T_x is a linear transformation (we call it translation by x). If χ is any character of G, then for all $x, y \in G$

$$(T_x\chi)(y) = \chi(x + y) = \chi(x)\chi(y).$$

This means that χ is an eigenvector of T_x with eigenvalue $\chi(x)$, for each $x \in G$. As we have a full basis of characters for $L^2(G)$, we see that they simultaneously diagonalize all the translation operators.

In fact we could have take this route to proving the existence of a full basis of characters[2] avoiding the use of the structure theorem for finite abelian groups. In this approach, we start with two observations:

(1) $T_x T_y = T_y T_x = T_{x+y}$ for any $x, y \in G$. This happens because the group is abelian.

(2) $(T_x)^* = T_{-x}$ since for any $f, g \in L^2(G)$,

$$\langle T_x f, g\rangle = \sum_{y\in G}(T_x f)(y)\overline{g(y)} = \sum_{y\in G} f(y+x)\overline{g(y)} = \sum_{z\in G} f(z)\overline{g(z-x)} = \langle f, T_{-x}g\rangle.$$

Consequently, $T_x T_x^* = T_x T_{-x} = T_0 = I$. That is, T_x is unitary.

An extension of the spectral theorem says that a commuting family of normal transformations (recall that T is normal if $T^*T = TT^*$) can be simultaneously diagonalized. Therefore, there is an orthonormal basis of $L^2(G)$, whose elements are eigenvectors for all T_x, $x \in G$.

Now suppose τ is a common eigenvector. Then $T_x\tau = \lambda_x \tau$ for all $x \in G$, for some $\lambda_x \in \mathbb{C}$ (in fact $|\lambda_x| = 1$ by unitarity of T_x). What this means is that $\tau(x+y) = \lambda_x\tau(y)$ for all $x, y \in G$. Since τ is not identically zero, this shows that $\tau(y) \neq 0$ for all y. Normalize the eigenvector so that $\tau(0) = 1$ (possible since $\tau(0) \neq 0$ to start with) to see that $\lambda_x = \tau(x)$. Thus, $\tau(x+y) = \tau(x)\tau(y)$, showing that τ is a character of G. As eigenvectors form a basis, we get a full basis of characters.

Remark 18. As an offshoot of this discussion, observe that if $T : L^2(G) \mapsto L^2(G)$ is any normal operator that commutes with all the translations, then each character of G is an eigenvector of T.

Convolution: For $f, g : G \mapsto \mathbb{C}$, define $f \star g : G \mapsto \mathbb{C}$ by $(f \star g)(x) = \sum_y f(y)g(x-y)$. This can also be written as $\sum_y f(x-y)g(y)$, hence $f \star g = g \star f$. When f and g are probability vectors on G, we have the interpretation of $f \star g$ as the probability distribution of $X + Y$ where X and Y are independent random variables with distributions f and g respectively.

Fourier transform converts convolution to product. For,

$$\begin{aligned}\widehat{(f\star g)}(\chi) = \sum_{x\in G}(f\star g)(x)\overline{\chi}(x) &= \sum_{x\in G}\sum_{y\in G} f(x-y)\; g(y)\; \overline{\chi(x-y)}\;\overline{\chi(y)} \\ &= \sum_{z\in G}\sum_{y\in G} f(z)g(y)\overline{\chi(z)\chi(y)} \\ &= \left(\sum_{z\in G} f(z)\overline{\chi(z)}\right)\left(\sum_{y\in G} g(y)\overline{\chi(y)}\right) = \hat{f}(\chi)\hat{g}(\chi).\end{aligned}$$

Exercise 19. For any $f : G \mapsto \mathbb{R}$ and any $\chi \in \hat{G}$, show that $f \star \chi$ is a multiple of χ, and find the multiplying factor. How is this related to the previous discussion of characters as eigenvectors?

[2] Thanks to Ritvik Radhakrishnan for pointing this out during the lecture.

7. Duality

Let G be a finite abelian group and $\hat{G}$ the set of its characters. We have seen that they have the same cardinality. So far $\hat{G}$ is just a set. We now endow it with a group structure.

Define products of characters pointwise, i.e., $(\chi_1\chi_2)(x) = \chi_1(x)\chi_2(x)$. As $\chi_1, \chi_2 \in \hat{G}$ if $\chi_i \in \hat{G}$ and $\chi^{-1} = \overline{\chi} \in \hat{G}$ if $\chi \in \hat{G}$, it follows that $\hat{G}$ becomes a group. It is in fact a finite abelian group.

If $G = \mathbb{Z}_n$, then $\hat{G} = \{\chi_k := e_{2\pi k/n} : 0 \le k \le n-1\}$. Observe that $\chi_k\chi_\ell = \chi_{k+\ell(\text{mod } n)}$. In other words, as a group $\hat{G}$ is isomorphic to G itself with $k \mapsto \chi_k$ being an isomorphism.

If G is isomorphic to $\mathbb{Z}_{n_1} \times \ldots \times \mathbb{Z}_{n_k}$ (as any finite abelian group is), then we have seen that $\hat{G}$ can be identified (as a set) with $\widehat{\mathbb{Z}_{n_1}} \times \ldots \times \widehat{\mathbb{Z}_{n_k}}$. But it is also clear that if $\chi = \chi_1 \otimes \ldots \otimes \chi_k$ and $\chi' = \chi_1' \otimes \ldots \otimes \chi_k'$, then the pointwise product $\chi\chi' = (\chi_1\chi_1') \otimes \ldots \otimes (\chi_k\chi_k')$. Thus, $\hat{G}$ is isomorphic to G as a group, with $(\ell_1, \ldots, \ell_k) \mapsto \chi_{\ell_1} \otimes \ldots \otimes \chi_{\ell_k}$ being a natural isomorphism.

Thus for any finite abelian group, $\hat{G}$ is isomorphic to G as a group. However, this isomorphism is not canonical, hence we do not emphasize it. The issue is that while G is isomorphic to a product of cyclic groups, this isomorphism is not canonical/natural. For example, consider the Klein-4 group $G = \{1, a, b, c\}$ with $a^2 = b^2 = c^2 = 1$ and $ab = c$, $bc = a$, $ca = b$. Then $G = \{1, a\} \times \{1, b\}$ but also $G = \{1, b\} \times \{1, c\}$. These in turn lead to different isomorphisms of G with $\hat{G}$.

In contrast, we shall now see that the double dual $\hat{\hat{G}}$ is naturally isomorphic to G. That they are isomorphic is already clear, $\hat{\hat{G}} \cong \hat{G} \cong G$, it is the naturalness that is important. One may compare this to the analogous fact about duals and double duals of finite dimensional vector spaces.

The duality: For $x \in G$, the evaluation mapping $\mathrm{ev}_x : \hat{G} \mapsto \mathbb{T}$ defined by $\mathrm{ev}_x(\chi) = \chi(x)$, is clearly a homomorphism on $\hat{G}$. In other words, $\mathrm{ev}_x \in \hat{\hat{G}}$. Further, $x \mapsto \mathrm{ev}_x$ from G to $\hat{\hat{G}}$ is an isomorphism. It is a homomorphism because

$$\mathrm{ev}_{xy}(\chi) = \chi(xy) = \chi(x)\chi(y) = \mathrm{ev}_x(\chi)\mathrm{ev}_y(\chi) = (\mathrm{ev}_x\mathrm{ev}_y)(\chi).$$

It is injective: if $\mathrm{ev}_x = 1$, then $\chi(x) = 1$ for all $\chi \in \hat{G}$, which implies that $x = 1$ (otherwise the characters would not be able to separate 1 from x). Since G and $\hat{\hat{G}}$ have the same cardinality (because both are equal to $|\hat{G}|$), the homomorphism is an isomorphism of the two groups.

Exercise 20. Let G be a finite abelian group. Show that for any $x, y \in G$,

$$\sum_{\chi \in \hat{G}} \chi(x)\overline{\chi}(y) = \begin{cases} |G| & \text{if } x = y, \\ 0 & \text{if } x \neq y. \end{cases}$$

These relations are complementary to the orthogonality of characters.

8. Fourier inversion

Let G be a finite abelian group. Let $f \in L^2(G)$ and consider the $\hat{\hat{f}} \in L^2(G)$ (by the isomorphism between $\hat{\hat{G}}$ and G), then

$$\hat{\hat{f}}(x) = \frac{1}{\sqrt{N}}\langle \hat{f}, \mathrm{ev}_x \rangle_{L^2(\hat{G})} = \frac{1}{\sqrt{N}} \sum_{\chi \in \hat{G}} \hat{f}(\chi)\overline{\mathrm{ev}_x(\chi)} = \frac{1}{\sqrt{N}} \sum_{\chi} \langle f, \chi \rangle \chi(-x) = f(-x).$$

Loosely speaking, Fourier transform is its own inverse, but that is not quite right because of the negative sign on x. Two correct statements:

(1) $\hat{\hat{\hat{\hat{f}}}} = f$. In other words, $\mathcal{F}^4 = I$. To write $\mathcal{F}^4$ is misleading (as is writing $\hat{f}$, since the domain of the Fourier transform is not indicated and is ambiguous if several groups are floating around), it is actually $\mathcal{F}_{\hat{G}} \circ \mathcal{F}_G \circ \mathcal{F}_{\hat{G}} \circ \mathcal{F}_G$.

(2) Define $\check{f}(\chi) = \langle f, \overline{\chi} \rangle$ (pronounced "f check"), then $\check{\hat{f}} = f$ as

$$\check{\hat{f}}(x) = \langle \hat{f}, \overline{\mathrm{ev}_x} \rangle_{L^2(\hat{G})} = \sum_{\chi \in \hat{G}} \hat{f}(\chi)\mathrm{ev}_x(\chi) = \sum_{\chi} \langle f, \chi \rangle \chi(x) = f(x).$$

9. Poisson summation formula

Let H be a subgroup of a finite abelian group G and let $q : G \mapsto G/H$ be the quotient map. Let $N = |G|$ and $M = |H|$ so that $|G/H| = N/M$. Given $F \in L^2(G)$, we create a function $f \in L^2(G/H)$ by summing F over the coset, i.e.,

$$f(c) = \sum_{x \in q^{-1}\{c\}} F(x).$$

If χ is a character of G/H, then $\chi \circ q$ is a character of G and

$$\begin{aligned}
\hat{F}(\chi \circ q) \ &= \ \frac{1}{\sqrt{N}} \langle F, \chi \circ q \rangle_{L^2(G)} = \frac{1}{\sqrt{N}} \sum_{c \in G/H} \sum_{x \in q^{-1}\{c\}} F(x)\overline{\chi}(q(x)) \\
&= \frac{1}{\sqrt{N}} \sum_{c \in G/H} \overline{\chi}(c) \sum_{x \in q^{-1}\{c\}} F(x) \\
&= \frac{1}{\sqrt{N}} \sum_{c \in G/H} f(c)\overline{\chi}(c) \ = \ \frac{\sqrt{N/M}}{\sqrt{N}} \frac{1}{\sqrt{N/M}} \langle f, \chi \rangle_{L^2(G/H)} \ = \ \frac{1}{\sqrt{M}} \hat{f}(\chi).
\end{aligned}$$

To be pedantic one must write $\mathcal{F}_G(F)(\chi \circ q) = \mathcal{F}_{G/H}(f)(\chi)$ but we don't do that unless necessary. Now we are ready to state the

Theorem 21 (Poisson summation formula)**.** *In the above setting, for any $F \in L^2(G)$,*

$$\frac{1}{\sqrt{M}} \sum_{x \in H} F(x) = \frac{1}{\sqrt{N/M}} \sum_{\chi \in \widehat{G/H}} \hat{F}(\chi \circ q).$$

Proof. Define f from F as before. Then,

$$\sum_{x\in q^{-1}\{c\}} F(x) = f(c) \ = \ \frac{1}{\sqrt{N/M}} \sum_{\chi \in \widehat{G/H}} \hat{f}(\chi)\chi(c) \ = \ \frac{\sqrt{M}}{\sqrt{N/M}} \sum_{\chi \in \widehat{G/H}} \hat{F}(\chi \circ q)\chi(c).$$

In particular, setting $c = 0$ (the identity in G/H), we get the claimed identity. ■

From the proof, it may seem that we could have stated the more general identity instead of setting $c = 0$. Actually the general case can be recovered from the special case by applying to $x \mapsto F(x + x_0)$ for some $x_0 \in q^{-1}(c)$.

Example 22. Suppose $n = rs$. Let $G = \mathbb{Z}_n$ and let $H = \{0, r, 2r, \dots, (s-1)r\} \cong \mathbb{Z}_s$ so that $G/H \cong \mathbb{Z}_r$. The quotient map is of course $q(k) = k \pmod r$. Then the Poisson summation formula says that

$$\frac{1}{\sqrt{s}} \sum_{j=0}^{s-1} F(jr) = \frac{1}{\sqrt{r}} \sum_{k=0}^{r-1} \hat{F}(ks)$$

and more generally,

$$\frac{1}{\sqrt{s}} \sum_{j=0}^{s-1} F(jr + d) = \frac{1}{\sqrt{r}} \sum_{k=0}^{r-1} \hat{F}(ks)e^{2\pi i k d/r}.$$

Above, we considered a function on G and created a function on G/H. In the exercise below, this is done in the reverse direction (but the operation is not the reversal of the above!).

Exercise 23. Let H be a subgroup of a finite abelian group G and let $f \in L^2(G/H)$. Composing with the quotient map $q : G \mapsto G/H$, we get a function $f \circ q =: F \in L^2(G)$. Show that $\hat{F}(\chi \circ q) = |H|\hat{f}(\chi)$ for any $\chi \in \widehat{G/H}$.

CHAPTER 2

Uncertainty principles and signal recovery problems

1. Introduction

Uncertainty principle is the meta-statement that a function and its Fourier transform cannot both be localized. For Fourier transforms on $\mathbb{R}$ to which we come later, it is a fact that if $\hat{g}(t)$ is the Fourier transform of $g(x)$, then $\hat{g}(t/a)$ is the Fourier transform of $ag(ax)$, for any $a > 0$. As $a \to \infty$, the function $ag(ax)$ gets more and more concentrated or localized around zero, but the Fourier transform spreads out. This is an illustration of the same principle (although a feature here, one should disregard the other direction: that is a function is spread out the Fourier transform must get localized. Both can be spread out.). In the context of finite groups, we saw that if $n = pq$ and a function on $\mathbb{Z}_n$ is constant on multiples of p and zero elsewhere, then the Fourier transform is constant on multiples of q and zero elsewhere. The product of support sizes is n - if one goes down, the other goes up. In this chapter we see certain uncertainty principles based on p-norms, in the setting of finite abelian groups.

2. Some uncertainty principles based on p-norms

Let G be a finite abelian group of order N and let $f \in L^2(G)$. Throughout this section, we shall assume that f is not identically zero.

Then $\hat{f} \in L^2(G)$, is not identically 0, and we see that since $|\chi(x)| = 1$ for $x \in G$ and $\chi \in \hat{G}$, we have

$$|\hat{f}(\chi)| \le \frac{1}{\sqrt{N}} \sum_{x \in G} |f(x)| \qquad \text{and} \qquad |f(x)| \quad \le \frac{1}{\sqrt{N}} \sum_{\chi \in \hat{G}} |\hat{f}(\chi)|.$$

Thus (we write $\|\cdot\|_p$ instead of $\|\cdot\|_{L^p(G)}$ and $\|\cdot\|_{L^p(\hat{G})}$ if it is not ambiguous)

$$\|\hat{f}\|_\infty \le \frac{1}{\sqrt{N}} \|f\|_1 \qquad \text{and} \qquad \|f\|_\infty \le \frac{1}{\sqrt{N}} \|\hat{f}\|_1$$

and multiplying the two, we get our first *uncertainty principle*[3]:

$$\frac{\|f\|_1}{\|f\|_\infty} \times \frac{\|\hat{f}\|_1}{\|\hat{f}\|_\infty} \ge N. \tag{1}$$

[3] We have taken much of the approach to uncertainty principles from a recent paper of A. Wigderson and Y. Wigderson titled *The uncertainty principle: variations on a theme* . Their primary message, which I found quite illuminating, is this: Take (1) as the starting point. And if you have a measure of spread of functions $H(f)$, try to show an inequality such as $H(g) \ge \varphi(\|g\|_1/\|g\|_\infty)$ for some increasing function $\varphi : \mathbb{R}_+ \mapsto \mathbb{R}_+$ (preferably unbounded). An immediate corollary is an uncertainty principle for this measure of spread: $H(f) \times H(\hat{f}) \ge \varphi(N)$.

Why do we call it an uncertainty principle? In general, an upper bound on a ratio like $\|g\|_p/\|g\|_q$ for $p < q$ says that g must be localized or concentrated. For example, if $\|g\|_1 = \|g\|_\infty$, then g is supported at one point (in other words $g = \delta_x$ for some x). Another example is that $\|g\|_1 \le \sqrt{N}\|g\|_2$ by Cauchy-Schwarz, with equality if and only if g is constant. If $g = \delta_x$ for some x, then $\|g\|_1 = \|g\|_2$. In this case too, we see that the ratio of L^1 to L^2 norms is small when the function is localized[4].

Thus, the inequality (1) says that in the L^1 to L^∞ measurement of concentration, both f and $\hat{f}$ cannot be too localized. Uncertainty principles are statements that say this, with different measures of localization/concentration. Another measure of localization is the cardinality of the support of the function.

For any g on a finite set, it is clear that $\|g\|_1 \le |S_g|\|g\|_\infty$, where $|S_g|$ is the cardinality of the support $S_g = \{x : g(x) \neq 0\}$. Writing this inequality for $f \in L^2(G)$ and for $\hat{f} \in L^2(G)$, we get the Donoho-Stark uncertainty principle

$$|S_f| \times |S_{\hat{f}}| \ge N \tag{2}$$

by applying (1). We know that equality can be achieved, for example, by $f(k) = \mathbf{1}_{k=0 \pmod p}$ on $\mathbb{Z}_{pq}$, in which case $\hat{f}(\ell) = c\mathbf{1}_{\ell=0 \pmod q}$ for a constant c (as usual we abuse notation and write $\hat{f}(\ell)$ instead of $\hat{f}(\chi_\ell)$ when working with the cyclic group). Thus $|S_f| = q$ and $|S_{\hat{f}}| = p$ and their product is the size of the group.

The rest of this section is optional. For our purposes, (2) and the idea of deriving it from (1) are sufficient.

Further analysis of (p, q) uncertainty principles: Although there is no L^p-norm for $0 \le p \le 1$, we may still use $\|g\|_p = (\sum_x |g(x)|^p)^{1/p}$ as a measure of the size of g. Not relevant to us, but one can also use it to get a metric $d(f, g) = \|f - g\|_p^p$. Observe that $\|g\|_p^p \to |S_g|$ as $p \to 0$. Thus the support size is a limiting case of the p-norms. In proving Donoho-Stark, we used the simple inequality $\|g\|_1 \le \|g\|_\infty |S_g|$. To extend it to general $p \in [0, 1]$, observe that

$$\|g\|_1 = \sum_x |g(x)| \le \|g\|_\infty^{1-p} \sum_x |g(x)|^p = \|g\|_\infty^{1-p}\|g\|_p^p.$$

Invoking (1), we get the following uncertainty principle

$$\left(\frac{\|f\|_p}{\|f\|_\infty}\right)^p \times \left(\frac{\|\hat{f}\|_p}{\|\hat{f}\|_\infty}\right)^p \ge N. \tag{3}$$

[4] In class I got into a twist by writing $\|g\|_2/\|g\|_1^2$ as an increasing transformation of the coefficient of variation of g, i.e., $\mathrm{var}(g)/\mathrm{mean}(g)^2$, and saying that the latter quantity is a measure of spread. This led to a confusion that we are looking at the ratio of a higher p-norm to a lower p-norm. But variance of g is the spread of the *values* of g, whereas what we are talking about is the spread on the *domain side*. For example, the constant function has zero variance whereas a delta-function has positive variance! Thanks to Chinmay S. I. for clearing up this point.

that interpolates between Donoho-Stark inequality (2) (case $p=0$) and (1) (case $p=1$). We can further extend to a (p,q) uncertainty principle with $0\le p\le 1$ and $1\le q\le\infty$ as follows: Write $\frac{1}{q}=\frac{p/q}{p}+\frac{(q-p)/q}{\infty}$ and use Hölder's inequality to write $\|f\|_q\le\|f\|_p^{p/q}\|f\|_\infty^{(q-p)/q}$ or equivalently

$$\frac{\|f\|_p}{\|f\|_q}\ge\left(\frac{\|f\|_p}{\|f\|_\infty}\right)^{(q-p)/p}.$$

Multiply with the analogous inequality for $\hat{f}$, raise to the power p, and use the (p,∞) uncertainty principle to get the (p,q)-uncertainty principle

$$\left(\frac{\|f\|_p}{\|f\|_q}\right)^p\times\left(\frac{\|\hat{f}\|_p}{\|\hat{f}\|_q}\right)^p\ge N^{1-\frac{p}{q}}. \tag{4}$$

When $p=0$ the left hand side should be interpreted as $|S_f|\times|S_{\hat{f}}|$, which is what one gets as $p\downarrow 0$.

Next consider $p\in(1,2)$ and let $p'\in(2,\infty)$ be the conjugate exponent defined by $\frac{1}{p}+\frac{1}{p'}=1$. For a linear transformation $T:U\mapsto V$, where U,V are normed vector spaces, the operator norm is defined as

$$\|T\|_{U\to V}:=\sup_{u\neq 0}\frac{\|Tu\|_V}{\|u\|_U}=\sup_{u\in U:\|u\|_U=1}\|Tu\|_V.$$

We say that T is a bounded operator is $\|T\|_{U\to V}<\infty$. This is always the case when U is finite dimensional. When the underlying spaces are $U=L^p(X)$ and $V=L^q(Y)$ (where X,Y are finite sets for now, but more generally they can be measure spaces) we write $\|T\|_{p\to q}$ for $\|T\|_{L^p(X)\to L^q(Y)}$.

Riesz-Thorin interpolation theorem: Assume $1\le p_0,p_1,q_0,q_1\le\infty$. For $0<\theta<1$ define p_θ,q_θ by $p_\theta^{-1}=(1-\theta)p_0^{-1}+\theta p_1^{-1}$ and $q_\theta^{-1}=(1-\theta)q_0^{-1}+\theta q_1^{-1}$. Then,

$$\|T\|_{p_\theta\to q_\theta}\le\|T\|_{p_0\to q_0}^{1-\theta}\|T\|_{p_1\to q_1}^{\theta}.$$

The theorem is true in general measure spaces, but one must say a few words first about T being bounded from L^{p_0} to L^{q_0} etc. We wish to apply this to the Fourier transform $\mathcal{F}:L^2(G)\mapsto L^2(\hat{G})$. We know that $\|\hat{f}\|_\infty\le\frac{1}{\sqrt{N}}\|f\|_1$ and $\|\hat{f}\|_2=\|f\|_2$. That is, $\|\mathcal{F}\|_{1\to\infty}=1/\sqrt{N}$ and $\|\mathcal{F}\|_2=1$. For $p\in(1,2)$, let $\theta=\frac{2}{p'}$ so that

$$\frac{1}{p}=\frac{1-\theta}{1}+\frac{\theta}{2}\qquad\text{and}\qquad\frac{1}{p'}=\frac{1-\theta}{\infty}+\frac{\theta}{2}.$$

From the Riesz-Thorin interpolation theorem, we get the *Hausdorff-Young inequality*:

$$\|\hat{f}\|_{p'}\le N^{\frac{1}{2p'}-\frac{1}{2p}}\|f\|_p.$$

Applying the same to the inverse Fourier transform, we get $\|f\|_{p'}\le N^{\frac{1}{2p'}-\frac{1}{2p}}\|\hat{f}\|_p$. Multiplying, we get the (p,p')-uncertainty principle

$$\frac{\|f\|_p}{\|f\|_{p'}}\times\frac{\|\hat{f}\|_p}{\|\hat{f}\|_{p'}}\ge N^{\frac{1}{p}-\frac{1}{p'}}.$$

Extend this to a (p,q) uncertainty principle for $p \in (1,2)$ and $q \in (2,p']$ to get

$$\frac{\|f\|_p}{\|f\|_q} \times \frac{\|\hat{f}\|_p}{\|\hat{f}\|_q} \geq N^{\frac{1}{p}-\frac{1}{q}}. \tag{5}$$

Exercise 1. Prove (5) using the (p,p')-uncertainty principle.

Summary: In summary, we have proved (p,q) uncertainty principles when $p \in [0,1]$ and $q \in [1,\infty]$ or $p \in [1,2]$ and $q \in [2,p']$. The inequalities (4) and (5) subsume all others proved in this section.

It may be worth recalling the standing assumption that $f \neq 0$. A useful way of stating each of the uncertainty principles is that if the conclusion is violated (e.g., if $|S_f| \times |S_{\hat{f}}| < N$), then $f = 0$.

3. Robust versions of the Donoho-Stark uncertainty principle

The support is a delicate thing. In the real world, no function can be said to be exactly zero at any point. Mathematical theorems which hold under certain conditions, but break down under the slightest perturbations, are generally not saying anything about the real world, because the hypotheses are never satisfied! In other words, we should look for theorems that are robust, or not too sensitive to the assumptions. For example, a robust version of the uncertainty principle would say that a function and its Fourier transform cannot both be *nearly* supported on small sets.

We prove two such versions in this section. First we need a definition for approximate support.

Definition 2. If $f : G \mapsto \mathbb{C}$, we say that $A \subseteq G$ is an (p,ϵ) support for f if $\|f\mathbf{1}_{A^c}\|_p \leq \varepsilon \|f\|_p$. Here $1 \leq p < \infty$ and $0 \leq \varepsilon \leq 1$.

The definition can be clearly made for L^p functions on an arbitrary measure space. A (p,ε) support always exists, for example G itself is always one. When $\varepsilon = 0$, the smallest (p,ε) support is the usual support. For $\varepsilon > 0$, in general there is no unique (p,ε) support. Note that we have not included any phrase like "the smallest set with...", but including that would not make it unique either.

Theorem 3. *Let G be a finite abelian group of order N. Assume that $f : G \mapsto \mathbb{R}$ is not identically zero. Suppose that $A \subseteq G$ are $B \subseteq \hat{G}$ are $(1,\varepsilon)$ and $(1,\delta)$ supports for f and $\hat{f}$ respectively. Then $|A| \times |B| \geq N(1-\varepsilon)(1-\delta)$.*

Proof. We observe that

$$\begin{aligned}\sqrt{N}\|\hat{f}\|_\infty \leq \|f\|_1 \;&\leq\; (1-\varepsilon)^{-1}\|f\mathbf{1}_A\|_1 \;\leq\; (1-\varepsilon)^{-1}\|f\|_\infty |A|,\\ \sqrt{N}\|f\|_\infty \leq \|\hat{f}\|_1 \;&\leq\; (1-\delta)^{-1}\|\hat{f}\mathbf{1}_B\|_1 \;\leq\; (1-\delta)^{-1}\|\hat{f}\|_\infty |B|.\end{aligned}$$

The first inequalities are from the definition of the Fourier transform (as $\hat{\hat{f}}(x) = f(-x)$, the first inequality in the second line follows from the same). The second inequalities are from the definition of A and B as approximate supports. The third inequality is obvious. Multiply the two inequalities to get the theorem. ■

The above theorem is due to Wigderson A., and Wigderson Y., inspired by the stronger Donoho-Stark theorem that states such an inequality for L^2 approximate supports.

Theorem 4 (Donoho-Stark)**.** *Let G be a finite abelian group of order N. Assume that $f : G \mapsto \mathbb{R}$ is not identically zero. Suppose that $A \subseteq G$ are $B \subseteq \hat{G}$ are $(2, \varepsilon)$ and $(2, \delta)$ supports for f and $\hat{f}$ respectively. Then $|A| \times |B| \geq N(1 - \varepsilon - \delta)^2$.*

In preparation for the proof, consider a linear operator $T : U \mapsto V$, where U and V are finite dimensional normed complex vector spaces. If T^*T has eigenvectors u_k with eigenvalues λ_k (which are non-negative), then for any $u \in U$,

$$\|Tu\|^2 = \langle T^*Tu, u\rangle = \sum_k \lambda_k |\langle u, u_k\rangle|^2 \leq (\max_k \lambda_k) \sum_k |\langle u, u_k\rangle|^2 \leq \mathrm{tr}(T^*T)\|u\|^2.$$

Here we bounded the maximum eigenvalue of T^*T by the sum of eigenvalues. Thus, $\|T\|_{U\to V} \leq \sqrt{\mathrm{tr}(T^*T)}$. The right side is also known as the Hilbert-Schmidt or Frobenius norm and denoted $\|T\|_F$ (or $\|T\|_{\mathrm{HS}}$). If we regard T as a matrix by fixing bases of U, V, then the Frobenius norm is just the Euclidean norm of the vector got by writing the matrix as a vector of dimension $\dim(U)\dim(V)$.

Proof of Theorem 4. Define the following projection operators on $L^2(G)$, for $A \subseteq G$ and $B \subseteq \hat{G}$:

$$P_A f = f\mathbf{1}_A, \qquad \hat{P}_B f = (\hat{f}\mathbf{1}_B)\check{} \text{ (or equivalently } \widehat{\hat{P}_B f} = \hat{f}\mathbf{1}_B).$$

In other words, P_A is restriction of f to A and $\hat{P}_B$ is restriction of the Fourier transform to B. The approximate support conditions and Plancherel relation mean that

$$\|(I - P_A)f\|_2 \leq \varepsilon\|f\|_2 \qquad \text{and} \qquad \|(I - \hat{P}_B)f\|_2 \leq \delta\|f\|_2$$

and hence

$$\begin{aligned}\|(I - \hat{P}_B P_A)f\|_2 &\leq \|(I - \hat{P}_B f)\|_2 + \|\hat{P}_B(I - P_A)f\|_2 \\ &\leq (\delta + \varepsilon)\|f\|_2.\end{aligned}$$

Therefore, $\|\hat{P}_B P_A f\|_2 \geq (1 - \delta - \varepsilon)\|f\|_2$. Therefore, $(1 - \delta - \varepsilon)^2 \leq \|\hat{P}_B P_A\|^2 \leq \|\hat{P}_B P_A\|_F^2$. As $P_A, \hat{P}_B$ are projections, we see that $(\hat{P}_B P_A)^* \hat{P}_B P_A = \hat{P}_B P_A$ and hence $\|\hat{P}_B P_A\|_F^2 = \mathrm{tr}(\hat{P}_B P_A)$. To compute the trace, we use the basis δ_k, $k \in G$. If $k \notin A$, then $\hat{P}_B P_A \delta_k = 0$. If $k \in A$, then $\hat{P}_B P_A \delta_k = \hat{P}_B \delta_k = (\mathbf{1}_B \hat{\delta_k})\check{}$. Hence the (k, k) entry of $\hat{P}_B P_A$ is equal to

$$\frac{1}{\sqrt{N}} \sum_{\chi \in \hat{G}} (\mathbf{1}_B \hat{\delta}_k)(\chi)\, \chi(k) \;=\; \frac{1}{N} \sum_{\chi \in B} \overline{\chi(k)}\chi(k) \;=\; \frac{1}{N}|B|.$$

Summing over $k \in A$, we get $\|\hat{P}_B P_A\|_F^2 = \frac{1}{N}|A| \times |B|$. Combining this with the lower bound for the Frobenius norm, we get $|A| \times |B| \geq (1 - \delta - \varepsilon)^2 N$. ■

Remark 5. Bounding the operator norm by the Frobenius norm may seem to be loose, since we bound the maximum eigenvalue by the sum. On the other hand, even with $\varepsilon = \delta = 0$, we know that equality can be achieved in Theorem 4. It is instructive to work out $\hat{P}_B P_A$ explicitly in that example to see how this comes about.

Here is a robust uncertainty principle mixing the L^1 and L^2 norms. The proof is not hard - it is closer to that of Theorem 3 than that of Theorem 4.

Exercise 6. Let G be a finite abelian group of order N. Let $f : G \mapsto \mathbb{C}$. Suppose A_1 is a $(1, \varepsilon)$ support for f and B_2 is a $(2, \delta)$ support for $\hat{f}$. Then, $|A_1| \times |B_2| \geq N(1-\varepsilon)^2(1-\delta)^2$.

Observe that $\hat{P}_B P_A f = f$ if and only if $\hat{f}$ is supported on B and f is supported on A. Hence, any upper bound on $\|\hat{P}_B P_A\|$ that is strictly less than 1 may be interpreted as an uncertainty principle. Donoho and Stark remark in their paper that while $\sqrt{|A| \cdot |B|/N}$ is a natural upper bound for $\|\hat{P}_B P_A\|$, it is the latter norm itself which is the key quantity of interest. The following exercise contains a quantitative refinement of Theorem 4.

Exercise 7. If $A \subseteq G$ and $B \subseteq \hat{G}$, show that $\|\hat{P}_B P_A\| = \sup\{\frac{\|\hat{g}\mathbf{1}_B\|_2}{\|g\|_2} : g \in L^2(G),\ S_g \subseteq A\}$. Deduce that if $|A| \times |B| < \alpha N$ for some $\alpha \in (0,1)$, then $\sum_{\chi \in B} |\hat{g}(\chi)|^2 \leq \alpha \sum_{\chi \in \hat{G}} |\hat{g}(\chi)|^2$ for any $g \in L^2(G)$ supported on A.

4. Applications of uncertainty principle to signal recovery problems

4.1. **Recovering a bandlimited signal from periodic measurements.** Consider $\mathbb{Z}_n$ and assume that $n = rs$. Recall the Poisson summation formula

$$\frac{1}{\sqrt{s}} \sum_{j=0}^{s-1} F(jr + d) = \frac{1}{\sqrt{r}} \sum_{k=0}^{r-1} \hat{F}(ks) e^{2\pi i k d/r}$$

for any $F \in L^2(\mathbb{Z}_n)$. Let $I_k = \{n-k+1, \dots, n-1, 0, \dots, k-1\}$ be the “interval” of length $2k-1$ in $\mathbb{Z}_n$ centered at 0. Suppose F is supported on I_p where $2p - 1 \leq r$. Then on the left side of the Poisson summation formula, only the term with $j = 0$ contributes, and hence we get

$$F(d) = \frac{\sqrt{s}}{\sqrt{r}} \sum_{k=0}^{r-1} \hat{F}(ks) e^{2\pi i k d/r} \quad \text{for } d \in I_r.$$

Thus, we can recover F from $\hat{F}(0), \hat{F}(s), \dots, \hat{F}((r-1)s)$. This is not too surprising, as the space of functions supported on I_p has dimension $2p - 1$, and we have $r \geq 2p - 1$ Fourier coefficients here. Reversing the role of F and $\hat{F}$, we can also say that $\hat{F}$ is supported on I_q where $2q - 1 \leq s$, then $\hat{F}$ can be recovered from $F(0), F(r), \dots, F((s-1)r)$.

$$\hat{F}(d) = \frac{\sqrt{r}}{\sqrt{s}} \sum_{k=0}^{s-1} F(kr) e^{2\pi i k d/s} \quad \text{for } d \in I_q.$$

By the Fourier inversion formula, this also recovers F. What we have proved is the discrete version of the Shannon-Nyquist formula.

Theorem 8 (Shannon-Nyquist, discrete version)**.** *Suppose $F : \mathbb{Z}_n \mapsto \mathbb{C}$ is a signal bandlimited to frequency q, then F can be recovered from its values sampled at points regularly spaced r distance apart, provided $r \le \frac{n}{2q-1}$.*

Strictly speaking, we have proved this only when n is a multiple of r. The general case is outlined in the following exercise.

Exercise 9. Assume that F is bandlimited to frequency q. Show that the linear transformation from $(\hat{F}(k))_{-q+1\le 0\le q-1}$ to $(F(jr))_{0\le j\le \lfloor \frac{n-1}{r}\rfloor}$ is injective if $r \le \frac{n}{2q-1}$ (compute the matrix of this linear transformation explicitly) and hence deduce Shannon-Nyquist theorem.

With the interpretation of $\mathbb{Z}_n$ as time and $\hat{\mathbb{Z}}_n$ as frequencies, it is the convention to say that a function supported on I_p is time-limited to p and that a function whose Fourier transform is supported on I_q is band-limited to q. Here we are using a special feature of $\hat{\mathbb{Z}}_n$: the characters $\chi_k(j) = e^{2\pi ijk/n}$, can be ordered in increasing order of $|k|$, and $|k|$ indicates the frequency (i.e., how rapidly the character changes from one point to the next). A band-limited signal is one that used only low frequencies. The real content of this theorem comes from the fact that band-limited assumption may actually be an assumption satisfied in reality.

4.2. **Recovering a bandlimited signal from incomplete noisy measurements.** Let $B = I_q = \{-q, \dots, q\}$ and $A = \{0, r, 2r, \dots, (s-1)r\}$ (here no relationship is assumed a priori on q, r, s, n). Make the assumption of bandlimitedness: That $\hat{f}$ is supported on B. Suppose we observe the signal f only on A^c and those observations are also noisy. That is, we observe $g(k) = (f(k) + \nu(k))\mathbf{1}_{k\notin A}$. The question is: Can we recover f from g? Of course, as there is noise, one does not expect exact recovery, what one asks for is an estimated signal close to the original one. The following theorem assures us that it can be done.

Theorem 10. *Assume that $|A| \times |B| < N$ and let $Q = (I - P_A\hat{P}_B)^{-1}$. Then $\|f - Qg\|_2 \le C_{A,B}\|\nu\|_2$ where $C_{A,B}^{-2} = 1 - \frac{1}{N}|A| \times |B|$.*

Proof. We know that $\|P_A\hat{P}_B\| \le \frac{1}{N}|A| \times |B|$, hence it follows that $I - P_A\hat{P}_B$ is invertible and in fact

$$Q = I + P_A\hat{P}_B + (P_A\hat{P}_B)^2 + \dots$$

Observe that $g = (I - P_A)f + \nu = (I - P_A\hat{P}_B)f + \nu$ since $\hat{P}_B f = f$. Therefore, $f = Q(g - \nu) = Qg - Q\nu$. Since $\|Q\nu\|_2 \le \|Q\|\|\nu\|_2$, using the bound for the norm of Q, we get the conclusion. ∎

Remark 11. Donoho and Stark remark that the expansion formula for Q can be used efficiently to numerically find Qg by starting with $f_0 = g$ and setting $f_{k+1} = g + P_A\hat{P}_B f_k$ for $k \ge 0$. As

$P_A\hat{P}_B f_k(x) = 0$ for $x \notin A$, observe that $f_k(x) = g(x)$ for all $x \notin A$ and for all k. It is the unobserved values that get updated at each iteration and finally converge to the fixed point of the equation $h = g + P_A\hat{P}_B h$ which is $h = Qg$.

CHAPTER 3

Social choice and Fourier analysis on $\mathbb{Z}_2^n$

1. Fourier analysis on $\mathbb{Z}_2^n$ revisited

For the product group $\mathbb{Z}_2^n$, we introduce a few other basic notions (I do not know if these can be generalised to other groups)[5].

Discrete derivative and influence: For $f : \mathbb{Z}_2^n \mapsto \mathbb{C}$ and index $j \in [n]$, define

$$D_j f(x) := \frac{1}{2}[f(x_1, \dots, x_{j-1}, 1, x_{j+1}, \dots, x_n) - f(x_1, \dots, x_{j-1}, -1, x_{j+1}, \dots, x_n)].$$

This measures the dependence of the value of f on the jth co-ordinate, when all other co-ordinates are fixed. It depends on the values of the other co-ordinates (but crucially, not on x_j itself). An overall measure of the *influence* of the jth co-ordinate is given by

$$\mathrm{Inf}_j(f) := \frac{1}{2^n}\|D_j f\|_2^2.$$

Of special interest are *Boolean functions* $f : \mathbb{Z}_2^n \mapsto \{-1, 1\}$ (sometimes the target space is written as $\{0, 1\}$). We may think of this as a voting rule, where there are two candidates, and each $x \in \mathbb{Z}_2^n$ denotes a particular voting pattern (x_i being the vote of the ith person). Each Boolean function f gives a different way in which the votes are pooled together to arrive at a decision. If f is a Boolean function, $D_j f(x) = \pm 1$ if the jth voter's vote changes the final decision (when all others fix their votes) and $D_j f(x) = 0$ if the decision can be made on others' votes, disregarding the jth voter. Therefore, the influence of voter j is the probability that his/her vote changes the decision:

$$\begin{aligned}\mathrm{Inf}_j(f) &= \frac{1}{2^n}\sum_{x \in \mathbb{Z}_2^n} \mathbf{1}_{f(x_1,\dots,x_{j-1},1,x_{j+1},\dots,x_n) \neq f(x_1,\dots,x_{j-1},-1,x_{j+1},\dots,x_n)} \\ &= \mathbf{P}\{f(x) \neq f(x^j)\}\end{aligned}$$

where $x_i^j = x_i$ if $i \neq j$ and $x_i^j = -x_i$ if $i = j$. We use the word probability to mean the uniform probability measure on $\mathbb{Z}_2^n$. In other words, each voter votes independently with equal probability to either candidate.

Example 1. Let $S \subseteq [n]$. Then $D_j \chi_S = 0$ if $j \notin S$ while $D_j \chi_S = \chi_{S \setminus \{j\}}$ if $S \ni j$. Note the formal similarity to the way we usually differentiate a polynomial (except that in $\mathbb{Z}_2^n$ the degree of any variable

[5] In everything we do on $\mathbb{Z}_2^n$, we follow Ryan O'donnell's book *Analysis of Boolean functions*. The book has a wealth of material and very good exposition - we make a small subselection.

is at most 1). In this case, $D_j\chi_S(x) = \pm 1$ for all x and hence $\text{Inf}_j(\chi_S) = 1$ for $j \in S$. Of course the influence is zero for $j \notin S$.

Characters χ_S are Boolean functions. Here are a few others.

Example 2. If n is odd, we define the *majority function* $M(x_1, \ldots, x_n) = \mathbf{1}_{x_1+\ldots+x_n>0} - \mathbf{1}_{x_1+\ldots+x_n<0}$ (here the addition is in $\mathbb{R}$, not in $\mathbb{Z}_2$). A *dictator function* is one of the form $x \mapsto x_j$, for some $j \in [n]$ (then j is called the dictator). A *recursive majority* function is defined as follows: Imagine a country that is divided into states, states divided into districts, districts divided into towns. In each town the majority vote is taken to decide which party candidate is elected. A majority vote among the town representatives decides the state representative and the majority among state representatives gives the overall decision.

Fourier expansion of derivatives: Let $f : \mathbb{Z}_2^n \mapsto \mathbb{R}$ and $j \in [n]$. We write $j = 1$ for simplicity of notation. If $S \ni 1$, then $\widehat{D_1 f}(S) = 0$ because $D_1 f(x)$ does not depend on x_1. If $S \not\ni 1$, then

$$\begin{aligned}\widehat{D_1 f}(S) &= \frac{1}{2}\sum_{x\in\mathbb{Z}_2^n}[f(1, x_2, \ldots, x_n) - f(-1, x_2, \ldots, x_n)]\prod_{i\in S} x_i \\ &= \sum_{x\in\mathbb{Z}_2^n} f(x_1, \ldots, x_n)x_1\prod_{i\in S} x_i \\ &= \hat{f}(S \sqcup \{1\})\end{aligned}$$

In general $D_j\hat{f}(S) = \hat{f}(S \cup \{j\})$ if $j \notin S$, and equal to zero otherwise. Equivalently,

$$D_j f(x) = \sum_{S:S\ni j} \hat{f}(S)\chi_{S\setminus\{j\}}(x)$$

One could also have arrived at this by applying D_j to the expression $f(x) = \sum_S \hat{f}(S)\chi_S$. Going further, this also shows that $\hat{f}(S) = \widehat{D_S f}(\emptyset)$ where $D_S = D_{j_1}\ldots D_{j_s}$ for $S = \{j_1, \ldots, j_s\}$. This is (a multiple of) the mean value of $D_S f$.

Exercise 3. Let $n = 2m + 1$ be odd and let M be the majority function. Show that

$$\hat{M}(S) = \begin{cases} 0 & \text{if } |S| \text{ is even,} \\ (-1)^k \frac{\binom{m}{k}\binom{2m}{m}}{2^{m-\frac{1}{2}}\binom{2m}{2k}} & \text{if } |S| = 2k+1. \end{cases}$$

[*Hint:* Take $S = \{1, \ldots, 2k+1\}$ and calculate $D_1 M(x)$. Use the relationship between Fourier coefficients of M and $D_1 M$.]

The heat semi-group: For $f \in L^2(\mathbb{Z}_2^n)$ and $t \geq 0$ (we interpret t as time), define

$$Q_t f = \frac{1}{2^{n/2}}\sum_{S\subseteq[n]} e^{-t|S|}\hat{f}(S)\chi_S.$$

In other symbols, $\widehat{Q_tf}(S) = e^{-t|S|}\hat{f}(S)$. Thus the mean value stays constant as $\widehat{Q_tf}(\emptyset) = \hat{f}(\emptyset)$, while every other Fourier coefficient decays exponentially fast. From the definition it immediately follows that $Q_{t+s} = Q_t \circ Q_s$, in other words $t \mapsto Q_t$ is a semi-group. It may also be seen that Q_t is self-adjoint:

$$\langle Q_tf, g\rangle \;=\; \langle \widehat{Q_tf}, \hat{g}\rangle \;=\; \sum_S e^{-t|S|}\hat{f}(S)\overline{\hat{g}(S)} \;=\,= \langle \hat{f}, \widehat{Q_tg}\rangle \;=\; \langle f, Q_tg\rangle.$$

One can think of Q_tf as interpolating between $f = Q_0(f)$ and the constant function $2^{-n/2}\hat{f}(\emptyset)\mathbf{1} = Q_\infty f$. If one thinks of f as giving the initial temperatures at points of $\mathbb{Z}_2^n$, what Q_t does is to smooth it (like heat flows from hotter to colder places, reducing gradient in temperature) all the way till every vertex is at the same temperature. To make the analogy with the heat equation closer, differentiate Q_tf w.r.t t to get

$$\begin{aligned}
\frac{d}{dt}Q_tf(x) &= -\frac{1}{2^{n/2}}\sum_{S\subseteq[n]} e^{-t|S|}|S|\hat{f}(S)\chi_S(x)\\
&= -\frac{1}{2^{n/2}}\sum_{S\subseteq[n]}\sum_{j\in S} e^{-t|S|}\hat{f}(S)\chi_S(x)\\
&= -\frac{1}{2^{n/2}}\sum_{j\in[n]}\sum_{S:S\ni j} e^{-t|S|}\hat{f}(S)\chi_S(x)\\
&= -\frac{1}{2^{n/2}}\sum_{j\in[n]} x_j \sum_{S:S\ni j} e^{-t|S|}\hat{f}(S)\chi_{S\setminus\{j\}}(x)\\
&= -LQ_tf(x)
\end{aligned}$$

where $L = \sum_{j=1}^n x_jD_j$. Comparing this with the usual heat equation, we may call L the Laplacian[6].

Exercise 4. (1) Find the Fourier coefficients of Lf in terms of Fourier coefficients of f. (2) Show that $Q_t = e^{-tL}$. (3) Find the spectral decomposition of L and of Q_t for any $t > 0$.

Remark 5. Some prefer to write $\rho = e^{-t}$ and define $T_\rho f$ by $\widehat{T_\rho f}(S) = \rho^{|S|}\hat{f}(S)$. This has the advantage of allowing ρ to be negative. Usually only $\rho \in [-1, 1]$ is considered. One quantity of interest is

$$(1) \qquad \langle T_\rho f, f\rangle = \frac{1}{2^n}\sum_S |\hat{f}(S)|^2\rho^{|S|} \;=\; \mathbf{E}[f(x)f(y)]$$

[6] To see why it is called Laplacian, observe that

$$x_jf(x) = \frac{1}{2}[f(x_1,\ldots,x_{j-1},x_j,x_{j+1},\ldots,x_n) - f(x_1,\ldots,x_{j-1},-x_j,x_{j+1},\ldots,x_n)]$$

and hence $Lf(x) = \frac{1}{2}\sum_{y\sim x}[f(y) - f(x)]$ where $y \sim x$ means that y and x differ in exactly one co-ordinate. In particular, $Lf(x) = 0$ if and only if $f(x)$ is the mean value of f at its neighbours in the Hamming cube (the graph with vertices $\{-1,1\}^n$ and edges between vertices that differ in a single co-ordinate). This is analogous to the mean value property of harmonic functions (which are functions that satisfy $\Delta u = 0$) and the formula for Lf is itself analogous to the fact that $\Delta u(x) = c_d \lim_{r\to 0}\int_{x+r\mathbb{S}^{d-1}}[u(y) - u(x)]d\sigma(y)$.

where $x, y \in \mathbb{Z}_2^n$ are random variables that satisfy (1) (x_k, y_k) are independent across k, (2) $\mathbf{E}[x_k] = \mathbf{E}[y_k] = 0$, (3) $\mathbf{E}[x_k y_k] = \rho$. In this case, expanding f in Fourier series,

$$\begin{aligned}\mathbf{E}[f(x)f(y)] &= \frac{1}{2^n}\sum_{S,T} \hat{f}(S)\hat{f}(T)\mathbf{E}[\chi_S(x)\chi_T(y)] \\ &= \frac{1}{2^n}\sum_{S,T} \hat{f}(S)\hat{f}(T)\ \mathbf{E}\left[\prod_{i\in S\setminus T} x_i \prod_{i\in T\setminus S} y_i \prod_{i\in S\cap T} x_i y_i\right] \\ &= \frac{1}{2^n}\sum_{S} |\hat{f}(S)|^2 \rho^{|S|}\end{aligned}$$

because the expectation factors over i. This is how we got the second equality in (1). The quantity $\mathbf{E}[f(x)f(y)]$ has the following interpretation for a Boolean function f: Start with x uniformly randomly chosen from $\mathbb{Z}_2^n$. For each co-ordinate k, with probability $\frac{1}{2}+\frac{1}{2}\rho$, keep x_k as it is, and with the remaining probability negate it. Let the resulting random vector be called y. Then, if f is a Boolean function, then $\frac{1}{2}+\frac{1}{2}\mathbf{E}[f(x)f(y)]$ is the probability that the value of f did not change. This is a measure of how stable f is to perturbation of a few co-ordinates, and is denoted $\text{Stab}_\rho(f)$. In this language, $\text{Stab}_\rho(f) = \langle T_\rho f, f\rangle$. For positive ρ, we may also write this as $\langle Q_t f, f\rangle$ with $e^{-t} = \rho$.

For those familiar with some probability, we explain how Q_t and L are related to a certain Markov chain.

The Markov chain interpretation: Let $X(t)$, $t \geq 0$, be a $\mathbb{Z}_2^n$-valued Markov chain in continuous time defined as follows: At each $j \in [n]$, there is a Poisson clock, meaning that there is a sequence of random times $0 < T_{j,1} < T_{j,2} < \ldots$ where $T_{j,r+1} - T_{j,r}$ are i.i.d. Exponential random variables with mean 1. At the time $T_{j,r}$, the jth co-ordinate of X_t is refreshed, meaning that it is reset to ± 1 with equal probability. If the process starts at $X(0) = x \in \mathbb{Z}_2^n$, then we claim that the expected value of $f(X_t)$ at time t is precisely $Q_t f(x)$. If you are not familiar with Markov chains, ignore this, otherwise take it as an exercise (basically the differential equation for $Q_t f$ is the Kolmogorov equation for a Markov chain). In particular, if $f = \delta_y$, then $Q_t f(x)$ is the probability that the Markov chain started from x at time 0 is at y at time t. In other words, Q_t is the transition matrix for the Markov chain. As we saw before, $-L = \frac{d}{dt}Q_t\big|_{t=0}$. In Markov chain literature, this is called the generator of the Markov chain. One fact that is clear from the Markov chain interpretation but not so obvious from the Fourier definition is the *positivity* of the operators Q_t: If $f(x) \geq 0$ for all x, then $Q_t f(x) \geq 0$ for all x.

Suppose the Markov chain starts at a uniform random point in $\mathbb{Z}_2^n$, i.e., $X_0 \sim \text{uniform}(\mathbb{Z}_2^n)$. What about X_t? It is easy to see that X_t is also uniform on $\mathbb{Z}_2^n$, but X_t is not independent of X_0. As the co-ordinates evolve independently, it is enough to look at what $(X_0(1), X_t(1))$. If there is no clock-ring between time 0 and time t (an event of probability e^{-t}), then $X_0(1) = X_t(1)$. Otherwise, $X_t(1)$ is an independent random choice of ± 1. Hence, $\mathbf{E}[X_0(1)X_t(1)] = e^{-t}$ denotes the correlation between

X_0 and X_t. As $t \to \infty$, the correlation decays rapidly. We see that

$$\langle Q_t f, f\rangle = \mathbf{E}[f(X_0)f(X_t)] \ =: \ \mathrm{Stab}_{e^{-t}}(f)$$

where we have introduced the notion of Stability of f at parameter $\rho \in [0,1]$ as $\mathbf{E}[f(x)f(y)]$ where (x_i, y_i) are independent across i and $\mathbf{E}[x_i] = \mathbf{E}[y_i] = 0$ and $\mathbf{E}[x_i y_i] = \rho$. For a Boolean function, it denotes the probability that the value of f changes if each input bit is refreshed with probability $1-\rho$.

By Plancherel's theorem, we see that for any $t \ge 0$ and any $f \in L^2(\mathbb{Z}_2^n)$,

$$\|Q_t f\|_2^2 = \sum_S |\hat{f}(S)|^2 e^{-t|S|} \ \le \ \sum_S |\hat{f}(S)|^2 \ = \ \|f\|_2^2.$$

Thus, $\|Q_t\|_{2\to 2} = 1$ (equal, since $Q_t\mathbf{1} = \mathbf{1}$). In the following exercise, you are asked to show the same for any L^p.

Exercise 6. Show that $\|Q_t\|_{p\to p} = 1$ for all t and all $p \in [1,\infty]$.

2. Voting between two candidates

We think of $x \in \mathbb{Z}_2^n$ as a pattern for voting, where the jth voter votes for candidate $x_j = \pm 1$. The votes can be combined together in any reasonable or unreasonable way using a Boolean function $f : \mathbb{Z}_2^n \mapsto \{-1,1\}$ to arrive at a final decision. For example, the majority function (well-defined if n is odd) maps x to $+1$ if more voters vote for $+1$ than for -1. Taking $f(x) = x_1$ is a dictatorial vote where the first person decides, ignoring everyone else. If f depends only of $x_1, \ldots, x_k$, then the first k individuals form a politburo that decides, ignoring everyone else.

What are desirable properties of such a function? Here are some. Say that f is

(1) *unanimous* is $f(\mathbf{1}) = 1$ and $f(-\mathbf{1}) = -1$,

(2) *monotone* if $f(x) \le f(y)$ whenever $x_i \le y_i$ for all i,

(3) *odd* if $f(-x) = -f(x)$ for all $x \in \mathbb{Z}_2^n$,

(4) *symmetric* if $f(x_{\pi(1)}, \ldots, x_{\pi(n)}) = f(x_1, \ldots, x_n)$ for all $\pi \in \mathcal{S}_n$ and all $(x_1, \ldots, x_n) \in \mathbb{Z}_2^n$.

The majority function (when n is odd) has all these properties and is the only one to do so!

Proposition 7. *Let $f : \mathbb{Z}_2^n \mapsto \{-1,1\}$ be monotone, symmetric, and odd. Then n is odd and f is the majority function.*

Proof. Since f is monotone and symmetric, it follows that $f(x) \le f(y)$ if $\sum_i x_i \le \sum_i y_i$. Indeed, given x and y, by permuting co-ordinates bring to the form $(1, \ldots, 1, -1, \ldots, -1)$, and two such vectors are comparable co-ordinatewise. Therefore, $f(x) = \mathbf{1}_{x_1+\ldots+x_n \ge t}$ for some t. For this to be odd, we must take $t = 0$ and to avoid the possibility of $x_1 + \ldots + x_n = 0$, we must also take n to be odd. ■

Observe that we did not assume that f is unanimous. Indeed, if f is monotone and odd, and not a constant, then it must also be unanimous. Dictator functions and recursive majority functions are not symmetric, although they are monotone and odd (and hence also unanimous).

Here is another nice thing about majority function. It is the one that maximizes the expected number of happy voters! Here we say that a voter is happy if the final decision agrees with his/her own vote. Of course, to talk of expected number, we bring in the assumption of independent random voting. If H is the number of happy voters, then $\sum_{i=1}^{n} \mathbf{E}[x_i f(x_i)] = \mathbf{E}[2H - n]$, hence our claim is captured by the following proposition.

Proposition 8. *Let n be odd and let $f : \mathbb{Z}_2^n \mapsto \{-1, 1\}$. Then $\sum_{i=1}^{n} \mathbf{E}[x_i f(x)]$ is maximized uniquely when f is the majority function.*

Proof. Since $|f(x)| = 1$, we have

$$\sum_{i=1}^{n} \mathbf{E}[x_i f(x)] \;=\; \mathbf{E}[f(x)(x_1 + \ldots + x_n)] \;\leq\; \mathbf{E}[|x_1 + \ldots + x_n|]$$

with equality if and only if $f(x)(x_1 + \ldots + x_n) = |x_1 + \ldots + x_n|$. When n is odd, the sum of x_is is not zero, hence we must have $f(x) = \text{sgn}(x_1 + \ldots + x_n)$. This is the majority function. ■

In the following exercise, you get to show that in any monotone, symmetric voting scheme, no voter can have a large influence on the outcome.

Exercise 9. If $f : \mathbb{Z}_2^n \mapsto \{-1, 1\}$ is monotone, show that $\text{Inf}_i(f) = \hat{f}(\{i\})$ and that $\mathbf{E}[H] = \frac{1}{2}n + \frac{1}{2}\sum_{i=1}^{n} \hat{f}(\{i\})$. If in addition f is symmetric, show that $\text{Inf}_i(f) \leq \frac{1}{\sqrt{n}}$ for all i.

3. Voting between three candidates

A question of interest is how to decide if there are three candidates A, B, C? One approach is to have three 2-way elections, between each pair of candidates. Let us assume that the same function $f : \mathbb{Z}_2^n \mapsto \{-1, 1\}$ is used to decide each pairwise election. If it so happens that A beats B and A beats C, then A may be declared the winner unambiguously (the election between B and C only serves to find the runner-up). Such a winner is called a *Condorcet winner*. The problematic situation is that there may be no such candidate who beats the other two.

Example 10. If there are three voters with preferences $A > B > C$ and $B > C > A$ and $C > A > B$, then A beats B while B beats C and C beats A. There is no Condorcet winner.

Is this likely? Is there a clever choice of the function f that ensures there will be Condorcet winner?

Suppose there is a single voter who prefers A to B and B to C and C to A. In this silly situation, there is no way to have a decision. Let us assume that at least at an individual level, the voters have clear preference: Namely they have a strict order of preference between the three candidates.

Then irrespective of the number of voters, if $f(x) = x_1$, i.e., 1 is a dictator, then we always have a Condorcet winner (the first preference of 1). It turns out that dictator functions are the only ones that have a Condorcet winner for every voting pattern!

Theorem 11 (Arrow's theorem)**.** *Assume that each voter has an unambiguous order of preference for the three candidates. If $f : \mathbb{Z}_2^n \mapsto \{-1, 1\}$ has a Condorcet winner for any $x \in \mathbb{Z}_2^n$, then $f(x) = x_k$ or $f(x) = -x_k$ for some k (i.e., a dictator or anti-dictator function).*

The proof that we give is probabilistic: Assume that all voters choose one of six possible preferences (permutations of A, B, C at random. Then we show that there is a positive probability that there is no Condorcet winner, unless f is a dictator function.

To use the language of Boolean functions, let $x, y, z \in \mathbb{Z}_2^n$ denote the voting between A and B, between B and C and between C and A, respectively. Here $y_i = 1$ indicates that ith voter prefers B to C and $y_i = -1$ denotes the opposite preference. By assumption, $(x_i, y_i, z_i) \notin \{(1, 1, 1), (-1, -1, -1)\}$ (check that all other 6-tuples lead to unambiguous ordering of the candidates).

Kalai's proof of Arrow's theorem. Choose (x_i, y_i, z_i) uniformly at random from the set $\{-1, 1\}^3 \setminus \{(1, 1, 1), (-1, -1, -1)\}$. Make the choices independently for $1 \leq i \leq n$. Fix $f \in \mathbb{Z}_2^n$. The event that there is no Condorcet winner is the event that $f(x) = f(y) = f(z)$. For three bits $u, v, w \in \{-1, 1\}$, observe that $uv + vw + wu = 3$ if $u = v = w$ and -1 otherwise. Hence,

$$\mathbf{P}\{f(x) = f(y) = f(z)\} = \frac{1}{4}\mathbf{E}\left[1 + f(x)f(y) + f(y)f(z) + f(z)f(x)\right] \; = \; \frac{1}{4} + \frac{3}{4}\mathbf{E}[f(x)f(y)].$$

Writing $f(x) = \frac{1}{\sqrt{2^n}} \sum_{S \subseteq [n]} \hat{f}(S)\chi_S(x)$, we see that

$$\mathbf{E}[f(x)f(y)] = \frac{1}{2^n} \sum_{S,T \subseteq \mathbb{Z}_2^n} \hat{f}(S)\hat{f}(T)\mathbf{E}[\chi_S(x)\chi_T(y)].$$

Now,

$$\mathbf{E}[\chi_S(x)\chi_T(y)] = \mathbf{E}\left[\prod_{i \in S \setminus T} x_i \prod_{i \in T \setminus S} y_i \prod_{i \in S \cap T} x_i y_i\right] = \begin{cases} 0 & \text{if } S \neq T, \\ (-1/3)^{|S|} & \text{if } S = T, \end{cases}$$

since the expectation factors over i, and $\mathbf{E}[x_i y_i] = -\frac{1}{3}$ while $\mathbf{E}[x_i] = \mathbf{E}[y_i] = 0$. Thus we arrive at

$$\begin{aligned}
\mathbf{P}\{f(x) = f(y) = f(z)\} &= \frac{1}{4} + \frac{3}{4 \times 2^n} \sum_{S \subseteq [n]} |\hat{f}(S)|^2 \left(-\frac{1}{3}\right)^{|S|} \\
&\geq \frac{1}{4} - \frac{1}{3} \times \frac{3}{4 \times 2^n} \sum_{S \subseteq [n]} |\hat{f}(S)|^2 \\
&= \frac{1}{4} - \frac{1}{4 \times 2^n} \sum_{x \in \mathbb{Z}_2^n} |f(x)|^2 \\
&= 0
\end{aligned}$$

since $f(x) = \pm 1$ for all x. This shows that $\mathbf{P}\{f(x) = f(y) = f(z)\} \geq 0$ (of course!) and that this probability is strictly positive unless $\hat{f}(S) = 0$ except for $|S| = 1$ (see the inequality step above).

In that case, we must have $f(x) = a_1x_1 + \ldots + a_nx_n$ for some $a_i \in \mathbb{C}$. If more than one $a_i \neq 0$, say a_1, a_2 are non-zero, then $a_1x_1 + a_2x_2$ can take the values $-a_1 - a_2$, $a_1 - a_2$, $a_2 - a_1$, $a_1 + a_2$ of which at least three are distinct (why?). which means that f takes at least three distinct values, contradicting that f is a Boolean function. Hence, we must have at most one non-zero coefficient, and if that is the k-th one, then $f(x) = x_k$ or $f(x) = -x_k$. ■

Observe that the proof shows that the probability to not have a Condorcet winner is equal to

$$\frac{1}{4} + \frac{3}{4 \times 2^n} \sum_{S \subseteq [n]} |\hat{f}(S)|^2 \left(-\frac{1}{3}\right)^{|S|}. \tag{2}$$

which is what we earlier denoted as $\frac{1}{4}(1 + 3\text{Stab}_{-\frac{1}{3}}(f))$. That is, the above quantity is equal to

$$\frac{1}{4} + \frac{3}{4}\mathbf{P}\{f(x) = f(y)\}$$

where x, y are uniformly random on $\mathbb{Z}_2^n$ with $\mathbf{E}[x_iy_i] = -\frac{1}{3}$. One way to generate them is to throw a fair die for each i, and set

$$x_i = \begin{cases} +1 & \text{if throw is } 1, 2, 3, \\ -1 & \text{if throw is } 4, 5, 6, \end{cases} \qquad y_i = \begin{cases} +1 & \text{if throw is } 1, 4, 5, \\ -1 & \text{if throw is } 2, 3, 6. \end{cases}$$

Exercise 12. For the majority function, *numerically* compute the probability that there is no Condorcet winner among three candidates and check that happens as $n \to \infty$, the probability converges to a non-zero number.

Extra: Can you show that in fact the probability converges to $1 - \frac{3}{2\pi}\arccos\left(-\frac{1}{3}\right) = 0.0877\ldots$? One way is to use the interpretation in terms of $\text{Stab}_{-\frac{1}{3}}(f)$ together with the central limit theorem, preferably in a justifiable manner.

The key point of the exercise above is that the probability does not approach zero as $n \to \infty$. Otherwise, we would ignore the positive but small probability of the unpleasant outcome, when the population size is large. The next section continues this discussion.

4. Robust version of Arrow's theorem

Arrow's theorem shows that if we demand a Condorcet winner 100% of the time, then we must choose bad systems like a dictator function. A question relevant to real applications would be: if we allow for a small positive probability of not getting a Condorcet winner, can we perhaps use one of the nicer systems like the Majority function or something reasonably equitable? Turns out, no! The following result is a robust or stability version of Arrow's theorem.

Theorem 13 (Kalai; Friedgut–Kalai–Naor)**.** *Suppose that for a Boolean function $f : \mathbb{Z}_2^n \mapsto \{-1, 1\}$, the probability of not getting a Condorcet winner in a 3-way election is at most δ. Then, there is some set $E \subseteq \mathbb{Z}_2^n$ with $|E| \leq C\delta 2^n$, and $k \in [n]$ such that on $\mathbb{Z}_2^n \setminus E$ we have either $f = \chi_{\{k\}}$ or $f = -\chi_{\{k\}}$.*

Here C is a constant, not depending on any parameters. The statement looks quite strong, since it says that f is actually equal to $\pm x_k$ on a set of probability close to 1. However, since f and x_k are Boolean, this is equivalent to measuring the distance in other ways. Indeed, for any two Boolean functions f, g, we have

$$|\{x \in \mathbb{Z}_2^n : f(x) \neq g(x)\}| = \frac{1}{4} \sum_{x \in \mathbb{Z}_2^n} |f(x) - g(x)|^2 \;=\; \frac{1}{4}\|f - g\|_2^2.$$

Proof. Taking inspiration from the proof of Arrow's theorem, set $g = 2^{-n/2} \sum_{k=1}^n \hat{f}(\{k\})\chi_{\{k\}}$. We show that f and g are close, and then that g (or $-g$) is close to one of the characters $\chi_{\{k\}}$.

Step 1: We show that f and g are close. From (2),

$$\begin{aligned}
\mathbf{P}\{\text{no Condorcet winner}\} &= \frac{1}{4} + \frac{3}{4 \times 2^n} \sum_{S \subseteq [n]} |\hat{f}(S)|^2 \left(-\frac{1}{3}\right)^{|S|} \\
&\geq \frac{1}{4} - \frac{1}{4 \times 2^n} \sum_{k=1}^n |\hat{f}(k)|^2 - \frac{1}{9 \times 4 \times 2^n} \sum_{3 \leq |S| \text{ odd}} |\hat{f}(S)|^2 \\
&\geq \frac{1}{4} - \frac{1}{4 \times 2^n} \sum_{k=1}^n |\hat{f}(k)|^2 - \frac{1}{9 \times 4} \left(1 - \frac{1}{2^n} \sum_{k=1}^n |\hat{f}(k)|^2\right) \\
&= \frac{2}{9} \left(1 - \frac{1}{2^n} \sum_{k=1}^n |\hat{f}(k)|^2\right).
\end{aligned}$$

Hence if this probability is at most δ, we get $\frac{1}{2^n} \sum_{k=1}^n |\hat{f}(k)|^2 \geq 1 - 5\delta$. Then it also follows that

$$\frac{1}{2^n}\|f - g\|_2^2 = \frac{1}{2^n} \sum_{|S| \neq 1} |\hat{f}(S)|^2 \leq 5\delta$$

and $\|g\|_2 \geq 1 - \sqrt{5\delta}$.

Step 2: We show that g^2 has small variance. This is because of hypercontractivity. More precisely, we use Exercise 17 in the next section, we get

$$\mathbf{E}[(g^2 - 1)^4] \leq (9\mathbf{E}[(g^2 - 1)^2])^2.$$

Observe that $\|g\|_2^2 \leq \|f\|_2^2 = 2^n$, hence $\mathbf{P}\{|g| > M\} \leq \frac{1}{18^2}$ if $M \geq 18$. Then,

$$\begin{aligned}
\mathbf{E}[(g^2 - 1)^2] &= \mathbf{E}[(g^2 - 1)^2 \mathbf{1}_{|g| \leq M}] + \mathbf{E}[(g^2 - 1)^2 \mathbf{1}_{|g| > M}] \\
&\leq (M + 1)^2 \mathbf{E}[|g - f|^2] + \sqrt{\mathbf{E}[(g^2 - 1)^4]} \sqrt{\mathbf{P}\{|g| > M\}}.
\end{aligned}$$

In the second summand we used Cauchy-Schwarz while in the first we wrote $g^2 - 1 = (g - 1)(g + 1)$ and used that $|g + 1| \leq M + 1$ and $|g - 1| \leq |g - f|$ (since f is Boolean).By the choice of M and the Bonami lemma conclusion obtained above, the second summand is bounded by $\frac{1}{2}\mathbf{E}[(g^2-1)^2]$. Transfer this to the other side to get

$$\begin{aligned}\mathbf{E}[(g^2-1)^2] &\le 2(M+1)^2\mathbf{E}[|g-f|^2]\\ &\le 10(M+1)^2\delta\end{aligned}$$

by invoking Step-1.

Step 3: Now we use the Fourier expansion of g,

$$g^2 = \left(\frac{1}{2^n}\sum_{i=1}^n |\hat{f}(\{i\})|^2\right) + \sum_{i<j}\frac{2\hat{f}(\{i\})\hat{f}(\{j\})}{2^n}\chi_{\{i,j\}}$$

to see that

$$\begin{aligned}\mathrm{Var}(g^2) &= \frac{1}{2^n}\sum_{i<j}\frac{4\hat{f}(\{i\})^2\hat{f}(\{j\})^2}{2^{2n}}\\ &= 2\left(1-\frac{1}{2^{2n}}\sum_{i=1}^n |\hat{f}(\{i\})|^4\right).\end{aligned}$$

Now we observe that $\mathrm{Var}(g^2) \le \mathbf{E}[(g^2-1)^2]$ which is bounded by $C\delta$ by Step-2, and that

$$\sum_{i=1}^n |\hat{f}(\{i\})|^4 \le \left(\max_i |\hat{f}(\{i\})|^2\right)\sum_{i=1}^n |\hat{f}(\{i\})|^2 = 2^n \max_i |\hat{f}(\{i\})|^2.$$

This leads to

$$\max_i \left|\frac{1}{2^{n/2}}\hat{f}(\{i\})\right|^2 \ge 1-\frac{1}{2}C\delta.$$

If k is the index that attains this maximum, then it follows that $\|f \pm \chi_{\{k\}}\|_2^2 \le C'\delta$. ■

5. Hypercontractivity

An exercise we gave earlier said that $\|Q_t\|_{p\to p} = 1$ for $1 \le p \le \infty$. It turns out that something more is true. For each $1 \le p \le q \le \infty$, we have $\|Q_t\|_{p\to q} = 2^{\frac{n}{q}-\frac{n}{p}}$, for sufficiently large t (how small depends on p and q). This property is called *hypercontractivity*. There are other abstract definitions of what it means for a random variable to be hypercontractive, but we limit ourselves to this.

The inequalities can be written more cleanly if we use the p-norms with respect to the uniform probability distribution on $\mathbb{Z}_2^n$ instead of the counting measure. Let $[\![f]\!]_p = \|f\|_p 2^{-n/p}$ denote this new norm.

Theorem 14 (The hypercontractivity theorem (Aline Bonami; Nelson, Gross))**.** *Let $1 \le p \le q \le \infty$. Then for $e^{-t} \le \frac{\sqrt{p-1}}{\sqrt{q-1}}$, we have $[\![Q_t]\!]_{p\to q} \le 1$, i.e., $[\![Q_t f]\!]_q \le [\![f]\!]_p$ for all $f : \mathbb{Z}_2^n \mapsto \mathbb{C}$.*

Why should it hold for large t but not small? First observe that $[\![f]\!]_p$ is increasing in p, for any f, and the increase is strict unless f is constant. Now consider the extreme cases:

(1) $t = 0$. Then $Q_0 f = f$, hence $[\![f]\!]_q \le [\![f]\!]_p$ holds in general only if $q \le p$.

(2) $t = \infty$. Then $Q_\infty f = 2^{-n/2}\hat{f}(\emptyset)\mathbf{1}$ is a constant function and hence its p-norm is the same for all p. Hence $\llbracket Q_\infty f \rrbracket_q \le \llbracket f \rrbracket_p$ holds for any p, q.

The hypercontractivity theorem gives something intermediate. In fact, turning the statement around, we see that for any finite t and any $1 \le p < \infty$, there is some $q > p$ (to be precise, any $q \le 1+(p-1)e^t$ will do) such that $\llbracket Q_t f \rrbracket_q \le \llbracket f \rrbracket_p$ for all f.

We can state this in another way: Remember that $\langle Q_t f, g\rangle = \mathbf{E}[f(x)g(y)]$ where x, y are e^{-t}-correlated (i.e., (x_i, y_i) are independent, $\mathbf{E}[x_i] = \mathbf{E}[y_i] = 0$ and $\mathbf{E}[x_i y_i] = e^{-t}$). Hence, for $q \le 1 + (p-1)e^t$ and q' its conjugate ($\frac{1}{q} + \frac{1}{q'} = 1$), we have

$$\mathbf{E}[f(x)g(y)] \le \llbracket Q_t f \rrbracket_q \llbracket g \rrbracket_{q'} \le \llbracket f \rrbracket_p \llbracket g \rrbracket_{q'}.$$

For $t = 0$, we have $x = y$ and $q = p$ and the inequality above is just Hölder's inequality $\llbracket fg \rrbracket_1 \le \llbracket f \rrbracket_p \llbracket g \rrbracket_q$. For $t = \infty$, we see that x and y are independent and $q = \infty$, and hence the inequality above says $\llbracket f \otimes g \rrbracket_1 \le \llbracket f \rrbracket_1 \llbracket g \rrbracket_1$. The inequality above gives an intermediate conclusion, when the correlation between x and y is neither 0 nor 1. To see that it is a strengthening of Hölder's inequality, observe that the latter would have given $\llbracket f \rrbracket_p \llbracket g \rrbracket_{p'}$ on the right, but $q' < p'$, hence the quantity $\llbracket f \rrbracket_p \llbracket g \rrbracket_{q'}$ is smaller.

We prove only very special cases of this inequality.

Proof of hypercontractivity for $p = 2$, $q = 4$. Write $f(x) = x_n g(x) + h(x)$ where

$$g(x) = \frac{1}{2^{n/2}} \sum_{S \subseteq [n-1]} \hat{f}(S \cup \{n\})\chi_S(x), \qquad h(x) = \frac{1}{2^{n/2}} \sum_{S \subseteq [n-1]} \hat{f}(S)\chi_S(x)$$

are functions of $x_1, \dots, x_{n-1}$. Let $\rho = e^{-t}$ and observe that $Q_t f(x) = \rho x_n Q_t g(x) + Q_t h(x)$. We use induction on n. Hence, assume that $\llbracket Q_t g \rrbracket_4 \le \llbracket |g \rrbracket_2$ and $\llbracket Q_t h \rrbracket_4 \le \rho \llbracket h \rrbracket_2$. Then,

$$\begin{aligned}
(Q_t f(x))^2 &= \rho^2 (Q_t g(x))^2 + (Q_t h(x))^2 + 2\rho x_n (Q_t g(x))(Q_t h(x)), \\
(Q_t f(x))^4 &= \rho^2 (Q_t g(x))^4 + (Q_t h(x))^4 + 6\rho^2 (Q_t g(x))^2 (Q_t h(x))^2 \\
&\quad + 4x_n (Q_t g(x))^3 (Q_t h(x)) + 4x_n (Q_t h(x))^3 (Q_t g(x)),
\end{aligned}$$

where we used $x_n^2 = 1$. When we sum over $x \in \mathbb{Z}_2^n$, since $Q_t g$ and $Q_t h$ depend only on $x_1, \dots, x_{n-1}$, the sum over x_n factors away and hence all terms with x_n factor (the last term in the first line and the last two terms in the second) vanish. We are left with

$$\begin{aligned}
\llbracket Q_t f \rrbracket_2^2 &= \rho^2 \llbracket Q_t g \rrbracket_2^2 + \llbracket Q_t h \rrbracket_2^2, \\
\llbracket Q_t f \rrbracket_4^4 &= \rho^4 \llbracket Q_t g \rrbracket_4^4 + \llbracket Q_t h \rrbracket_4^4 + 6\rho^2 \llbracket (Q_t g)^2 \cdot (Q_t h)^2 \rrbracket_1 \\
&\le \rho^4 \llbracket Q_t g \rrbracket_4^4 + \llbracket Q_t h \rrbracket_4^4 + 6\rho^2 \llbracket Q_t g \rrbracket_4^2 \llbracket Q_t h \rrbracket_4^2 \ \text{(Cauchy-Schwarz)} \\
&\le \rho^4 \llbracket Q_t g \rrbracket_2^4 + \llbracket Q_t h \rrbracket_2^4 + 6\rho^2 \llbracket Q_t g \rrbracket_2^2 \llbracket Q_t h \rrbracket_2^2 \ \text{(inductive hypothesis).}
\end{aligned}$$

If $6\rho^2 \leq 2$, then $\rho \leq 1$ and the last quantity is bounded from above by

$$[\![Q_t g]\!]_2^4 + [\![Q_t h]\!]_2^4 + 2[\![Q_t g]\!]_2^2[\![Q_t h]\!]_2^2 = [\![Q_t f]\!]_2^4.$$

Thus the induction closes and we get the conclusion $[\![Q_t f]\!]_4 \leq [\![f]\!]_2$, for $\rho \leq \frac{1}{\sqrt{3}}$.

Of course, it remains to check the base case $n = 1$. In this case, we have $f(x) = a + bx$ and $Q_t f(x) = a + b\rho x$ for $x \in \mathbb{Z}_2$, and direct calculation gives

$$\begin{aligned} [\![f]\!]_2^4 = [(a+b)^2 + (a-b)^2]^2 &= 4a^4 + 4b^4 + 8a^2b^2, \\ [\![Q_t f]\!]_4^4 = (a+b\rho)^4 + (a-b\rho)^4 &= 2a^4 + 2\rho^4 b^4 + 12\rho^2 a^2 b^2. \end{aligned}$$

Therefore, if $12\rho^2 \leq 8$ (certainly satisfied if $\rho \leq \frac{1}{\sqrt{3}}$), then by comparing coefficients we see that $[\![Q_t f]\!]_4 \leq [\![f]\!]_2$. ■

Exercise 15. Use the fact that $[\![Q_t]\!]_{2\to 4} \leq 1$ for $e^{-t} \leq \frac{1}{\sqrt{3}}$ to show that $[\![Q_t]\!]_{\frac{4}{3}\to 2} \leq 1$ for the same values of t.

Here is a useful corollary.

Corollary 16. *If $f : \mathbb{Z}_2^n \mapsto \mathbb{C}$ is homogenous of degree k (i.e., $\hat{f}(S) = 0$ unless $|S| = k$), then $[\![f]\!]_4 \leq 3^{\frac{k}{2}}[\![f]\!]_2$.*

Proof. Immediate from the $(2,4)$-hypercontractivity, since $Q_t f = \rho^k f$ if $e^{-t} = \rho$. ■

But in fact, the same conclusion holds even without the assumption of homogeneity.

Exercise 17. [Bonami's lemma] Show that if f has degree at most k (i.e., $\hat{f}(S) = 0$ if $|S| > k$), then $[\![f]\!]_4 \leq 3^{\frac{k}{2}}[\![f]\!]_2$. [*Hint:* Imitate the proof of the $(2,4)$-hypercontractivity theorem. If you can actually deduce this exercise from that theorem, I would like to know how that can be done.]

5.1. **Sahasranand's proof of Bonami's lemma from** $(2,4)$**-hypercontractivity.** As we have seen, for a homogenous polynomial f, the inequality $[\![f]\!]_4 \leq \rho^k[\![f]\!]_2$ for $\rho = 1/\sqrt{3}$ follows from the $(2,4)$-hypercontractivity. For general f of degree at most k, one is tempted to decompose it as $f = f_0 + \ldots + f_k$, where $f_i = \frac{1}{2^{n/2}}\sum_{|S|=i} \hat{f}(S)\chi_S$ are the homogenous components of f. Then with $e^{-t} = \rho$, we have $[\![Q_t f_i]\!]_4 \leq [\![f_i]\!]_2$ for each i and $Q_t f = Q_t f_0 + \ldots + Q_t f_k$. But it is not clear how to relate $[\![f]\!]_4$ with $[\![Q_t f]\!]_4$ and the individual $[\![Q_t f_i]\!]_4$ (all sorts of "cross terms" enter). Sahasranand's[7] trick is to notice that such difficulties disappear on the 2-norm side, by orthogonality.

[7] Sahasranand Kodinthirapully Ramanadhan, student of ECE department, IISC.

To use this observation, let $g_i = \rho^{-i} f_i$ and $g = g_0 + \ldots + g_k$. Then $Q_t g_i = f_i$ and $Q_t g = f$. By the $(2,4)$-hypercontractivity we have $[\![f]\!]_4 = [\![Q_t g]\!]_4 \leq [\![g]\!]_2$. On the other hand, g_i are orthogonal, hence

$$\begin{aligned}
[\![f]\!]_4^2 \ \leq \ [\![g]\!]_2^2 &= [\![g_0]\!]_2^2 + \ldots + [\![g_k]\!]_2^2 \\
&= [\![f_0]\!]_2^2 + \rho^{-2}[\![f_1]\!]_2^2 + \ldots + \rho^{-2k}[\![f_k]\!]_2^2 \\
&\leq \rho^{-2k}\left([\![f_0]\!]_2^2 + [\![f_1]\!]_2^2 + \ldots + [\![f_k]\!]_2^2\right) \\
&= \rho^{-2k}[\![f]\!]_2^2
\end{aligned}$$

by the orthogonality of the f_is. As $\rho = 1/\sqrt{3}$, we have proved that $[\![f]\!]_4 \leq 3^{k/2}[\![f]\!]_2$. ■

CHAPTER 4

Dirichlet's theorem on primes in arithmetic progressions

This is a famous theorem of Dirichlet, often referred to as the starting point of analytic number theory[8].

1. The theorem

Theorem 1 (Dirichlet)**.** *Let a, d be co-prime natural numbers. Then the arithmetic progression $a, a+d, a+2d, \ldots$ contains infinitely many prime numbers.*

It is clear that the condition of a, d being co-prime is necessary, otherwise, there is at most one prime number in the sequence. For special cases, it is possible to prove this theorem along elementary lines like that of Euclid's proof that there are infinitely many prime numbers.

Example 2. Let $a = 3$ and $d = 4$, so the sequence is the set of number that are $-1 \pmod 4$. If $p_1, \ldots, p_k$ are all the prime numbers in this sequence, then $4p_1 \ldots p_k - 1$ is not divisible by any of them, and must have a prime factor other than $2, p_1, \ldots, p_k$. If all its prime factors were of the form $4j + 1$, their product would also be $1 \pmod 4$, but the number constructed is $-1 \pmod 4$.

In general, apparently there are no such proofs[9] and the proof of Dirichlet's theorem is not analogous to Euclid's proof. It is much closer to Euler's proof of divergence of $\sum_p \frac{1}{p}$, and indeed, the proof proceeds by showing that $\sum_{p \equiv a (\text{mod } d)} \frac{1}{p} = \infty$, or what is the same, $\sum_{p \equiv a (\text{mod } d)} \frac{1}{p^s} \to \infty$ as $s \downarrow 1$.

2. Euler's proof that $\sum \frac{1}{p} = \infty$

We ignore convergence issues (which will be justified later in greater generality) and recall this proof. The starting point is Euler's product formula, valid for $s > 1$ (one can also take complex s with $\text{Re}(s) > 1$, but we don't need it here)

$$\zeta(s) := \sum_{n=1}^{\infty} \frac{1}{n^s} = \prod_p \left(1 + \frac{1}{p^s} + \frac{1}{p^{2s}} + \ldots\right) = \prod_p \frac{1}{1 - \frac{1}{p^s}}.$$

[8] Many books have the proof. Serre's *A course in arithmetic*, Apostol's *Introduction to analytic number theory* and Stein and Shakarchi's *Fourier analysis: an introduction*, all have superb expositions. Because of this, our notes will be quite brief.

[9] But there are "elementary proofs" such as one by Selberg.

The last equality is clear, and the previous one is an expression of the fundamental theorem of arithmetic that every natural number other than 1 can be written as a product of prime powers in a unique way.

Now take logarithm of both sides to write $\log \zeta(s) = -\sum_p \log(1 - \frac{1}{p^s})$. Since $\log(1-x) = -x - \frac{1}{2}x^2 - \frac{1}{3}x^3 - \dots$ for $|x| < 1$, we can write

$$\log \zeta(s) = \sum_p \frac{1}{p^s} + \frac{1}{2}\sum_p \frac{1}{p^{2s}} + \frac{1}{3}\sum_p \frac{1}{p^{3s}} + \dots$$

For $k \geq 2$ and $s > 1$, we see that

$$\sum_p \frac{1}{p^{ks}} \leq \sum_{n\geq 2} \frac{1}{n^k} \leq \int_2^\infty \frac{1}{x^k} dx = \frac{1}{(k-1)2^k}.$$

Thus

$$\frac{1}{2}\sum_p \frac{1}{p^{2s}} + \frac{1}{3}\sum_p \frac{1}{p^{3s}} + \dots \leq \sum_{k\geq 2} \frac{1}{k(k-1)2^k}.$$

Therefore, $\log \zeta(s) = \sum_p \frac{1}{p^s} + O(1)$ as $s \downarrow 1$. But $\log \zeta(s) \to \infty$ as $s \downarrow 1$, since $\sum_n \frac{1}{n}$ diverges. Consequently,

$$\lim_{s\downarrow 1} \sum_p \frac{1}{p^s} = \infty.$$

As the left side is bounded from above by $\sum_p \frac{1}{p}$, it follows that $\sum_p \frac{1}{p} = \infty$.

3. Dirichlet L-functions

Let $\mathbb{Z}_d^* = \{0 \leq j \leq d-1 : (j,d) = 1\}$. This is a group under multiplication modulo d, and the cardinality of this group is denoted $\varphi(d)$. If $\chi \in \widehat{\mathbb{Z}_d^*}$, then we extend it to a function on $\mathbb{N}$ by setting $\chi(n) = 0$ if $(n,d) \neq 1$ and $\chi(n) = \chi(\bar{n})$ if $(n,d) = 1$. Here $\bar{n}$ is the residue of n modulo d. Of course, $|\chi(n)| \leq 1$.

Exercise 3. Explicitly find all characters of $\mathbb{Z}_d^*$ for $d = 6, 8$.

To such a character, we associate the *Dirichlet* L-function

$$L_\chi(s) = \sum_{n\geq 1} \frac{\chi(n)}{n^s}$$

which is clearly absolutely convergent for $s > 1$.

The trivial character: Let χ_0 denote the trivial character of $\mathbb{Z}_d^*$. Note that when extended as a function on $\mathbb{N}$, it is not identically 1, but $\chi_0(n) = \mathbf{1}_{(n,d)=1}$. Hence, if $p_1, \dots, p_k$ are the distinct primes that divide d, then by inclusion-exclusion,

$$\begin{aligned} L_{\chi_0}(s) &= \sum_{n\geq 1}\frac{1}{n^s} - \sum_{i=1}^{k}\sum_{m\geq 1}\frac{1}{(mp_i)^s} + \sum_{i<j\leq k}\sum_{m\geq 1}\frac{1}{(mp_ip_j)^s} - \dots \\ &= \zeta(s)\prod_{j=1}^{k}(1-\frac{1}{p_j^s}). \end{aligned}$$

Hence, we also see that $L_{\chi_0}(s)\uparrow\infty$ as $s\downarrow 1$. We can make this more precise. Observe that

$$\zeta(s) - \frac{1}{s-1} = \sum_{n\geq 1}\int_n^{n+1}\left(\frac{1}{n^s}-\frac{1}{x^s}\right)dx \leq \sum_{n\geq 1}\frac{s}{n^{s+1}}$$

as $n^{-s}-x^{-s}=(x-n)\frac{s}{t^{s+1}}$ by the intermediate value theorem, for some $t\in(n,x)$. The last series converegs uniformly on (δ,∞) for any $\delta>0$. This shows that

(1) $\zeta(s)=\frac{1}{s-1}+\tilde{\zeta}(s)$ where $\tilde{\zeta}$ is continuous on $(0,\infty)$.

(2) $\zeta(s)\sim\frac{1}{s-1}$ as $s\downarrow 1$.

(3) $L_{\chi_0}(s)\sim\frac{\prod_{j=1}^{k}(1-p_j^{-1})}{s-1}$ as $s\downarrow 1$.

Non-trivial characters: Now suppose $\chi\neq\chi_0$. Then L_χ converges for $s>0$ and is continuous there. This follows from the following more general lemma, as χ is a d-periodic sequence and $\sum_{j=1}^{d}\chi(j)=\langle\chi,\chi_0\rangle_{L^2(\mathbb{Z}_d^*)}=0$ (which shows that the partial sums of χ take at most d distinct values).

Lemma 4. *Let $a:\mathbb{N}\mapsto\mathbb{C}$. Assume that its partial sums $A_n=a_1+\ldots+a_n$ are uniformly bounded. Then $\sum_n a_n n^{-s}$ is convergent for $s>0$ and uniformly convergent for $s>\delta$ for any $\delta>0$.*

Proof. Let $k<\ell$ and consider

$$\begin{aligned} \sum_{n=k}^{\ell}\frac{a_n}{n^s} &= \sum_{n=k}^{\ell}\frac{A_n-A_{n-1}}{n^s} \\ &= \sum_{n=k}^{\ell-1}A_n\left(\frac{1}{n^s}-\frac{1}{(n+1)^s}\right)+\frac{A_\ell}{\ell^s}-\frac{A_{k-1}}{k^s}. \end{aligned}$$

Hence if $|A_n|\leq M$, then

$$\begin{aligned} |\sum_{n=k}^{\ell}\frac{a_n}{n^s}| &\leq M\sum_{n=k}^{\ell-1}\left(\frac{1}{n^s}-\frac{1}{(n+1)^s}\right)+\frac{M}{\ell^s}+\frac{M}{k^s} \\ &\leq \frac{3M}{k^s}. \end{aligned}$$

Thus if k is large enough, then all partial sums of the above form are small. By Cauchy criterion, the proof is complete. ■

Exercise 5. If $(a_n)_n$ is bounded, then show that $s \mapsto \sum_{n\geq 1} a_n n^{-s}$ is smooth on $(1,\infty)$. If partial sums of a_n are bounded, show that $s \mapsto \sum_{n\geq 1} a_n n^{-s}$ is smooth on $(0,\infty)$.

4. Product formulas

Now let us state a more general lemma that we prove rigorously.

Lemma 6. *Let $a : \mathbb{N} \mapsto \mathbb{C}$ be a completely multiplicative function, i.e., $a(mn) = a(m)a(n)$ for all $m, n \in \mathbb{N}$. Assume that a is uniformly bounded. Then, for $s > 1$ we have*

$$\sum_n \frac{a(n)}{n^s} = \prod_p \frac{1}{1-\frac{a(p)}{p^s}}.$$

Proof. Order the primes as $p_1 < p_2 < \dots$. Consider the finite product

$$\prod_{j=1}^k \frac{1}{1-\frac{a(p_j)}{p_j^s}} = \prod_{j=1}^k \left(1 + \frac{a(p_j)}{p_j^s} + \frac{a(p_j)^2}{p_j^{2s}} + \dots\right)$$
$$= 1 + \sum_{m_1,\dots,m_k \geq 0} \frac{a(p_1)^{m_1}\dots a(p_k)^{m_k}}{p_1^{sm_1}\dots p_k^{sm_k}}$$

is justified. Implicitly there is a rearrangement of terms here, which is okay because of absolute convergence. By the fundamental theorem of arithmetic, $p_1^{m_1}\dots p_k^{m_k}$ cover all numbers up to p_k. Hence as $k \to \infty$, the last sum converges to $\sum_n \frac{a_n}{n^s}$. Therefore the limit of the left side must also exist. ■

Corollary 7. *For any $d \geq 1$ and any $\chi \in \widehat{\mathbb{Z}_d^*}$, and for $s > 1$,*

$$L_\chi(s) = \prod_p \left(1 - \frac{\chi(p)}{p^s}\right)^{-1}.$$

5. Logarithm

For $z \in \mathbb{C}$ with $|z-1| < 1$, define

$$\log z := -\sum_{k\geq 1} \frac{1}{k}(z-1)^k.$$

Exercise 8. Show that $e^{\log z} = z$ whenever $|z-1| < 1$. Show that for $|w| < \frac{1}{2}$, both $\log(1-w)$ and $\log\frac{1}{1-w}$ are both well-defined and negatives of each other.

If $\chi \in \widehat{\mathbb{Z}_d^*}$, then for any prime p and any $s > 1$, we have $|\chi(p)p^{-s}| \leq \frac{1}{2}$. Hence by the above exercise,

$$\exp\left\{\sum_p \log\left(1-\frac{\chi(p)}{p^s}\right)^{-1}\right\} = \prod_p \exp\left\{\log\left(1-\frac{\chi(p)}{p^s}\right)^{-1}\right\} = \prod_p \left(1-\frac{\chi(p)}{p^s}\right)^{-1} = L_\chi(s).$$

But the exponent on the left is

$$\begin{aligned}\sum_p \log\left(1-\frac{\chi(p)}{p^s}\right)^{-1} &= -\sum_p \log\left(1-\frac{\chi(p)}{p^s}\right)\\ &= \sum_p \sum_{m\geq 1} \frac{\chi(p)^m}{mp^{ms}}\\ &= \sum_p \frac{\chi(p)}{p^s} + \sum_p \sum_{m\geq 2} \frac{\chi(p)^m}{mp^{ms}}.\end{aligned}$$

The second summand can be bounded in absolute value by

$$\sum_p \sum_{m\geq 2} \frac{1}{p^{ms}} = \sum_p \frac{1}{p^s(p^s-1)} \leq \sum_p \frac{2}{p^2}.$$

Thus,

$$\sum_p \log\left(1-\frac{\chi(p)}{p^s}\right)^{-1} = \sum_p \frac{\chi(p)}{p^s} + O(1)$$

as $s \downarrow 1$. Here the $O(1)$ is uniform over $s > 1$. Hence

$$L_\chi(s) = \exp\left\{\sum_p \frac{\chi(p)}{p^s} + O(1)\right\} \qquad \text{as } s \downarrow 1.$$

Therefore, we see that the following are equivalent[10]

(1) $\sum_p \frac{\chi(p)}{p^s}$ stays bounded as $s \downarrow 1$.

(2) $L_\chi(s)$ does not converge to 0 or ∞ as $s \downarrow 1$.

For the trivial character, the second condition does not hold, in fact we have seen that $L_{\chi_0}(s) \sim C/(s-1)$ as $s \downarrow 1$, hence

$$\sum_p \frac{\chi_0(p)}{p^s} = \infty.$$

As $\chi_0(p) = 1$ for all but the finitely many primes that divide d, this implies (with all the rigour added in) Euler's theorem that $\sum_p \frac{1}{p}$ diverges.

For non-trivial characters, we know that $L_\chi(s) \to L_\chi(1)$, a finite number as $s \downarrow 1$. Hence there is no divergence to infinity. The following Lemma is crucial.

Lemma 9. *If $\chi \neq \chi_0$, then $L_\chi(1) \neq 0$.*

Assuming the lemma, we see that $\sum_p \frac{\chi(p)}{p^s}$ stays bounded as $s \downarrow 1$, for any $\chi \neq \chi_0$.

[10] At this point, we have simplified the presentation of Stein and Shakarchi a bit (which makes one suspicious if we are making a mistake!). In their book, they further define the logarithm of the $L_\chi(s)$, show that that definition agrees with the sum of logarithms of $1-\chi(p)p^{-s}$, and then take logarithm on both sides of the product formula. We seem to be able to avoid some of these steps.

Remark 10. That the number theoretic question of divergence of $\sum_{\bar{p}=a} \frac{1}{p}$ is reduced to the vanishing or non-vanishing of certain continuous (even analytic) functions at certain points is marvellous! Now such connections have become routine, but Dirichlet was the first. Later, the prime number theorem, that says that the number of primes less than x is asymptotic to $\log x/x$, was derived by showing that $\zeta(s)$ does not vanish anywhere on the line $\operatorname{Im} s = 1$.

6. Proof of Dirichlet's theorem

Now we can put together all the ingredients to get a proof of Dirichlet's theorem. Fix d, a such that $(a, d) = 1$ and consider the Fourier expansion of $\mathbf{1}_{\{a\}}$ in $\mathbb{Z}_d^*$ (and then extend to all integers):

$$\mathbf{1}_{n=a \pmod d} = \frac{1}{\varphi(d)} \sum_{\chi \in \widehat{\mathbb{Z}_d^*}} \overline{\chi(a)}\chi(n).$$

Hence,

$$\sum_{p:n=a \pmod d} \frac{1}{p^s} = \frac{1}{\varphi(d)} \sum_{\chi \in \widehat{\mathbb{Z}_d^*}} \overline{\chi(a)} \sum_{p:n=a \pmod d} \frac{\chi(p)}{p^s}.$$

On the right, only the summand with $\chi = \chi_0$ has a series that blows up as $s \downarrow 1$. For all other χ, the series stays bounded as $s \downarrow 1$. Therefore,

$$\sum_{p:n=a \pmod d} \frac{1}{p^s} \to \infty$$

as $s \downarrow 1$. This shows that there must be infinitely many primes that are congruent to a modulo d. In fact, we get more information: Letting $\bar{p}$ denote the image of p in $\mathbb{Z}_d^*$,

$$\sum_{p:\bar{p}=a} \frac{1}{p^s} \sim \frac{1}{\varphi(d)} \sum_{p} \frac{\chi_0(p)}{p^s} \sim \frac{1}{\varphi(d)} \log \frac{1}{s-1}.$$

The last asymptotic follows from Euler's proof, the behaviour of $\zeta(s)$ near $s = 1$, and the fact that $\chi_0(p) = 1$ for all but finitely many primes.

Exercise 11. Show that $\sum\limits_{\bar{p}=a,\ p\le x} \frac{\log p}{p} \sim \frac{1}{\varphi(d)} \sum\limits_{p \le x} \frac{\log p}{p}$, as $x \to \infty$.

What one would like is to show that asymptotically there are an equal number of primes in each of the congruence classes modulo d (of course only those congruence classes $a \pmod d$ for a co-prime to d. That would mean showing that

$$\sum_{p\le x,\ \bar{p}=a} 1 \sim \frac{1}{\varphi(d)} \sum_{p \le x} 1,$$

as $x \to \infty$. The above exercise shows a statement in the same spirit, with the weight $\log n/n$ in place of the counting function 1. This exercise shows that in terms of the weight $\log n/n$, the

7. Towards non-vanishing of Dirichlet-L functions at 1

The only remaining step is to prove that $L_\chi(1) \neq 0$ for any non-trivial $\chi \in \widehat{\mathbb{Z}_d^*}$ (i.e., Lemma 9). The following lemma is a key step towards this.

Fix $d \geq 2$ and for any prime p that does not divide d, let $f(p)$ denote its order in $\mathbb{Z}_d^*$. This is the smallest integer k such that $p^k \equiv 1 \pmod{d}$. Let $g(p) = \varphi(d)/f(p)$.

Lemma 12. *Fix $d \geq 2$. Then for $s > 1$,*

$$\prod_{\chi \in \widehat{\mathbb{Z}_d^*}} L_\chi(s) = \prod_{p : p=(p,d)=1} \left(1 - \frac{1}{p^{sf(p)}}\right)^{g(p)}.$$

In particular,

$$\prod_{\chi \in \widehat{\mathbb{Z}_d^*}} L_\chi(s) \geq 1.$$

Proof. If $w, w^2, \ldots, w^{f(p)} = 1$ are the $f(p)$-th roots of unity, and if $\chi(p) = w$, then $\chi^2(p) = w^2, \ldots \chi^{f(p)}(p) = w^p$. From this, we see that the number of χ that map p to a particular $f(p)$th root of unity is the same for all these roots (and of course $\chi(p)$ has to be one of these, since $p^{f(p)} \equiv 1 \pmod{d}$). Hence p is mapped to each of them by $g(p)$ distinct character. This shows that

$$\prod_{\chi}(1 - \chi(p)z) = \left(\prod_{j=1}^{f(p)}(1 - w^j z)\right)^{g(p)} = (1 - z^{f(p)})^{g(p)}.$$

Set $z = p^{-s}$ and take product over all primes co-prime to d to get the lemma. ■

We can use the stronger conclusion with a bit of complex analysis or the weaker conclusion with a longer analysis (but avoiding any holomorphic functions) to prove Lemma 12. The former route is taken in Serre's book and the latter in the books of Stein and Shakarchi and of Apostol. First we present the real-variables approach.

Proof of Lemma 9. Recall that

$$L_{\chi_0}(s) = \prod_{p : p|d} (1 - p^{-s}) \left(\frac{1}{s-1} + \tilde{\zeta}(s)\right)$$

where $\tilde{\zeta}$ is continuous on $(0, \infty)$. For $\chi \neq \chi_0$, we know that L_χ is continuous on $(0, \infty)$. By Exercise 5 it is differentiable at 1, and hence, if $L_\chi(1) = 0$, then $L_\chi(s) = (s-1)L'_\chi(1)(1 + o(1))$ as $s \to 1$. Combining with the above fact for L_{χ_0}, we see that as $s \to 1$,

$$\prod_{\chi \in \widehat{\mathbb{Z}_d^*}} L_\chi(s) = \begin{cases} O(1) & \text{if } L_\chi(1) = 0 \text{ for some } \chi, \\ o(1) & \text{if } L_\chi(1) = 0 \text{ for at least two distinct } \chi. \end{cases}$$

By Lemma 12, the product on the left is bounded below by 1, hence the second possibility is ruled out. If χ is a character, so is $\bar{\chi}$ and $L_{\bar{\chi}} = \overline{L_{\chi}}$, hence if $L_{\chi}(1) = 0$, then $L_{\bar{\chi}}(1) = 0$. Therefore, unless $\chi = \bar{\chi}$, we cannot have $L_{\chi}(1) = 0$.

This leaves the one case when $L_{\chi}(1) = 0$ for one single real character (i.e., $\chi = \bar{\chi}$). We prove this in the next section as it is slightly longer. ■

8. Non-vanishing of the L-functions of real characters

In this case, $\chi : \mathbb{N} \mapsto \{-1, 0, 1\}$. We consider the sum

$$\sum_{(k,\ell):k\ell\leq N} \frac{\chi(k)}{\sqrt{k\ell}} = \sum_{r=1}^{N} \frac{1}{\sqrt{r}} \sum_{k:k|r} \chi(k).$$

If $r = p_1^{a_1} \dots p_m^{a_m}$, then the inner sum is

$$\begin{aligned}\sum_{0\leq b_i\leq a_i} \chi(p_1^{b_1} \dots p_m^{b_m}) &= \sum_{0\leq b_i\leq a_i} \chi(p_1)^{b_1} \dots \chi(p_m)^{a_m} \\ &= \prod_{\ell=1}^{m} (1 + \chi(p_\ell) + \dots + \chi(p_\ell)^{a_\ell}).\end{aligned}$$

Consider $1 + \chi(p_\ell) + \dots + \chi(p_\ell)^{a_\ell}$. Each term is 0 or ± 1, and always starts with a 1. Therefore, the sum is non-negative, and in fact strictly positive unless $\chi(p_\ell) = -1$ and a_ℓ is odd. In particular, if all a_ℓ are even (same as saying r is a perfect square), then the sum is at least 1, and so is the product. Thus, writing $r = t^2$ (other terms are non-negative and dropped)

$$\sum_{(k,\ell):k\ell\leq N} \frac{\chi(k)}{\sqrt{k\ell}} \geq \sum_{t\leq\sqrt{N}} \frac{1}{t} = \log\sqrt{N} + O(1). \tag{1}$$

On the other hand, we can write the sum on the left as

$$\sum_{\ell\leq N} \sum_{k\leq\sqrt{N}} \frac{\chi(k)}{\sqrt{k\ell}} + \sum_{\ell\leq\sqrt{N}} \sum_{\sqrt{N}<k\leq\frac{N}{\ell}} \frac{\chi(k)}{\sqrt{k\ell}}. \tag{2}$$

Referring back to the proof of Lemma 4, we recall that if partial sums of $(a_n)_n$ are bounded, then $|\sum_{j=m}^{n} a_j j^{-s}| \leq Cm^{-s}$. Therefore, we can bound the second summand in (2) by

$$\sum_{\ell\leq\sqrt{N}} \frac{1}{\sqrt{\ell}} \Big| \sum_{\sqrt{N}<k\leq\frac{N}{\ell}} \frac{\chi(k)}{\sqrt{k}} \Big| \leq \sum_{\ell\leq\sqrt{N}} \frac{1}{\sqrt{\ell}N^{\frac{1}{4}}} = O(1)$$

by bounding the sum of $1/\sqrt{\ell}$ by the integral of $1/\sqrt{x}$ over the appropriate range.

The first summand in (2) can be rewritten as (for some c, use Exercise 13)

$$\begin{aligned}\sum_{k\leq\sqrt{N}}\frac{\chi(k)}{\sqrt{k}}\sum_{\ell\leq\frac{N}{k}}\frac{1}{\sqrt{\ell}} &= \sum_{k\leq\sqrt{N}}\frac{\chi(k)}{\sqrt{k}}\left(2\sqrt{\frac{N}{k}}+c+O\left(\sqrt{\frac{k}{N}}\right)\right)\\ &= 2\sqrt{N}\sum_{k=1}^{\sqrt{N}}\frac{\chi(k)}{k}+c\sum_{k=1}^{N}\frac{\chi(k)}{\sqrt{k}}+O\left(\sqrt{\frac{k}{N}}\sum_{k=1}^{N}\frac{1}{\sqrt{k}}\right)\\ &= 2\sqrt{N}(L_\chi(1)-O(\frac{1}{\sqrt{N}}))+O(1)+O(1)\end{aligned}$$

where we have repeatedly used the above quoted fact from the proof of Lemma 4, and of course that $L_\chi(1)=\sum_{k=1}^{\infty}\frac{\chi(k)}{k}$. Plugging all this back into (2), we see that

$$\sum_{(k,\ell):k\ell\leq N}\frac{\chi(k)}{\sqrt{k\ell}}=2\sqrt{N}L_\chi(1)+O(1).$$

If $L_\chi(1)=1$, this contradicts (1). Hence $L_\chi(1)\neq 0$. ■

Exercise 13. If $0<s<1$, show that $\sum_{k=1}^{n}\frac{1}{k^s}=\frac{n^{1-s}}{1-s}-\frac{1+s}{2(1-s)}+O(n^{-s})$.

The proof of non-vanishing of L-functions for real characters was achieved by summing the function $\chi(k)/\sqrt{k\ell}$ on lattice points $(k,l)\in\mathbb{N}^2$ that lie under the hyperbola $xy=N$, in two different ways. This idea can be used for other functions to obtain useful number theoretical information. For example, a well-known arithmetic function is $d(n)$, the number of distinct divisors of n (e.g., $d(6)=4$). This function does not have a regular behaviour as $n\to\infty$, since $d(p)=2$ for all primes, but $d(2^n)=n+1$ can be made arbitrarily large. However, on average, it does have a regular behaviour as the following exercise shows. One can summarize it as saying that a typical large number n has about $\log n$ divisors, on average.

Exercise 14. Show that $\frac{1}{N}\sum_{k=1}^{N}d(k)\sim\log N$ (meaning that the ratio of the two sides goes to 1 as $N\to\infty$). [*Hint:* Sum an appropriate function on the lattice $\mathbb{N}^2$ under the hyperbola $xy=N$ in two different ways.]

CHAPTER 5

Fourier analysis on the circle group

1. Introduction

$T = \{e^{it} : 0 \leq t < 2\pi\}$ is an abelian group under multiplication. It is sometimes also written as $[0, 2\pi]/0 \sim 2\pi$ or as $\mathbb{R}/\mathbb{Z}$. Functions on T will be written as $f(e^{it})$ or as $f(t)$ with $t \in [0, 2\pi)$ - there should be no confusion. It also has a topology (the standard one, inherited from the complex plane) and the group operations are compatible with the topology in that $(x, y) \mapsto xy$ from $T \times T \mapsto T$ and $x \mapsto x^{-1}$ from $T \mapsto T$ are continuous. Any group with a topology w.r.t which these maps are continuous, is called a topological group.

By definition, a character is a continuous homomorphism from a topological group into T. When our group is T, we have the characters $e_m(t) := e^{2\pi imt}$, $m \in \mathbb{Z}$. We leave it as an exercise to show that there are no other characters. If we use the inner product $\langle f, g\rangle = \frac{1}{2\pi}\int_0^{2\pi} f(e^{it})\overline{g(e^{it})}dt$, then the characters form an orthonormal set as

$$\langle e_n, e_m\rangle = \frac{1}{2\pi}\int_0^{2\pi} e^{it(n-m)}dt = \delta_{n-m}.$$

The first main question is whether $\{e_n : n \in \mathbb{Z}\}$ an orthonormal basis for $L^2(T)$. The answer is yes, but unlike with finite abelian groups, dimension considerations are of no help here in showing this. What we need to show is that the span of the characters is dense in $L^2(T)$.

2. Fejér's theorem

Fix $f \in L^2(T)$. To show that it can be approximated by finite linear combinations of characters (these are called *trigonometric polynomials*), it is natural to consider its projection in $L^2(T)$ to the span of e_k, $|k| \leq n$. If we denote this projection operator by S_n, then

$$\begin{aligned} S_n f(t) &= \sum_{k=-n}^{n} \langle f, e_k\rangle e_k(t) \\ &= \frac{1}{2\pi}\int_T f(s) \sum_{k=-n}^{n} e_k(t)\overline{e_k(s)}\; dt \\ &= \frac{1}{2\pi}\int_T f(s)D_n(t-s)\; ds \;=\; (f \star D_n)(t) \end{aligned}$$

where the convolution $(f_1 \star f_2)(t) := \int_T f_1(s) f_2(t-s) \frac{ds}{2\pi}$ (whenever $f_1(s)f_2(t-s)$ is integrable) and the Dirichlet kernel

$$\begin{aligned} D_n(u) = \sum_{k=-n}^{n} e_k(u) \;&=\; e^{-inu} \frac{e^{i(2n+1)u} - 1}{e^{iu} - 1} \quad \text{(sum the geometric series)} \\ &= \frac{e^{i(n+\frac{1}{2})u} - e^{-i(n+\frac{1}{2})u}}{e^{i\frac{1}{2}u} - e^{-i\frac{1}{2}u}} \\ &= \frac{\sin[(n+\frac{1}{2})u]}{\sin[\frac{1}{2}u]}. \end{aligned}$$

$\{e_n\}_{n\in\mathbb{Z}}$ is an orthonormal basis if and only if $S_n f \to f$ in $L^2(T)$ for all $f \in L^2(T)$. How to show this? As usual in analysis, one good way is to show it for a convenient dense subset of $L^2(T)$. For example, $C(T)$ or $C^\infty(T)$. For such functions, perhaps one can try to show that $S_n f \to f$ uniformly on T, which is stronger than L^2 convergence. It is not true that $S_n f \to f$ uniformly for $f \in C(T)$. Although it is true for $f \in C^1(T)$, I do not know how to show that without first showing that $\{e_n\}$ is an orthonormal basis for $L^2(T)$.

Fejér was the one who solved the problem, by showing that $\sigma_n f = \frac{1}{n+1}(S_0 f + \ldots + S_n f)$ converges uniformly to f, for any $f \in C(T)$. Since $\sigma_n f$ is also a trigonometric polynomial, this shows that $\operatorname{span}\{e_n : n \in \mathbb{Z}\}$ is dense in $C(T)$ in sup-norm, and consequently also dense in $L^2(T)$.

Observe that $\sigma_n f = f \star K_n$, where

$$K_n(u) = \frac{1}{n+1}(D_0(u) + \ldots + D_n(u)) \;=\; \frac{1}{(n+1)\sin\frac{u}{2}} \sum_{k=0}^{n} \sin\left((k+\frac{1}{2})u\right).$$

The series can be written as

$$\operatorname{Im}\left\{\sum_{k=0}^{n} e^{i(k+\frac{1}{2})u}\right\} = \operatorname{Im}\left\{e^{i\frac{u}{2}} \frac{e^{i(n+1)u} - 1}{e^{iu} - 1}\right\} \;=\; \operatorname{Im}\left\{\frac{e^{i(n+1)u} - 1}{e^{i\frac{u}{2}} - e^{-i\frac{u}{2}}}\right\}.$$

As the denominator is $2i\sin(u/2)$, this becomes

$$\frac{1}{2\sin(u/2)} \operatorname{Re}\{1 - e^{i(n+1)u}\} \;=\; \frac{1-\cos((n+1)u)}{\sin(u/2)} \;=\; \frac{\sin^2(\frac{1}{2}(n+1)u)}{\sin^2(\frac{1}{2}u)}.$$

Plugging this back into the expression for K_n, we arrive at

$$K_n(u) = \frac{\sin^2(\frac{1}{2}(n+1)u)}{(n+1)\sin^2(\frac{1}{2}u)}$$

This is known as *Fejér's kernel*. Another useful expression for the Fejér kernel is

$$(1) \qquad K_n(u) = \frac{1}{n+1} \sum_{k=0}^{n} D_k(u) \;==\; \sum_{j=-n}^{n} \frac{n+1-|j|}{n+1} e_j(u)$$

by writing $D_k = e_{-k} + \ldots + e_k$ and interchanging the sums.

The contrast between the Dirichlet and Fejér kernels can be seen in Figure 2. The key observations

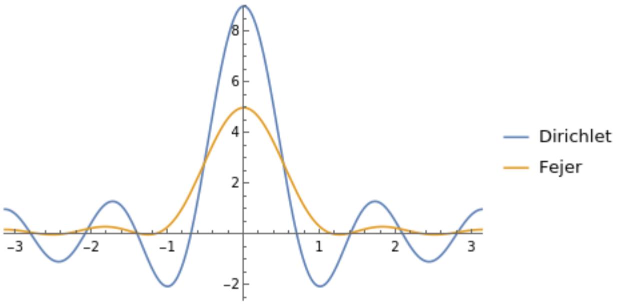

about the Fejér kernel are as follows:

(1) $K_n(u) \geq 0$ for all u, (2) $\int\limits_T K_n(u)\frac{du}{2\pi} = 1$, (3) $\int\limits_{T\setminus[-\delta,\delta]} K_n(u)\frac{du}{2\pi} \leq \frac{1}{n+1}\frac{1}{\sin^2(\delta/2)}$.

In probability language, $K_n(\cdot)$ is a probability density on T which puts most of its mass near 0 (for large n). In analysis, a sequence of functions satisfying these three conditions is called an *approximate identity*. They approximate the "Dirac delta-function", which does not exist as a function but is the putative identity for the convolution product: $f \star \delta_0 = f$ for all f. Hence we expect that $f \star K_n \to f$ and that is the gist of the proof of Fejér's theorem.

Theorem 1 (Fejér). *If $f \in C(S^1)$, then $\sigma_n f \to f$ uniformly on T. As a consequence, span$\{e_n : n \in \mathbb{Z}\}$ is dense in $L^2(T)$.*

Proof. Fix $\varepsilon > 0$ and find $\delta > 0$ so that $|f(t) - f(s)| \leq \varepsilon$ whenever $|e^{it} - e^{is}| \leq 2\delta$. As seen above $\sigma_n f(t) = \frac{1}{2\pi}\int_T f(s)K_n(t-s)ds$. Hence, with $J_\delta = \{s : |e^{it} - e^{is}| \leq \delta\}$, we have

$$\begin{aligned}|\sigma_n f(t) - f(t)| &\leq \int_{J_\delta} |f(t) - f(s)|K_n(t-s)\frac{ds}{2\pi} + \int_{T\setminus J_\delta} |f(t) - f(s)|K_n(t-s)\frac{ds}{2\pi} \\ &\leq \varepsilon \int_{J_\delta} K_n(t-s)\frac{ds}{2\pi} \; + \; 2\|f\|_{\sup}\frac{1}{n+1}\frac{1}{\sin^2(\delta/2)} \\ &\leq \varepsilon \; + \; 2\|f\|_{\sup}\frac{1}{n+1}\frac{1}{\sin^2(\delta/2)}.\end{aligned}$$

Choose $n > \frac{4\|f\|_{\sup}}{\varepsilon\sin^2(\delta/2)}$, then $\|\sigma_n f - f\|_{\sup} < 2\varepsilon$. This proves the first statement.

Given $g \in L^2(T)$, find $f \in C(T)$ such that $\|g - f\|_{L^2(T)} < \varepsilon$ and n such that $\|\sigma_n f - f\|_{L^2(T)} \leq \|\sigma_n f - f\|_{\sup} < \varepsilon$. Then $\|g - \sigma_n f\|_{L^2(T)} \leq \varepsilon$. This proves the second statement. ■

In the following exercise, derive Weierstrass' approximation and theorem from Fejér's theorem.

Exercise 2. Let $f \in C_{\mathbb{R}}[0,1]$.

(1) Construct a function $g : [-\pi, \pi] \to \mathbb{R}$ such that (a) g is even, (b) $g = f$ on $[0,1]$ and (c) g vanishes outside $[-2,2]$.

(2) Invoke Fejér's theorem to get a trigonometric polynomial T such that $\|T - g\|_{\sup} < \varepsilon$.

(3) Use the series $e^z = \sum_{k=0}^{\infty}\frac{1}{k!}z^k$ to replace the exponentials that appear in T by polynomials. Be clear about the uniform convergence issues.

(4) Conclude that there exists a polynomial P with *real* coefficients such that $\|f-P\|_{\sup} < 2\varepsilon$..

The following exercise abstracts the main features of an approximate identity.

Exercise 3. Let $g_n \in L^1(T)$ be a sequence of functions such that
(1) $\sup_n \int\limits_T |g_n(u)|\,du < \infty$, (2) $\int\limits_T g_n(u)\frac{du}{2\pi} = 1$, (3) $\int\limits_{T\setminus[-\delta,\delta]} g_n(u)\,du \to 0$ as $n\to\infty$.
Show that $f \star g_n \to f$ uniformly as $n\to\infty$ for any $f\in C(T)$.

2.1. **The problem with the Dirichlet kernel.** What is the shortcoming of the Dirichlet kernel that we cannot show that $f\star D_n \to f$ uniformly for $f\in C(T)$? One difference with the Fejér kernel is that D_n takes negative values too. But that is in itself not the main issue as shown by Exercise 3 which does not require positivity of g_n. In fact, the problem is that $\int_T |D_n|$ is not bounded.

Write $\lambda = n+\frac{1}{2}$ and observe that $\sin(\lambda u)$ vanishes in $[-\pi,\pi]$ at $u_k = \pi k/\lambda$ for $|k|\le\lambda$. The function $\sin(\lambda u)$ has constant sign on $[u_k, u_{k+1}]$ and further, in the middle-half of this interval $|\sin(\lambda u)| \ge \sin(\pi/4) > \frac{1}{2}$. Therefore,

$$\int_{u_k}^{u_{k+1}} D_n(u)\,du \;\ge\; \frac{1}{\sin(u_k/2)}\int_{u_k}^{u_{k+1}} |\sin(\lambda u)|\,du \;\ge\; \frac{\pi}{4\lambda\sin(\pi k/\lambda)}.$$

As $\sin x \le x$, the last quantity is at least $\frac{1}{4k}$. Summing over $1\le k\le n-1$ (recall that $\lambda = n+\frac{1}{2}$), we see that

$$\int_T |D_n(u)|\,du \ge \frac{1}{4}\sum_{k=1}^{n-1}\frac{1}{k} \;\ge\; \frac{1}{4}\sum_{k=1}^{n-1}\int_k^{k+1}\frac{1}{x}dx \;=\; \frac{1}{4}\log n.$$

Thus the L^1-norms of D_n are unbounded.

Exercise 4. Show that there is a function $f\in C(T)$ such that $(f\star D_n)(0)$ is unbounded. Conclude that $f\star D_n$ need not converge to f even in point-wise sense.

3. Fourier coefficients, Plancherel and inversion

For $f\in L^2(T)$, define its Fourier transform[11] as $\hat{f}:\mathbb{Z}\mapsto\mathbb{C}$, defined by

$$\hat{f}(n) = \langle f, e_n\rangle \;=\; \int_T f(t)e^{-int}\frac{dt}{2\pi}.$$

However, $L^2(T)$ is not the natural domain of functions on which to define Fourier transform. If $f\in L^1(T)$ then the integral is well-defined (since e_n is bounded) and hence $\hat{f}:\mathbb{Z}\mapsto\mathbb{C}$ is well-defined. As $L^1(T)\supseteq L^2(T)$, this extends the domain of the Fourier transform. Going further, for any measure μ on T, we can define its Fourier transform $\hat{\mu}:\mathbb{Z}\mapsto\mathbb{C}$ by $\hat{\mu}(n) = \int_T e_n d\mu$. For us measures are positive measures, but as any complex measure μ can be written uniquely as $\mu_1-\mu_2+i\mu_3-i\mu_4$ where μ_i are positive measures, this also extends the definition of the Fourier transform to complex

[11] On the circle group, it is customary to use the term "Fourier series", but we just use the common term "Fourier transform" for any group.

Borel measures on T. At that level, it becomes a further extension of the Fourier transform, since each $f \in L^1(T)$ may be identified with the complex measure $d\mu_f(t) = f(t)\frac{dt}{2\pi}$ on T, and $\widehat{\mu_f} = \hat{f}$. There are even more general objects (distributions) to which one can extend the Fourier transform, but we do not go into that here.

3.1. **The L^2 theory.** As $\{e_n\}$ form an orthonormal basis for $L^2(T)$, we can recover f from $\hat{f}$ by the "inversion formula"

$$f(t) = \sum_{n\in\mathbb{Z}} \hat{f}(n)e_n(t)$$

where the convergence of the series on the right is only in the sense of $L^2(T)$. It need not be pointwise in general. From general Hilbert space theory, we also have the Parseval-Plancherel relations:

$$\langle f, g\rangle_{L^2(T)} = \sum_{n\in\mathbb{Z}} \hat{f}(n)\overline{\hat{g}(n)} \;=\; \langle \hat{f}, \hat{g}\rangle_{L^2(\mathbb{Z})}$$

where $L^2(\mathbb{Z})$ is w.r.t the counting measure on $\mathbb{Z}$. In particular,

$$\|f\|^2_{L^2(T)} = \|\hat{f}\|^2_{L^2(\mathbb{Z})}.$$

One may worry that this does not look like the inversion formula in the finite abelian case: the forward formula is an integral and the inverse formula is a sum. The two sides of the Plancherel relationship also look different superficially. The root of all this is that one feature of the finite abelian case breaks down in general. It is no longer true that $\hat{G}$ is isomorphic to G. Instead $\hat{T} = \mathbb{Z}$ and $\hat{\mathbb{Z}} = T$, as we shall see. To make this precise, first we investigate Fourier analysis on $\mathbb{Z}$.

4. Fourier transform on $\mathbb{Z}$

If $\chi : \mathbb{Z} \mapsto T$ is a character, then $\chi(1) = e^{it}$ for some $t \in [0, 2\pi)$. As $\mathbb{Z}$ is cyclic, that fixes χ, since $\chi(n) = e_n(t)$. This are indeed characters, since $e_{n+m}(t) = e_n(t)e_m(t)$. Denote this character as ev_t or ev_{e^it} (as it is the evaluation of the characters $\{e_n\}$ of T at the point e^{it}). Thus we may identify the $\hat{\mathbb{Z}}$ with T.

A point to note is that the characters of $\mathbb{Z}$ are not in $L^2(\mathbb{Z})$, hence there is no sense in which they are orthonormal. Nevertheless, the purely formal statement of orthogonality

$$\langle \mathrm{ev}_t, \mathrm{ev}_s\rangle = \sum_{n\in\mathbb{Z}} e^{in(t-s)} \overset{?}{=} \begin{cases} 1 & \text{if } t = s, \\ 0 & \text{if } t \neq s. \end{cases}$$

is of value in developing intuition, but will have to be approached indirectly.

Exercise 5. Consider $H_n = L^2(\{-n, \dots, n\})$ with respect to normalised counting measure. Show that $\langle \mathrm{ev}_t, \mathrm{ev}_s\rangle_{H_n} \to \delta_{t,s}$ as $n \to \infty$.

This also leads to subtleties in definition of Fourier transform. We cannot simply define

$$\hat{f}(e^{it}) = \langle f, \mathrm{ev}_{e^{it}} \rangle_{L^2(\mathbb{Z})} = \sum_{n\in\mathbb{Z}} f(n)e^{-int}$$

for $f \in L^2(\mathbb{Z})$. However, the final sum does converge absolutely if we assume $f \in L^1(\mathbb{Z})$. Hence we define the Fourier transform for L^1 functions. Observe that finite complex measures on $\mathbb{Z}$ are the same as $\sum_n f(n)\delta_n$ with $f \in L^1(\mathbb{Z})$ (also $f \geq 0$ if we want positive measures), hence the Fourier transform is also well-defined for them. But what about $L^2(\mathbb{Z})$?

Let $f \in L^1(\mathbb{Z}) \cap L^2(\mathbb{Z})$. Then,

$$\begin{aligned}\int_T |\hat{f}(e^{it})|^2 \frac{dt}{2\pi} &= \int_T \sum_{m,n\in\mathbb{Z}} f(n)\overline{f(m)}e^{i(n-m)t}\ \frac{dt}{2\pi} \\ &\overset{?}{=} \sum_{m,n\in\mathbb{Z}} f(n)\overline{f(m)} \int_T e^{i(n-m)t}\ \frac{dt}{2\pi} \\ &= \sum_{n\in\mathbb{Z}} |f(n)|^2.\end{aligned}$$

The interchange of integral and sum marked with '?' is justified by Fubini's theorem, since the function $(n, m, t) \mapsto f(n)f(m)e^{i(n-m)t}$ is absolutely integrable on $\mathbb{Z}\times\mathbb{Z}\times T$, as $\int_T \sum_{m,n} |f(n)||f(m)|\ \frac{dt}{2\pi} = \|f\|^2_{L^1(\mathbb{Z})}$. Thus, $\hat{f} \in L^2(T)$ and $\|\hat{f}\|_{L^2(T)} = \|f\|_{L^2(\mathbb{Z})}$. As $L^1(\mathbb{Z}) \cap L^2(\mathbb{Z})$ is dense in $L^2(\mathbb{Z})$ (even functions on $\mathbb{Z}$ that vanish outside a finite set form a dense set in $L^2(\mathbb{Z})$). Therefore, the Fourier transform extends to an isometry from $L^2(\mathbb{Z})$ into $L^2(T)$ (surjectivity is not claimed yet, but is true and will follow shortly).

For those who have not seen this kind of argument,

Exercise 6. Let X, Y be metric spaces. Assume that Y is complete. If D is a dense subset of X and $f : D \mapsto Y$ is uniformly continuous, then show that there is a unique $g : X \mapsto Y$ that is continuous and extends f (i.e., $g(x) = f(x)$ for $x \in D$).

Use this to deduce the above statement about the extension of Fourier transform to $L^2(\mathbb{Z})$.

Since $L^p(\mathbb{Z}) \supseteq L^2(\mathbb{Z})$, this also shows that Fourier transform is well-defined on $L^p(\mathbb{Z})$. In fact, the image of $L^p(\mathbb{Z})$ under the Fourier transform is in $L^{p'}(T)$, where p' is the conjugate exponent. To see this, observe that for $f \in L^1(\mathbb{Z})$, we have $\|\hat{f}\|_{L^\infty(T)} \leq \|f\|_{L^1(\mathbb{Z})}$. Since we also have $\|\hat{f}\|_{L^2(T)} = \|f\|_{L^2(\mathbb{Z})}$, by the Riesz-Thorin interpolation theorem, the Fourier transform extends in a unique way to map $L^p(\mathbb{Z})$ into $L^{p'}(T)$ where $\frac{1}{p} + \frac{1}{p'} = 1$. An alternate way is outlined below.

Exercise 7. Fix $p \in (1, 2)$. If $f \in L^p(\mathbb{Z})$, show that there exist $g \in L^1(\mathbb{Z})$ and $h \in L^2(\mathbb{Z})$ such that $f = g+h$. Set $\hat{f} = \hat{g}+\hat{h}$. Show that this is a valid definition, and maps $L^p(\mathbb{Z})$ into $L^{p'}(T)$ and agrees with the original definition for $f \in L^1(\mathbb{Z})$.

5. T and $\mathbb{Z}$ are dual to each other

We have seen that $\hat{T} = \mathbb{Z}$ and $\hat{\mathbb{Z}} = T$, first as sets. The group structures are also consistent: If we define multiplication of characters as point-wise multiplication, then $\hat{T}$ is the same as the group $(\mathbb{Z}, +)$, since $e_n(t)e_m(t) = e_{n+m}(t)$. Similarly $\hat{\mathbb{Z}}$ is the same as the circle group, since $\mathrm{ev}_{e^{it}}(n)\mathrm{ev}_{e^{is}}(n) = \mathrm{ev}_{e^{i(t+s)}}(n)$. For the moment we ignore the question of topologies[12]. Then,

$$\hat{T} = \mathbb{Z} \qquad \text{and} \qquad \hat{\mathbb{Z}} = T.$$

To keep distinctions clear, momentarily denote the Fourier transform on G by $\mathcal{F}_G$.

Theorem 8 (Fourier inversion formula)**.** *$\mathcal{F}_{\mathbb{Z}}\mathcal{F}_T f(e^{it}) = f(e^{-it})$ for $f \in L^2(T)$ and $\mathcal{F}_T\mathcal{F}_{\mathbb{Z}} g(n) = g(-n)$ for $g \in L^2(\mathbb{Z})$. In particular, $\mathcal{F}_{\mathbb{Z}}$ is an isometry from $L^2(\mathbb{Z})$ onto $L^2(T)$ and $\mathcal{F}_T$ is an isometry from $L^2(T)$ onto $L^2(\mathbb{Z})$.*

We have already seen this, starting from T and expanding in the orthonormal basis of characters. Any unclear details are left as exercise.

Thus $\mathcal{F}_T : L^2(T) \mapsto L^2(\mathbb{Z})$ and $\mathcal{F}_{\mathbb{Z}} : L^2(\mathbb{Z}) \mapsto L^2(T)$ are almost inverses of each other. What is precisely true is that $\mathcal{F}_T^{-1} = \mathcal{F}_{\mathbb{Z}}\mathcal{F}_T\mathcal{F}_{\mathbb{Z}}$ and $\mathcal{F}_{\mathbb{Z}}^{-1} = \mathcal{F}_T\mathcal{F}_{\mathbb{Z}}\mathcal{F}_T$. The Plancherel theorem is the statement that $\mathcal{F}_{\mathbb{Z}}$ and $\mathcal{F}_T$ are unitary.

6. Fourier transforms of measures on T

The L^2-version of Fourier inversion formula does not hold for measures. Let $\mathcal{M}(T)$ denote the set of finite Borel measures on T. First we summarize the key results about Fourier transform on $\mathcal{M}(T)$.

Theorem 9. *Let $\mu, \nu \in \mathcal{M}(T)$. If $\hat{\mu} = \hat{\nu}$, then $\mu = \nu$.*

This shows that the Fourier transform is injective on $\mathcal{M}(T)$. Since positive L^1 functions are densities of measures, this also shows the injectivity on $L^1(T)$. Further, the proof will give inversion formulas to recover μ from $\hat{\mu}$. As a corollary, we shall also deduce the following.

Corollary 10. *Suppose $\mu \in \mathcal{M}(T)$ and $\hat{\mu} \in L^1(\mathbb{Z})$. Then μ is absolutely continuous and has bounded density $\check{\hat{\mu}}(-t) = \sum_n \hat{\mu}(n)e^{int}$ w.r.t. the normalized Lebesgue measure on T.*

This is the first of a general feature of Fourier transform that relates decay of the Fourier transform to the smoothness of μ. For example, if we assume that faster decay of $\hat{\mu}$, then we can deduce that the density of μ must be correspondingly smooth. The converse, interpreted qualitatively, is also true: smoothness of μ implies decay of $\hat{\mu}$. We later cover some of these aspects in exercises, but for now here is one useful lemma.

[12] The question is: why should we take discrete topology on $\mathbb{Z}$ and the standard topology on T? More precisely, if we start with a group G (say $\mathbb{Z}$), what topology does one impose on $\hat{G}$ (in this case T)? The answer is that it is the smallest topology on $\hat{G}$ that makes all the evaluations $\mathrm{ev}_x : \hat{G} \mapsto \mathbb{C}$, for $x \in G$, continuous.

Lemma 11. *If $f \in L^1(T)$, then $\hat{f}(n) \to 0$ as $|n| \to \infty$. In particular, if $\mu \in \mathcal{M}(T)$ is absolutely continuous, then $\hat{\mu}(n) \to 0$ as $|n| \to \infty$.*

One question remains unanswered above. We know that Fourier transform maps $L^2(T)$ onto $L^2(\mathbb{Z})$. What is the range of the Fourier transforms on $L^1(T)$ and $\mathcal{M}(T)$? There is no explicit answer to the first question, but there is one for the second! To state it, we make a definition:

Definition 12. A function $\varphi : \mathbb{Z} \mapsto \mathbb{C}$ is said to be positive definite if

$$\sum_{j,k=1}^{n} c_j \bar{c}_k \varphi(m_j - m_k) \geq 0$$

for any $n \geq 1$, any $m_1, \ldots, m_n \in \mathbb{Z}$ and any $c_1, \ldots, c_n \in \mathbb{C}$. Equivalently, finite principal sub-matrices of $(\varphi(j-k))_{j,k\in\mathbb{Z}}$ are positive semi-definite.

Theorem 13 (Herglotz)**.** *A function $\varphi : \mathbb{Z} \mapsto \mathbb{C}$ is equal to $\hat{\mu}$ for some $\mu \in \mathcal{M}(T)$ if and only if φ is positive definite.*

As already mentioned, the range of $L^1(T)$ under the Fourier transform has no such explicit characterization, although there are necessary and sufficient conditions one can give (for example, Lemma 11 gives a necessary condition).

Before we proceed to the proofs, we recall and generalize the important notion of convolution.

Definition 14. If $\mu, \nu \in \mathcal{M}(T)$, define $\mu \star \nu \in \mathcal{M}(T)$ by $(\mu \star \nu)(A) = \int_T \mu(A-s)d\nu(s)$ for $A \in \mathcal{B}(T)$. If $d\mu(t) = f(t)\frac{dt}{2\pi}$, then $d(\mu \star \nu)(t) = (f \star \nu)(t)\frac{dt}{2\pi}$ where $(f \star \nu)(t) := \int_T f(t-s)d\nu(s)$. If in addition, $d\nu(t) = g(t)\frac{dt}{2\pi}$, then $(f \star \nu)(t) := \int_T f(t-s)g(s)\frac{ds}{2\pi}$, agreeing with the definition of convolutions of functions that we gave earlier.

If μ, ν are probability distributions, then $\mu \star \nu$ is the probability distribution of $e^{i(x+Y)}$ where e^{iX} and e^{iY} are independent random variables having distributions μ and ν respectively.

Exercise 15. Show that $\widehat{\mu \star \nu}(n) = \hat{\mu}(n)\hat{\nu}(n)$ for $n \in \mathbb{Z}$.

Proof. Let K_n be the Fejér kernel and let $f_n(t) = (\mu \star K_n)(t) = \int_T K_n(t-s)d\mu(s)$. From the expression (1) for the Fejér kernel,

$$f_n(t) = \sum_{k=-n}^{n} \left(1 - \frac{|k|}{n+1}\right) \hat{\mu}(k)e_k(t)$$

and with $I = [\alpha, \beta] \subseteq [0, 2\pi]$

$$\begin{aligned}\int_I f_n(t)\frac{dt}{2\pi} &= \sum_{k=-n}^{n} \left(1 - \frac{|k|}{n+1}\right) \hat{\mu}(k) \int_\alpha^\beta e_k(t)\frac{dt}{2\pi} \\ &= \hat{\mu}(0)(\beta - \alpha) + \sum_{k\in[-n,n]\setminus\{0\}} \left(1 - \frac{|k|}{n}\right) \hat{\mu}(k)\frac{e^{ik\beta} - e^{ik\alpha}}{ik}.\end{aligned}$$

We now claim that $\int_I f_n(t)\frac{dt}{2\pi} \to \mu(\alpha,\beta)+\frac{1}{2}\mu\{\alpha,\beta\}$. Since the right hand side of the above equality is expressed entirely in terms of $\hat{\mu}$, this shows that from $\hat{\mu}$, we can recover $\mu(\alpha,\beta)+\frac{1}{2}\mu\{\alpha,\beta\}$ for all $0\le\alpha<\beta<2\pi$. In particular, for any arc I whose end-points are not atoms of μ, we recover $\mu(I)$. From this it is clear that we can recover μ. In particular, if $\hat{\mu}=\hat{\nu}$, then $\mu=\nu$.

To prove the claim, we write

$$\begin{aligned}\int_I f_n(t)\frac{dt}{2\pi} &= \int_I\int_T K_n(t-s)\frac{ds}{2\pi}\,\frac{dt}{2\pi}\\ &= \int_T\left[\int_I K_n(t-s)\frac{dt}{2\pi}\right]\frac{ds}{2\pi}.\end{aligned}$$

Because of the approximate identity property of K_n, we see that

$$\int_I K_n(t-s)\frac{dt}{2\pi}\to\begin{cases}1 & \text{if } s\in(\alpha,\beta),\\ 0 & \text{if } s\notin[\alpha,\beta],\\ \frac{1}{2} & \text{if } s\in\{\alpha,\beta\}.\end{cases}$$

In the last case, we use the symmetry $K_n(t)=K_n(-t)$. Further, the integral here is bounded by 1. Hence by the dominated convergence theorem,

$$\int_I f_n(t)\frac{dt}{2\pi}\to\int_T(\mathbf{1}_{(\alpha,\beta)}(s)+\frac{1}{2}\mathbf{1}_{\{\alpha,\beta\}}(s))d\mu(s) \;=\; \mu(\alpha,\beta)+\frac{1}{2}\mu\{\alpha,\beta\}.$$

This proves the claim. ■

One can extract more from the proof. One is the generalized inversion formula

$$(2)\qquad \mu(\alpha,\beta)+\frac{1}{2}\mu\{\alpha,\beta\}=\lim_{n\to\infty}\left\{\hat{\mu}(0)(\beta-\alpha)+\sum_{k\in\mathbb{Z}\setminus\{0\}}(1-\frac{|k|}{n})_+\,\hat{\mu}(k)\frac{e^{ik\beta}-e^{ik\alpha}}{ik}\right\}.$$

Proof of Corollary 10. Consider (2) and observe that $e^{ik\beta}-e^{ik\alpha}=ik(\beta-\alpha)e^{ik\gamma_k}$ for some $\gamma_k\in(\alpha_k,\beta_k)$, because of which the summand $\hat{\mu}(k)\frac{e^{ik\beta}-e^{ik\alpha}}{2\pi ik}$ is dominated by $|\hat{\mu}(k)|(\beta-\alpha)/2\pi$. By DCT,

$$\mu(\alpha,\beta)+\frac{1}{2}\mu\{\alpha,\beta\}=\hat{\mu}(0)(\beta-\alpha)+\sum_{k\in\mathbb{Z}\setminus\{0\}}\hat{\mu}(k)\frac{e^{ik\beta}-e^{ik\alpha}}{2\pi ik}.$$

The right side can be written as (again the interchange of sum and integral is justified by the summability of $\hat{\mu}$)

$$\int_\alpha^\beta\sum_{k\in\mathbb{Z}}\hat{\mu}(k)e^{ikt}\,\frac{dt}{2\pi}$$

which shows that μ has density given by the integrand. ■

Next we prove the "converse" statement, that smoothness of μ implies decay of the Fourier transform.

Proof of Lemma 11. Let $f \in L^1(T)$. First assume that $f \in C^1(T)$. Then $f' \in C(T)$ and integrate by parts to get

$$\begin{aligned}\hat{f'}(n) = \int_T f'(t)e^{-int}\frac{dt}{2\pi} &= f(t)e^{-int}|_0^{2\pi} + in\int_0^{2\pi} f(t)e^{-int}\ dt\\ &= in\hat{f}(n).\end{aligned}$$

As $f' \in C(T)$, we know that $\|\hat{f'}\|_{L^\infty(\mathbb{Z})} \le \|f'\|_{L^1(T)}$. Therefore, $\hat{f}(n) = O(1/|n|)$, which is more than saying that $\hat{f}(n) \to 0$ as $|n| \to \infty$.

Now take any $f \in L^1(T)$ and choose $g \in C^1(T)$ such that $\|f - g\|_{L^1(T)} < \varepsilon$. This is possible, for example by taking $g = f \star K_n$ for a large n. Then

$$\begin{aligned}|\hat{f}(n)| &\le |\widehat{(f-g)}(n)| + |\hat{g}(n)|\\ &\le \|f-g\|_{L^1(T)} + \frac{\|g'\|_{L^\infty(T)}}{|n|}.\end{aligned}$$

Letting $n \to \pm\infty$, we see that $\limsup_{n\to\infty} |\hat{f}(n)| \le \varepsilon$, for any $\varepsilon > 0$. ■

As for Herglotz's theorem, we only prove the easy part.

Proof of the easy half of Herglotz's theorem. For any $p \ge 1$ and any $c_1, \ldots, c_p \in \mathbb{C}$, and any $m_1, \ldots, m_p \in \mathbb{Z}$

$$\begin{aligned}\sum_{j,k=1}^{p} \bar{c}_j c_k \hat{\mu}(m_j - m_k) &= \int_T \sum_{j,k=1}^{p} \bar{c}_j c_k e^{i(m_j - m_k)\theta}\ d\mu(\theta)\\ &= \int_T |\sum_{k=1}^{p} c_k e^{-im_k\theta}|^2 d\mu(\theta)\\ &\ge 0.\end{aligned}$$

Thus, the positive semi-definiteness of $(\varphi(j-k))_{j,k\in\mathbb{Z}}$ is necessary for φ to be the Fourier transform of a measure. ■

The following exercises further amplify the statement that the smoothness of a measure or a function is equivalent to the decay of its Fourier transform.

Exercise 16. Suppose $f \in C^p(T)$. Show that $\hat{f}(n) = o(n^{-p})$ as $n \to \pm\infty$.

Exercise 17. Suppose $n^p\hat{\mu}(n)$ is (absolutely) summable, where $p \ge 0$ is an integer. Show that μ has a density $f \in C^p(T)$ and that the derivatives up to order p are bounded. Express the derivatives of f in terms of $\hat{\mu}$.

7. Locally compact abelian groups

We just outline the general theory[13].

▶ A *topological group* is a group G with a Hausdorff topology w.r.t. which the group operations $(g,h) \mapsto gh$ and $g \mapsto g^{-1}$ are continuous. Here the first map is from $G \times G$ to G, and the topology on $G \times G$ is the product topogy.

▶ If the group is abelian, and the topology is locally compact (every point has an open neighbourhood whose closure is compact), then we say that G is an LCA group.

▶ $\mathbb{Z}^d$, T^d, $\mathbb{R}^d$, finite abelian groups are all LCA groups, as are their direct products such as $\mathbb{Z} \times T \times \mathbb{R}^2$. However $\mathbb{Q}$ is not an LCA, as it is not locally compact.

▶ A character is a continuous homomorphism from G into T. The set of all characters is denoted $\hat{G}$. It is not empty, as there is at least the trivial character.

We already know that $\hat{T} = \mathbb{Z}$, $\hat{\mathbb{Z}} = T$, $\hat{\mathbb{R}} = \hat{\mathbb{R}}$. Hence $\mathbb{Z}^m \times T^p \times \mathbb{R}^q$ has dual $T^m \times \mathbb{Z}^p \times \mathbb{R}^q$.

▶ Pointwise multiplication, $\chi_1\chi_2(x) = \chi_1(x)\chi_2(x)$ makes $\hat{G}$ an abelian group.

▶ For compact $K \subseteq G$, $r > 0$, $\chi \in \hat{G}$, let $V_{K,r}(\chi) := \{\chi' \in \hat{G} : \|\chi' - \chi\|_{L^\infty(K))} < r\}$. We endow $\hat{G}$ with the smallest topology with respect to which $V_{K,r}(\chi)$ are all open.

It may be easier to understand the special case when G is σ-compact, i.e., there exist compact sets K_n that increase to G. Then, the topology above is the same as the one given by the metric on $\hat{G}$ defined by $d(\chi, \chi') = \sum_n \|\chi - \chi'\|_{K_n} 2^{-n}$. In fact this can be used to define a metric on $C_b(G)$, the space of bounded continuous functions from G to $\mathbb{C}$. In this metric, $f_n \to f$ if and only if f_n converges uniformly to f on every compact set.

Note that all our examples, $\mathbb{Z}^d, \mathbb{R}^d, T^d$ and finite products of these, are σ-compact.

▶ With the above multiplication and topology, $\hat{G}$ becomes an LCA group.

▶ For each $x \in G$, the evaluation $\text{ev}_x(\chi) = \chi(x)$ defines a character on $\hat{G}$. These are all the characters, and hence $\hat{\hat{G}} = G$. This is the *Pontryagin duality*.

▶ G is compact if and only if $\hat{G}$ is discrete (and vice versa). For example $\hat{T} = \mathbb{Z}$ and $\hat{\mathbb{Z}} = T$.

▶ To go further and define Fourier transform, we need a measure to integrate against. To respect the group structure, what we need is a measure μ on the Borel sigma-algebra of G that is regular ($\mu(K) < \infty$ for compact K; $\mu(A) = \sup\{\mu(K) : A \supseteq K \text{ compact }\}$; $\mu(A) = \inf\{\mu(G) : A \subseteq G \text{ open}\}$) and *invariant* ($\mu(A + x) = \mu(A)$ where $A + x = \{a + x : a \in A\}$). On any LCA group G, such a measure exists and is unique up to multiplication by positive constants. It is called the *Haar measure* and we denote it as m_G (an arbitrary choice of the scalar multiple is made). Everywhere below $L^p(G)$ will mean $L^p(G, m_G)$.

▶ For $f \in L^1(G, \mu)$, its Fourier transform is $\hat{f} : \hat{G} \mapsto \mathbb{C}$ defined by $\hat{f}(\chi) = \int_G f(x)\overline{\chi}(x) d\mu_G(x)$. For $\mu \in \mathcal{M}(G)$, the space of finite Borel measures on G, define $\hat{\mu}(\chi) = \int_G \chi d\mu$.

[13] The first chapter of Rudin's *Fourier analysis on groups* is an excellent self-contained introduction with all the proofs.

- ▶ For $f \in L^1(G) \cap L^2(G)$ (if G is compact then $L^1 \cap L^2 = L^2$), one can show that $\hat{f} \in L^2(\hat{G})$ and $\|\hat{f}\|_{L^2(\hat{G})} = \kappa_G \|f\|_{L^2(G)}$. Here κ_G is a constant, which is not necessarily 1 because we arbitrarily fixed the Haar measures $m_G, m_{\hat{G}}$. One can of course change the Haar measure on $\hat{G}$ to $\kappa_G m_{\hat{G}}$, in which case the constant changes to 1.

Hence, the Fourier transform can be extended to an isometry of $L^2(G)$ into $L^2(\hat{G})$ (*Plancherel relation*). This isomorphism is also surjective, as seen next.

- ▶ For $\mu \in \mathcal{M}(G)$ and $\nu \in \mathcal{M}(\hat{G})$, we have the *Parseval relation*: $\int_{\hat{G}} \hat{\mu}(\chi) d\nu(\chi) = \int_G \hat{\nu}(x) d\mu(x)$. To see this integrate $(x, \chi) \mapsto \chi(x)$ w.r.t. $\mu \otimes \nu$ in two ways.

- ▶ Injectivity of the Fourier transform on $\mathcal{M}(G)$ and on $L^1(G)$ are true. Further, when $\hat{\mu}$.

CHAPTER 6

Fourier analysis on $\mathbb{R}$

1. Self-duality

The characters of $(\mathbb{R},+)$ are precisely $\{e_t : t \in \mathbb{R}\}$, where $e_t(x) = e^{itx}$. Since $e_t(x)e_s(x) = e_{t+s}(x)$, this shows that the dual $\hat{\mathbb{R}}$ is also $(\mathbb{R},+)$. This makes the theory a bit more symmetrical compared to that on the circle group, but when one keeps in mind the more general situation of locally compact abelian groups, it is better to have in mind two separate copies of $\mathbb{R}$, one for the original, one for the dual.

In many ways Fourier analysis on $\mathbb{R}$ will look similar to that of T, with various sums replaced by integrals, but in other ways the similarities with $\mathbb{Z}$ is even closer. Both $\mathbb{R}$ and $\mathbb{Z}$ are non-compact groups, and their characters are not L^2 functions, in particular there is no orthogonality. But as in the case of $\mathbb{Z}$, the approximate orthonormality

$$\frac{1}{2L}\int_{-L}^{L} e_t(x)\overline{e_s(x)}dx \to \delta_{t,s} \qquad \text{as } L \to \infty,$$

provides valuable intuition and also route to various proofs.

2. Fourier transform

For $f \in L^1(\mathbb{R})$, define its Fourier transform $\hat{f} : \mathbb{R} \mapsto \mathbb{C}$ by $\hat{f}(t) = \int_{\mathbb{R}} f(x)\overline{e_t(x)}dx$. For $\mu \in \mathcal{M}(\mathbb{R})$ (finite Borel measures on $\mathbb{R}$), similarly define $\hat{\mu}(t) = \int_{\mathbb{R}} \overline{e_t(x)}d\mu(x)$. The two definitions are consistent in that if μ has density f then $\hat{\mu} = \hat{f}$.

Example 1. If $\mu = \delta_0$, then $\hat{\mu}(t) = 1$. If μ is uniform on $[-1,1]$, then $\hat{\mu}(t) = \frac{1}{2}\int_{-1}^{1} e^{-itx}dx = \frac{\sin t}{t}$. If $d\mu(x) = \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}x^2}dx$ is the Gaussian measure, then $\hat{\mu}(t) = e^{-\frac{1}{2}t^2}$.

Like in $\mathbb{Z}$, here too L^2 is not contained in L^1 (nor is there a containment in the reverse direction). The way to define Fourier transform for L^2 functions is similar to the way we did in $\mathbb{Z}$. We elaborate on this after we see basic properties of the Fourier transform on L^1.

3. Properties of the Fourier transform

Let us list various properties of the Fourier transform[14].

[14] We have given many references already. In addition, volume 2 of Feller's *An introduction to probability theory and its applications* is highly recommended for Fourier transforms of measures.

▶ For $\mu \in \mathcal{M}(\mathbb{R})$, $\hat{\mu}$ is a bounded, uniformly continuous function. Indeed, $|\hat{\mu}(t) - \hat{\mu}(s)| \le \int_{\mathbb{R}} |e^{i(t-s)x} - 1| d\mu(x)$. As $t - s \to 0$, apply DCT (the integrand is bounded by 2) to get uniform continuity. It is also obvious that $|\hat{\mu}(t)| \le \mu(\mathbb{R})$.

▶ For $f \in L^1(\mathbb{R})$, $\hat{f}$ is a bounded continuous function that vanishes at infinity. The boundedness and continuity can be argued as above. To see vanishing at infinity, first assume that $f \in C_c^1$. Integrating by parts, we see that $\hat{f'}(t) = it\hat{f}(t)$. Since $\hat{f'}$ is bounded, it must be the case that $\hat{f}(t) = O(1/|t|)$. For general $f \in L^1$, find $g \in C_c^1$ such that $\|f - g\|_1 < \varepsilon$. Then $|\hat{f}(t) - \hat{g}(t)| < \varepsilon$ for all t. Let $t \to \pm\infty$ to see that $\limsup |\hat{f}(t)| \le \varepsilon$ as $t \to \pm\infty$.

As a particular case, if $\mu \in \mathcal{M}(\mathbb{R})$ has a density, then $\hat{\mu}(t) \to 0$ as $t \to \pm\infty$.

▶ Inversion formula: If $\mu \in \mathcal{M}(\mathbb{R})$, then

$$\mu(a,b) + \frac{1}{2}\mu\{a,b\} = \lim_{L\to\infty} \frac{1}{2\pi} \int_{-L}^{L} \hat{\mu}(t) \frac{e^{ibt} - e^{iat}}{it} \left(1 - \frac{|t|}{L}\right)_+ dt.$$

The proof is similar to the one we gave on the circle group, and can be found in many books. We omit it.

▶ In particular, if $\hat{\mu} = \hat{\nu}$, then $\mu = \nu$. In addition, if $\hat{\mu} \in L^1$, then we can apply DCT above to get

$$\mu(a,b) + \frac{1}{2}\mu\{a,b\} = \frac{1}{2\pi} \int_{-L}^{L} \hat{\mu}(t) \frac{e^{ibt} - e^{iat}}{it} dt$$

from which it follows that μ must have density given by $\frac{1}{2\pi}\hat{\hat{\mu}}(-x)$.

▶ For $f \in L^1$, if $\hat{f} \in L^1$ then $f(x) = \frac{1}{2\pi}\int_{\mathbb{R}} \hat{f}(t) e^{itx} dx$ for a.e. x. In particular, f can be modified to be a continuous function vanishing at infinity. This inversion formula can also be written as $\hat{\hat{f}}(x) = 2\pi f(-x)$.

▶ One component in the omitted proofs and of great importance in general, is convolution. For $\mu, \nu \in \mathcal{M}(\mathbb{R})$, the convolution $\mu \star \nu(A) := \int_A \mu(A - x) d\nu(x)$ for $A \in \mathcal{B}_{\mathbb{R}}$, defines another element of $\mathcal{M}(\mathbb{R})$. For $f, g \in L^1$, $(f \star g)(x) = \int f(x - t) g(t) dt$ is defined for a.e. t and $f \star g \in L^1$. The definitions are consistent in the sense that of $d\mu(x) = f(x)dx$ and $d\nu(x) = g(x)dx$ then $d(\mu \star \nu)(x) = (f \star g)(x)dx$. In fact, $\mu \star \nu$ has a density if just one of μ or ν does (why?).

One important point: $\widehat{(\mu \star \nu)}(t) = \hat{\mu}(t)\hat{\nu}(t)$ and $\widehat{(f \star g)}(t) = \hat{f}(t)\hat{g}(t)$. This is at the heart of why Fourier transforms are useful in probability theory, when studying sums of independent random variables. If μ, ν are probability measures, then $\mu \star \nu$ is the distribution of a sum of independent random variables drawn from these two distributions.

▶ Suppose $f \in L^1 \cap C^1$ and $f' \in L^1$. Then we can integrate by parts to see that $\hat{f'}(t) = it\hat{f}(t)$. Since $\hat{f'}$ vanishes at infinity, it follows that $\hat{f}(t) = o(1/|t|)$. Continuing, show that if $f \in L^1 \cap C^k$, and that $f^{(j)} \in L^1$ for all $j \le k$, then $\widehat{f^{(k)}}(t) = (it)^k \hat{f}(t)$ and in particular, $\hat{f}(t) = o(|t|^{-k})$. All this can be summarized by the slogan "Smoothness of a function implies the decay of the Fourier transform".

▶ Next we state the slogan "Decay of the function implies the smoothness of the Fourier transform". Indeed, start from $\hat{f}(t) = \int f(x) e^{-itx} dx$ and formally differentiate under the integral to get

$\hat{f}^{(k)}(t) = (-i)^k \int x^k f(x)dx$. Can this be justified. It suffices to consider the case $k = 1$ and proceed inductively. Start with

$$\frac{\hat{f}(t+h) - \hat{f}(t)}{h} = \int f(x)e^{-itx}\frac{e^{-ihx} - 1}{h}dx.$$

When $h \downarrow 0$, the integrand on the right converges to $-ixf(x)e^{-itx}$. Further, as $e^{-ihx} - 1 = -ihe^{-ihy}$ for some y between 0 and x, and hence the integrand is bounded by $|xf(x)|$. Therefore, if we assume that $xf(x) \in L^1$, then the formal calculationis justified and we get $(\hat{f})'(t) = -i(\widehat{xf(x)})(t)$. For the kth derivative formula, it suffices to assume that $x^k f(x) \in L^1$.

▶ Because of the inversion formulas, we get for free two additional slogans: "Smoothness of the Fourier transform implies the decay of the function" and "Decay of the Fourier transform implies the smoothness of the function". We leave as exercise to write down the precise statements. It may also be observed that the statements are not exact converses: Assuming $f \in C^k$ gives $\hat{f}(t) = o(t^{-k})$ but to get $f \in C^k$ we need to assume $t^k\hat{f}(t) \in L^1$. This was also seen above: If μ has density, then $\hat{\mu}$ decays at infinity. To prove that μ has density we had to assume that $\hat{\mu} \in L^1$.

▶ A function $\varphi : \mathbb{R} \mapsto \mathbb{C}$ is said to be positive definite if $\varphi(-t) = \overline{\varphi(t)}$ and $\sum_{j,k=1}^{n} c_j\overline{c_k}\varphi(t_j - t_k) \geq 0$ for any $n \geq 1$ and $c_1, \ldots, c_n \in \mathbb{C}$ and $t_1, \ldots, t_n \in \mathbb{R}$. The relevance of this definition is as follows:

Bochner's theorem: Let $\varphi : \mathbb{R} \mapsto \mathbb{C}$. Then $\varphi = \hat{\mu}$ for some $\mu \in \mathcal{M}(\mathbb{R})$ if and only if φ is continuous and positive definite.

▶ The proof of the easy side of Bochner's theorem is similar to that in the circle group. If $\mu \in \mathcal{M}(\mathbb{R})$, then we have seen that $\hat{\mu}$ is continuous. Further,

$$\sum_{j,k=1}^{n} c_j\overline{c_k}\hat{\mu}(t_j - t_k) = \int_{\mathbb{R}} \Big|\sum_{j=1}^{n} c_j e^{-ijx}\Big|^2 d\mu(x) \; \geq 0.$$

▶ There is no analogous theorem characterizing the range of the Fourier transform on L^1. We only have necessary and sufficient conditions (in terms of smoothness, as we have seen).

Exercise 2. Show that the Fourier transform of $\frac{1}{\pi(1+x^2)}$ is $e^{-|t|}$.

4. Fourier transform on $L^2(\mathbb{R})$

First approach.

(1) Show the Plancherel relation $\|\hat{f}\|_2^2 = 2\pi\|f\|_2^2$ for $f \in L^1 \cap L^2$. This is indicated in the next section.

(2) Using the density of $L^1 \cap L^2$ in L^2, extend the Fourier transform to L^2. The Plancherel relation continues to hold.

From the fact that $\|\hat{f}\|_\infty \leq \|f\|_1$ for $f \in L^1$ and $\|\hat{f}\|_2 = \sqrt{2\pi}\|f\|_2$, using Riesz-Thorin interpolation we can extend the Fourier transform to L^p for $1 < p < 2$ and see that it maps into L^q, where q is the conjugate of p and satisfies $\|\hat{f}\|_q \leq (2\pi)^{\theta/2}\|f\|_p$ if $\frac{1}{p} = \frac{1-\theta}{1} + \frac{\theta}{2}$.

Second approach. We describe another way to define Fourier transform for L^2 functions[15]. We start with the Gaussian $\varphi(x) = \frac{1}{\sqrt{2\pi}}e^{-x^2/2}$ whose Fourier transform is $\hat{\varphi}(t) = \sqrt{2\pi}\varphi(t)$. We shall see in the next section that

$$\widehat{f'}(t) = it\hat{f}(t) \qquad \text{and} \qquad (\hat{f})'(t) = -i(\widehat{xf(x)})(t)$$

and therefore, defining $\mathcal{A} = -\frac{d}{dx} + x$, we have $\widehat{\mathcal{A}f} = -i\mathcal{A}\hat{f}$. In particular, if we define $h_k = \mathcal{A}^k\varphi$, then we get $\hat{h_k} = (-i)^k\sqrt{2\pi}h_k$. Thus, formally we may say that h_k are eigenfunctions of the Fourier transform with eigenvalue $(-i)^k\sqrt{2\pi}$. By direct computation, we see that

$$h_1(x) = 2x\varphi(x), \qquad h_2(x) = (4x^2-2)\varphi(x), \qquad h_3(x) = (8x^3-12x)\varphi(x), \ldots$$

In general, it is easy to see that $h_k = H_k\varphi$, where H_k is a polynomial of degree equal to k. It has real coefficients, leading coefficient equal to 2^n, and in fact all its coefficients can be computed explicitly. They are known as Hermite polynomials, and h_k are called Hermite functions. From this structure, we see that span$\{h_k : k \geq 0\} = \{p(x)\varphi(x) : p \text{ is a polynomial}\}$. The latter is dense in $L^2(\mathbb{R})$ (why?). We now claim that h_k are orthogonal in $L^2(\mathbb{R})$. To see this, let $k > \ell \geq 0$ and integrate by parts to get

$$\int h_k(x)h_\ell(x)dx = \int \varphi(x)\mathcal{A}^{*k}h_\ell \, dx$$

where $\mathcal{A}^* = \frac{d}{dx} + x$.

Exercise 3. Show that $\mathcal{A}^*h_k = c_kh_{k-1}$ (where $h_{-1} = 0$) for an explicit constant c_k

Observe that $\mathcal{A}$ raises the index of Hermite functions by 1 and $\mathcal{A}^*$ decreases it by 1. From the exercise, we see that $\mathcal{A}^{*k}h_\ell = 0$ if $\ell < k$. Further, if $k = \ell$, then $\mathcal{A}^kh_k = c_kc_{k-1}\ldots c_1$ (constant function) and hence $\|h_k\|_2^2 = c_kc_{k-1}\ldots c_1$. In short, $\{h_k : k \geq 0\}$ is an orthogonal basis for $L^2(\mathbb{R})$.

Therefore, for $f \in L^2(\mathbb{R})$, we have the L^2 expansion $f = \sum_{k\geq 0}\frac{1}{\|h_k\|_2^2}\langle f, h_k\rangle h_k$. This makes it natural to define the Fourier transform as

$$\hat{f} := \sqrt{2\pi}\sum_{k\geq 0}\frac{1}{\|h_k\|_2^2}\langle f, h_k\rangle(-i)^kh_k.$$

Immediately we get the Plancherel relation

$$\|\hat{f}\|_2^2 = 2\pi\|f\|_2^2.$$

Of course, things are not satisfactory till one proves that for $f \in L^1 \cap L^2$ this definition of Fourier transform agrees with the original one. To see this, observe that if $f \in \text{span}\{h_0, h_1, \ldots\} = \{p(x)\varphi(x) : p \text{ is a polynomial}\}$ the two definitions

(1) $\hat{f}(t) = \int f(x)e^{-itx}dx$ and

(2) $\hat{f} = \sqrt{2\pi}\sum_{k\geq 0}\frac{1}{\|h_k\|_2^2}\langle f, h_k\rangle(-i)^kh_k$.

[15] Taken from chapter 1 of Thangavelu's book *An introduction to the uncertainty principle*.

obviously agree. The first definition is a uniformly continuous map from L^1 to L^∞ and the second is uniformly continuous from L^2 to L^2.

Exercise 4. Show that

(1) $H_n(x) = e^{\frac{1}{2}x^2} \frac{d^n}{dx^n} e^{-\frac{1}{2}x^2}$.

(2) Find the constants c_k above and show that $\|h_n\|_2^2 = n!$.

(3) Show that $\sum_{k=0}^{\infty} h_k(x) \frac{t^k}{k!} = \frac{1}{\sqrt{2\pi}} e^{xt - \frac{1}{2}t^2}$.

5. Poisson summation formula

Let $f \in L^1(\mathbb{R})$ and define $g : T \mapsto \mathbb{C}$ (in this section we shall write T as $[0, 2\pi)$) by $g(x) = \sum_{n \in \mathbb{Z}} f(x - 2\pi n)$. Observe that

$$\sum_{n \in \mathbb{Z}} \int_T |f(x - 2\pi n)| dx = \sum_{n \in \mathbb{Z}} \int_{2\pi n}^{2\pi(n+1)} |f(x)| dx = \int_{\mathbb{R}} |f(x)| dx = \|f\|_{L^1(\mathbb{R})}.$$

Therefore, the series defining g is absolutely convergent for a.e. x and $g \in L^1(T)$. Hence the above integration can be done without absolute values and shows that $\int_T g = \int_{\mathbb{R}} f$. Actually more is true.

$$\begin{aligned} \int_T g(x) e^{-ikx} dx &= \int_T \sum_{n \in \mathbb{Z}} f(x - 2\pi n) e^{-ikx} \, dx \\ &= \sum_{n \in \mathbb{Z}} \int_0^{2\pi} f(x - 2\pi n) e^{-ikx} dx \end{aligned}$$

where the application of Fubini's theorem is justified by the earlier proof that $g \in L^1$. Now change variables $y = x - 2\pi n$ in the inner integral and observe that $e^{2\pi i k n} = 1$ to see that it is $\int_{2\pi n}^{2\pi(n+1)} f(x) e^{-kx} dx$. Summing up, we arrive at

$$2\pi \hat{g}(k) = \hat{f}(k) \quad \text{for } k \in \mathbb{Z}.$$

As customary, we have used the hat to denote Fourier transform, but on the right side it is Fourier transform on $\mathbb{R}$ and on the left side it is Fourier transform on the circle group (where we define it with a factor of $1/2\pi$ in the integral).

To proceed further, assume that $f \in C^1(\mathbb{R})$ and that $f(x)$ and $f'(x)$ are both bounded by C/x^2. Then the series $\sum_n f'(x - 2\pi n)$ and $\sum_n f(x - 2\pi n)$ both converge uniformly on $[0, 2\pi)$. By a standard lemma one learns in basic analysis, this shows that $g \in C^1(T)$ and $g'(x) = \sum_n f'(x - 2\pi n)$. For $C^1(T)$ functions, the Fourier series converges uniformly to the function. Hence, we have

$$g(x) = \sum_{n \in \mathbb{Z}} \hat{g}(k) e^{ikx}.$$

Now apply the definition of g and the relationship between $\hat{g}$ and $\hat{f}$ to get

$$2\pi \sum_{n \in \mathbb{Z}} f(x - 2\pi n) = \sum_{n \in \mathbb{Z}} \hat{f}(k) e^{ikx} \quad \text{for all } x \in [0, 2\pi).$$

In particular, setting $x = 0$ we get

$$2\pi \sum_{n\in\mathbb{Z}} f(x - 2\pi n) = \sum_{k\in\mathbb{Z}} \hat{f}(k).$$

This is known as the *Poisson summation formula*

CHAPTER 7

Bernoulli convolution problem

Let $0 < \lambda < 1$ and define the Bernoulli convolution

$$\nu_\lambda = (\frac{1}{2}\delta_{-\lambda} + \frac{1}{2}\delta_\lambda) \star (\frac{1}{2}\delta_{-\lambda^2} + \frac{1}{2}\delta_{\lambda^2}) \star (\frac{1}{2}\delta_{-\lambda^3} + \frac{1}{2}\delta_{\lambda^3}) \dots$$

An equivalent description of ν_λ is that it is the distribution of the random variable $X_\lambda = \sum_{n=1}^{\infty} \varepsilon_n \lambda^n$, where ε_n are independent random variables taking the values ± 1 with equal probability. Yet another way to characterise it is via the Fourier transform:

$$\hat{\nu}_\lambda(t) = \mathbf{E}[e^{itX_\lambda}] = \prod_{n=1}^{\infty} \cos(\lambda^n t).$$

The product on the right converges uniformly over t in compact sets, as $1-\cos(\lambda^n t) = 2\sin^2(\lambda^n t/2) \leq \lambda^{2n} t^2/2$ is summable, uniformly over t in compact sets. This also shows that $\hat{\nu}_\lambda(t) \neq 0$ unless $t = \pi(m + \frac{1}{2})\lambda^{-n}$ for some $m \in \mathbb{Z}$ and $n \geq 1$ (in general, if $\sum_n |a_n| < \infty$ and $a_n \neq 1$ for all n, then $\prod_n (1 - a_n) \neq 0$).

Example 1. If $\lambda = \frac{1}{2}$, then ν_λ is the normalized Lebesgue measure on $[-1, 1]$. If $\lambda = \frac{1}{3}$, then ν_λ is the Cantor measure, supported on the standard $\frac{1}{3}$-Cantor set (except that we do the middle-third deletion starting from $[-1, 1]$ instead of $[0, 1]$).

Like the Lebesgue measure and Cantor measure, the measures ν_λ has an important self-similarity property.

Self-similarity: If X_λ has distribution ν_λ and ε is an independent symmetric Bernoulli random variable, then $\varepsilon + \lambda X_\lambda$ has the same distribution as X_λ. This is clear from the series expansion of X_λ.

We now claim that ν_λ is the only probability measure for which this distributional equality holds. That is, if $X \sim \nu$ and symmetric Bernoulli ε are independent, and $\varepsilon + \lambda X$ also has distribution ν, then $\nu = \nu_\lambda$. To see this first take expectation over ε to get

$$\begin{aligned}\mathbf{E}\left[e^{it(\varepsilon+\lambda X)}\right] &= \frac{1}{2}\mathbf{E}[e^{it(1+\lambda X)}] + \frac{1}{2}\mathbf{E}[e^{it(-1+\lambda X)}] \\ &= (\cos t)\mathbf{E}[e^{it\lambda X}]\end{aligned}$$

which means that that $\hat{\nu}(t) = (\cos t)\hat{\nu}(t\lambda)$. Continuing, we see that $\hat{\nu}(t) = \hat{\nu}(t\lambda^N)\prod_{n=1}^{N-1}\cos(t\lambda^n)$. As $N \to \infty$, the product converges to $\hat{\nu}_\lambda(t)$, while $\hat{\nu}(t\lambda^N) \to \hat{\nu}(0) = 1$, showing that $\hat{\nu} = \hat{\nu}_\lambda$. Therefore $\nu = \nu_\lambda$.

The main question of Bernoulli convolution is whether ν_λ is absolutely continuous to Lebesgue measure on $[-1,1]$ or whether it is singular. A priori, it could be neither, having a non-zero absolutely continuous part and a non-zero singular part, but that does not happen.

Lemma 2 (Jessen's law of pure types)**.** *For any $\lambda \in (0,1)$, the measure ν_λ is either absolutely continuous or singular.*

Proof. To see this, we observe that both the singular and absolutely continuous parts of ν_λ must satisfy the same self-similarity as ν_λ (why?). Therefore, if one of them is non-zero, then it must be a multiple of ν_λ. This shows that exactly one of them can be non-zero. ■

We have already seen that $\nu_{\frac{1}{2}}$ is the normalized Lebesgue measure on $[-1,1]$. We also said that $\nu_{\frac{1}{3}}$ is the Cantor measure and hence singular. In fact, the same holds for any $\lambda < \frac{1}{2}$.

Claim 3. *ν_λ is singular for $\lambda < \frac{1}{2}$.*

Proof. To see this, observe that the series beyond the nth term is

$$|\sum_{k=n}^{\infty} \varepsilon_k \lambda^k| \le \sum_{k\ge n} \lambda^k = \frac{\lambda^n}{1-\lambda}.$$

Further, there are only 2^{n-1} different possible values of $\varepsilon_1\lambda + \ldots + \varepsilon_{n-1}\lambda^{n-1}$. Hence, the support of ν_λ is covered by 2^{n-1} intervals each of length at most $2\frac{\lambda^n}{1-\lambda}$. The total Lebesgue measure of these intervals is $(2\lambda)^n/(1-\lambda)$ which can be made arbitrarily small by choosing n large. Hence ν_λ must be singular. ■

Exercise 4. Show that if $\frac{1}{2} \le \lambda < 1$, then the support of ν_λ is the interval $[-(1-\lambda)^{-1}, (1-\lambda)^{-1}]$.

Recall that the support of a measure is the smallest *closed set* whose complement has zero measure. Hence, the above exercise does not imply by any means that ν_λ is absolutely continuous for $\lambda > \frac{1}{2}$. For example, if rational numbers are enumerated as $r_1, r_2, \ldots$ and $\nu = \sum_j 2^{-j}\delta_{r_j}$, then ν has support equal to $\mathbb{R}$. Of course, ν_λ is not this bad - it has no atoms (why?), but it could be a singular continuous measure. However, the above exercise, the case $\lambda = \frac{1}{2}$ and a wish to see a natural progression in λ may lead one to guess that ν_λ ought to be absolutely continuous for $\lambda \ge \frac{1}{2}$.

We shall use the Riemann-Lebesgue lemma to see that there are $\lambda \in (\frac{1}{2}, 1)$ for which ν_λ is singular! To state this amazing discovery of Paul Erdös, we recall some notions.

We say that $\theta \in \mathbb{C}$ is called an *algebraic integer* if it is the root of a monic polynomial with integer coefficients. In that case, there is a unique such polynomial of minimal degree, called the *minimal polynomial* of θ. The minimal polynomial is irreducible, and its other roots are called the (Galois) conjugates of θ. A *Pisot-Vijayaraghavan number* or PV number is a real algebraic integer greater than 1, all of whose conjugates are of absolute value less than 1.

Example 5. The minimal polynomial of $(1+\sqrt{5})/2$ is $x^2 - x - 1$ and the other root of the minimal polynomial is $(1-\sqrt{5})/2 = -0.618\ldots$ and hence $(1+\sqrt{5})/2$ is a PV number.

Theorem 6 (Erdös (1939))**.** *Suppose $\lambda = \frac{1}{\theta}$ where $1 < \theta < 2$ is a PV number. Then ν_λ is singular.*

Proof. We claim that $\hat{\nu}_\lambda(2\pi\theta^k) \not\to 0$. Recall that $\lambda = 1/\theta$ to write

$$\begin{aligned}\hat{\nu}_\lambda(2\pi\theta^k) &= \prod_{j=1}^{k-1} \cos(2\pi\theta^j) \prod_{j=k+1}^{\infty} \cos(2\pi\lambda^{j-k}) \\ &= \hat{\nu}_\lambda(2\pi) \prod_{j=1}^{k-1} \cos(2\pi\theta^j).\end{aligned}$$

Observe that $1 - \cos(2\pi x) = 2\sin^2(\pi x) \leq 2\pi^2 x^2$. By the evenness and periodicity of cosine, we can write this as $1 - \cos(2\pi x) < 20\llbracket x \rrbracket^2$, where $\llbracket x \rrbracket$ is the distance from x to the closest integer. If the conjugates of θ are $\tau_1, \ldots, \tau_m$, then $\theta^j + \tau_1^j + \ldots + \tau_m^j \in \mathbb{Z}$, as it can be written as sums of products of coefficients of the minimal polynomial[16]. If $u = \max_i |\tau_i| < 1$, then this shows that $\llbracket \theta^j \rrbracket \leq m u^j$. Therefore, $1 - \cos(2\pi\theta^j)$ is summable, and hence $\prod_{j\geq 1} \cos(2\pi\theta^j)$ converges. As θ is irrational, $\cos(2\pi\theta^j) \neq 0$ for all j, showing that $\hat{\nu}_\lambda(\theta^k)$ converges to a non-zero constant as $k \to \infty$ (recall that $\hat{\nu}_\lambda(2\pi) \neq 0$). Thus, $\hat{\nu}_\lambda$ does not vanish at infinity and hence by the Riemann-Lebesgue lemma, ν_λ is not absolutely continuous. By the law of pure types, it must be singular. ■

One may now swing to the other direction and wonder if ν_λ is singular for all $\lambda > \frac{1}{2}$. It is not, by the following result of Wintner.

Theorem 7 (Wintner (1935))**.** *If $\lambda = 2^{-\frac{1}{k}}$, then ν_λ is absolutely continuous and has a C^{k-2} density.*

Proof. Fix a k and write integers modulo k to see that

$$\hat{\nu}_\lambda(t) = \prod_{r=0}^{k-1} \prod_{m=1}^{\infty} \cos(2^{-m} t 2^{-\frac{r}{k}}) \;=\; \prod_{r=0}^{k-1} \hat{\nu}_{\frac{1}{2}}(t 2^{-\frac{r}{k}}).$$

But $\nu_{\frac{1}{2}}$ is the normalized Lebesgue measure on $[-1,1]$ and $t \mapsto \hat{\nu}(t2^{-r/k})$ is the Fourier transform of the normalized Lebesgue measure on $[-2^{-r/k}, 2^{-r/k}]$. From the above formula, ν_λ is a convolution of k of these measures, and therefore has density that is C^{k-2}. ■

If the last line of the proof is not clear, take as an exercise to prove that if μ has a C^k density and ν has a C^ℓ density, then $\mu \star \nu$ has $C^{k+\ell}$ density. In Wintner's paper he observes that $\hat{\nu}(t) = \sin t/t$ to write

$$|\hat{\nu}_\lambda(t)| = \prod_{r=0}^{k-1} \left| \frac{\sin(t2^{-\frac{r}{k}})}{t2^{-\frac{r}{k}}} \right| \;=\; O(|t|^{-k})$$

[16] If $P(x) = x^{m+1} + a_1 x^m + \ldots + a_m$ is the minimal polynomial, then $p_j := \theta^j + \tau_1^j + \ldots + \tau_m^j$ is equal to $-a_1$ for $j = 1$, equal to $a_1^2 - 2a_2$ for $j = 2$, and so on. More precisely, inductively one can show that p_j is an integer, based on *Newton's identities:* $\sum_{i=1}^{k} a_{k-i} p_i = -k a_k$ where $a_0 = 1$.

and asserts that this implies that ν_λ has a C^{k-1} density. I don't see how (consider $k = 1$). Perhaps he means piecewise C^{k-1}, but to be on the safe side I have proved a slightly weaker statement.

Now we have a countable set of $\lambda \in (\frac{1}{2}, 1)$ (reciprocals of PV numbers) for which ν_λ is singular and a countable set of λ (reciprocals of roots of 2) for which it is absolutely continuous (in fact with a certain amount of smoothness). What about all the other λ? The problem is still open, but results like the following are known.

Fact 8 (Solomyak)**.** ν_λ is absolutely continuous for a.e. $\lambda \in (\frac{1}{2}, 1)$ and the density is in L^2.

This was conjectured by Garsia, after a weaker result of Erdös that stated that ν_λ is absolutely continuous for a.e. $\lambda \in (1 - \delta, 1)$ for some $\delta > 0$.

Exercise 9. Use Solomyak's result and deduce that the density of ν_λ is C^k for a.e. $\lambda \in (1 - \delta_k, 1)$ for some $\delta_k > 0$. The result was stated this way in Erdös' paper.

In the proof of Erdös' theorem, we showed that $[\![\theta^j]\!]$ decays exponentially, but what was needed subsequently was only that it is square summable. One may wonder if that gives room to find more examples of λ for which ν_λ is singular. Actually no!

Result: (Pisot). If $\theta > 1$ and $\sum_j [\![t\theta^j]\!]^2 < \infty$ for some t, then θ is a PV number.

In fact, using this Salem showed that $\hat{\nu}_\lambda$ vanishes at infinity except when λ is the reciprocal of a PV number (observe that this is also true for $\lambda < \frac{1}{2}$). This is a somewhat harder exercise (optional).

Exercise 10. Show that $\hat{\nu}_\lambda(t) \to 0$ as $t \to \pm\infty$ if $\frac{1}{\lambda}$ is not a PV number.

CHAPTER 8

Equidistribution

1. Equidistribution on an interval

In this chapter, we shall write the circle as $T = [0, 1)$ with the identification $x \leftrightarrow e^{2\pi i x}$. A sequence $x = (x_n)_{n\geq 1}$ taking values in T is said to be equidistributed[17] if $\frac{1}{N}\sum_{k=1}^{n} \mathbf{1}_{x_k \in I} \to b - a$ for any interval $I = [a, b] \subseteq [0, 1]$.

Lemma 1. *Let $x = (x_n)_{n\geq 1}$ take values in T. The following are equivalent.*

(1) x is equistributed.

(2) $\frac{1}{N}\sum_{k=1}^{N} f(x_k) \to \int_0^1 f(x)$ for all $f \in C(T)$.

(3) $\frac{1}{N}\sum_{k=1}^{N} e_m(x_k) \to 0$ for all $m \in \mathbb{Z} \setminus \{0\}$ (as always, $e_m(x) = e^{2\pi i m x}$).

If one is familiar with the notion of convergence in distribution (weak convergence of probability measures), then all these are easily seen to be equivalent to the weak convergence of $\frac{1}{N}\sum_{k=1}^{N} \delta_{x_k}$ to the Lebesgue measure on $[0, 1]$. But as a direct argument is easy to give, we do that.

Proof. Two observations that will allow us to carry out the required approximations.

(a) If $\|f - g\|_{\text{sup}} < \varepsilon$, then $\left|\frac{1}{N}\sum_{k=1}^{N} f(x_k) - \frac{1}{N}\sum_{k=1}^{N} g(x_k)\right| < \varepsilon$ and $\left|\int_0^1 f(x)dx - \int_0^1 g(x)dx\right| < \varepsilon$.

(b) If $g \leq f \leq h$, then $\frac{1}{N}\sum_{k=1}^{N} g(x_k) \leq \frac{1}{N}\sum_{k=1}^{N} f(x_k) \leq \frac{1}{N}\sum_{k=1}^{N} h(x_k)$ and $\int_0^1 g(x)dx \leq \int_0^1 f(x)dx \leq \int_0^1 h(x)dx$.

Note that the definition of equidistribution is equivalent to the statement that $\frac{1}{N}\sum_{k=1}^{N} f(x_k) \to \int_0^1 f(x)$ for all step functions f.

Assume (1). Given any $f \in C(T)$, there exists a step function g such that $\|f - g\| < \varepsilon$. By the first observation, letting $N \to \infty$ we see that the limit point of $\frac{1}{N}\sum_{k=1}^{N} f(x_k)$ are within 2ε of $\int_0^1 f$. Hence (2) follows.

Assume (3). Then $\frac{1}{N}\sum_{k=1}^{N} f(x_k) \to \int_0^1 f(x)$ for all trigonometric polynomials f. By Fejér's theorem, they are dense in $C(T)$, hence again by the first observation we conclude (2).

Assume (2). Then (3) is obvious as $e_m \in C(T)$. To conclude (1), we use the second observation above. If $f = \mathbf{1}_{[a,b]}$, we may find $g, h \in C(T)$ such that $g \leq f \leq h$ and $\int(h - g) \leq \varepsilon$. By the second observation, letting $N \to \infty$, we see that the limit points of $\frac{1}{N}\sum_{k=1}^{N} g(x_k)$ are within 2ε of $\int_0^1 f$. Thus (1) follows. ■

[17] Equidistribution, however interpreted, is a large subject. What we cover in the first few sections (and much more) can be found in the book *Uniform distribution of sequences* by Kuipers and Niederreiter (John Wiley & Sons (1974)).

From the point of view of proving that a sequence is equidistributed, the third condition (called Weyl's criterion) is the most convenient, as it involves the least checking, and that too with particularly nice functions. This idea is at the heart of many things, including the use of characteristic functions to prove weak convergence (CLT for example) in probability theory.

As we shall be using this criterion to show that various sequences are equidistributed, let us start with an example that is not. Let us write $\overline{x}$ for $x \pmod 1$.

Example 2. If θ is a PV number, we saw that $[\![\theta^n]\!] \to 0$ as $n \to \infty$. Therefore, $\overline{\theta^n}$ is far from equidistributed. In contrast $\overline{t^n}$ is equidistributed for a.e. $t > 1$.

Exercise 3. A sequence $x = (x_n)_n$ in $T^d = [0,1)^d$ is said to be equidistributed if $\frac{1}{n}\sum_{k=1}^n f(x_k) \to \int_{T^d} f(x)dx$ for all $f \in C(T^d)$. Show that this is equivalent to either of the following statements:

(1) The convergence in the definition holds for $f = e_m$, $m \in \mathbb{Z}^d \backslash \{0\}$ where $e_m(x) = e^{2\pi i(m_1x_1+\ldots+m_dx_d)}$.

(2) $\frac{1}{n}\sum_{k=1}^n \mathbf{1}_{x_k\in I} \to \prod_{j=1}^d (b_j - a_j)$ for any rectangle $I = [a_1, b_1] \times \ldots \times [a_d, b_d] \subseteq [0,1]^d$.

2. Linear sequences

Theorem 4. *$(\overline{n\alpha})_{n\geq 1}$ is equidistributed if and only if $\alpha \notin \mathbb{Q}$.*

Proof. If $\alpha = \frac{p}{q} \in \mathbb{Q}$, then $\overline{n\alpha} = \overline{m\alpha}$ whenever $n - m$ is divisible by q. Therefore, the sequence takes only finitely many values periodically. Not equidistributed.

If $\alpha \notin \mathbb{Q}$, then fix $m \in \mathbb{Z} \setminus \{0\}$ and consider

$$\begin{aligned}\frac{1}{N}\sum_{k=1}^N e_m(\overline{\alpha n}) &= \frac{1}{N}\sum_{k=1}^N e^{2\pi i m\overline{k\alpha}} \\ &= \frac{1}{N}\sum_{k=1}^N e^{2\pi i m k\alpha} \\ &= \frac{1}{N}\frac{e^{2\pi i m(N+1)\alpha} - e^{2\pi i m\alpha}}{1 - e^{2\pi i m\alpha}}\end{aligned}$$

where we used the fact that $e^{2\pi i m\alpha} \neq 1$ as α is irrational. Clearly the last quantity is bounded by $\frac{2}{N|1-e^{2\pi i m\alpha}|}$ which goes to zero. By Weyl's criterion, equidistribution holds. ∎

3. Polynomial sequences

Let $P(x) = \alpha_d x^d + \ldots + \alpha_1 x + \alpha_0$ be a polynomial with real coefficients. Is $(\overline{P(n)})_{n\geq 1}$ equidistributed in $[0,1]$? We did not include the constant coefficient because it makes no difference to the equidistibution (just shifts by α_0 mod 1).

Theorem 5 (Weyl)**.** *The sequence $(\overline{P(n)})_{n\geq 1}$ equidistributed in $[0,1]$ if and only at least one of $\alpha_1, \ldots, \alpha_d$ is irrational.*

Observe that α_0 being irrational is of no help, as it only induces a shift in the sequence (modulo 1). In other words, if $(P(n))_{n\geq 1}$ is equidistributed if and only if $(P(n)-\alpha_0)_{n\geq 1}$ is equidistributed.

Further, one side of the theorem is easy. If $\alpha_1,\ldots,\alpha_d$ are rational, if $N\alpha_j \in \mathbb{Z}$ for all j for some N, and hence

$$\begin{aligned} P(kN+r) &= \alpha_d(kN+r)^d + \ldots + \alpha_1(kN+r) \\ &\equiv \alpha_d r^d + \ldots + \alpha_1 r \mod 1. \end{aligned}$$

Thus the sequence $(\overline{P(n)})_{n\geq 1}$ is N-periodic and cannot be equidistributed. It is the converse direction that is non-trivial and interesting. The key step is the following lemma that allows to reduce the degree.

Lemma 6 (van der Korput)**.** *Let $x = (\overline{x_n})_{n\geq 1}$ and for $h \geq 1$, let $x_h = (\overline{x_{n+h} - x_n})_{n\geq 1}$. If x_h is equidistributed in $[0,1]$ for all h, then x is equidistributed.*

Proof of Weyl's theorem assuming van der Korput's lemma. Let $x_n = P(n)$ so that $x(n+h) - x(n) = Q(n)$, where $Q_h(\cdot) = P(\cdot + h) - P(\cdot)$ is a polynomial of degree at most $d-1$. Write $P(x) = \alpha_d x^2 + \ldots + \alpha_1 x + \alpha_0$ and $Q_h(x) = \beta_{d-1}x^{d-1} + \ldots + \beta_1 x + \beta_0$. Choose $1 \leq \ell \leq d$ such that α_ℓ is irrational, but α_j is rational for $\ell < j \leq d$. Then it is easy to see that $\beta_{\ell-1}$ is also irrational. We must divide into two cases.

Case 1: If $\ell \geq 2$, then $\ell - 1 \geq 1$, hence Q_h also satisfies the conditions of the theorem. Inductively (the base case $d = 1$ of linear polynomials was taken care of earlier), we know that $(Q_h(n))_{n\geq 1}$ is equidistributed. As this applies for all h, by van der Korput's lemma we conclude that $(P(n))_{n\geq 1}$ is equidistributed.

Case 2: If $\ell = 1$, we can only conclude that the constant coefficient of Q_h is irrational, and it is of no use. Instead, we write $x_n = y_n + z_n$ where $y_n = P(n) - \alpha_1 n$ and $z_n = \alpha_1 n$. Observe that y is a periodic sequence (application of a polynomial with rational coefficients to natural numbers) and that z is equidistributed (as α_1 is irrational). From Exercise 7 below, it follows that x is equidistributed. ■

Exercise 7. Let $x_n = y_n + z_n$ where y is periodic and z is equidistributed. Then show that x is equidistributed. [*Hint:* You may use Weyl's criterion, or argue directly from the definition of equidistribution.]

Proof of van der Korput's lemma. We use Weyl's criterion again. Fix $m \in \mathbb{Z} \setminus \{0\}$ and let $v_k = e_m(x_k) = e_m(\overline{x_k})$. For fixed h,

$$\begin{aligned} \Big|\sum_{k=1}^{n} v_{k+h} - \sum_{k=1}^{n} v_k\Big| &= \Big|\sum_{k=n+1}^{n+h} v_k - \sum_{k=1}^{h} v_k\Big| \\ &\leq 2h. \end{aligned}$$

Average over $1 \leq h \leq H$ to get

$$\left|\frac{1}{H}\sum_{h=1}^{H}\sum_{k=1}^{n} v_{k+h} - \sum_{k=1}^{n} v_k\right| \leq H+1.$$

Hence,

$$\begin{aligned}
\left|\frac{1}{N}\sum_{k=1}^{n} v_k\right| &\leq \left|\frac{1}{HN}\sum_{h=1}^{H}\sum_{k=1}^{n} v_{k+h}\right| + \frac{H+1}{N} \\
&= \left|\frac{1}{N}\sum_{k=1}^{n}\frac{1}{H}\sum_{h=1}^{H} v_{k+h}\right| + \frac{H+1}{N} \\
&\leq \left(\frac{1}{N}\sum_{k=1}^{N}\left|\frac{1}{H}\sum_{h=1}^{H} v_{k+h}\right|^2\right)^{\frac{1}{2}} + \frac{H+1}{N} \qquad (1)
\end{aligned}$$

by Cauchy-Schwarz inequality. The quantity in the first summand under the square-root is

$$\frac{1}{NH^2}\sum_{k=1}^{N}\sum_{h_1,h_2=1}^{H} v_{k+h_1}\overline{v}_{k+h_2} = \frac{1}{NH^2}\sum_{h_1,h_2=1}^{H}\sum_{k=1}^{N} v_{k+h_1}\overline{v}_{k+h_2}.$$

For the H pairs (h_1, h_2) with $h_1 = h_2$, the inner summand is N. For $h_1 < h_2$, denoting $h = h_2 - h_1$, we see that

$$\begin{aligned}
\sum_{k=1}^{N} v_{k+h_1}\overline{v}_{k+h_2} &= \sum_{k=1}^{N} v_k\overline{v}_{k+h} + \sum_{k=N+1}^{N+h_1} v_k\overline{v}_{k+h} - \sum_{k=1}^{h_1} v_k\overline{v}_{k+h} \\
&= \sum_{k=1}^{N} v_k\overline{v}_{k+h} + R_h
\end{aligned}$$

where $|R_h| \leq 2H$. For $h_2 < h_1$, we get the same, except that it is conjugated. Therefore

$$\begin{aligned}
\frac{1}{NH^2}\sum_{h_1,h_2=1}^{H}\sum_{k=1}^{N} v_{k+h_1}\overline{v}_{k+h_2} &= \frac{1}{NH^2}\left\{HN + 2\sum_{h=1}^{H-1}(H-h)\left(\mathrm{Re}\left[\sum_{k=1}^{N} v_k\overline{v}_{k+h}\right] + R_h\right)\right\} \\
&\leq \frac{1}{H} + \frac{4H}{N} + \frac{2}{NH}\sum_{h=1}^{H}\left|\sum_{k=1}^{N} v_k\overline{v}_{k+h}\right|.
\end{aligned}$$

We simply bounded $H - h$ by H and the second term comes by summing up all the inequalities $|R_h| \leq 2H$. We plug this back into (1) while observing that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ and that $\frac{H+1}{N} \leq \frac{2H}{N}$ and $\frac{H}{N} \leq \sqrt{\frac{H}{N}}$.Therefore,

$$\begin{aligned}
\left|\frac{1}{N}\sum_{h=1}^{H} v_k\right| &\leq \sqrt{\frac{2}{NH}\sum_{h=1}^{H}\left|\sum_{k=1}^{N} v_k\overline{v}_{k+h}\right|} + \frac{1}{\sqrt{H}} + \frac{2\sqrt{H}}{\sqrt{N}} + \frac{H+1}{N} \\
&\leq \sqrt{\frac{2}{NH}\sum_{h=1}^{H}\left|\sum_{k=1}^{N} v_k\overline{v}_{k+h}\right|} + \frac{1}{\sqrt{H}} + \frac{4\sqrt{H}}{\sqrt{N}} \qquad (2)
\end{aligned}$$

This is known as van der Korput inequality and is the key technical tool. It is valid for any sequence v_k with $|v_k| = 1$.

Now substitute $v_k = e_m(x_k) = e_m(\overline{x}_k)$ for some $m \in \mathbb{Z} \setminus \{0\}$. We get

$$\Big|\frac{1}{N}\sum_{k=1}^{N} e_m(\overline{x}_k)\Big| \le \sqrt{\frac{2}{H}\sum_{h=1}^{H}\Big|\frac{1}{N}\sum_{k=1}^{N} e_m(\overline{x}_k - \overline{x}_{k+h})\Big|} + \frac{1}{\sqrt{H}} + \frac{4\sqrt{H}}{\sqrt{N}}.$$

Fix H and let $N \to \infty$. All the inner sums in the first summand converge to 0, by the assumption that x_h is equidistributed. Hence

$$\limsup_{N\to\infty}\Big|\frac{1}{N}\sum_{k=1}^{N} e_m(\overline{x}_k)\Big| \le \frac{1}{\sqrt{H}}.$$

Now let $H \to \infty$ to conclude that equidistribution holds for x. ■

4. Equidistribution of a few other elementary sequences

What are the eqidistribution properties of $(n^q)_{n\ge1}$ (for $0 < q < 1$ say) and $(\log n)_{n\ge1}$ and $(\theta^n)_{n\ge1}$.

▶ We have already seen that there are $\theta > 1$ for which $(\theta^n)_{n\ge1}$ is not equidistributed (even if we assume that θ^n is never an integer). It is apparently known that for almost every $\theta > 1$, this sequence is equidistributed, but we shall not go into that here.

▶ The sequence $(\log n)_{n\ge1}$ is not equidistributed. This is not hard to see and is a given in the problem set with a hint (one can use Weyl's criterion, or just the direct definition of equidistribution).

▶ Let $0 < q < 1$. Then $(n^q)_{n\ge1}$ is equidistributed. This follows from the following more general theorem. It has the flavour of the differencing trick, but is in fact more elementary. The assumption on differences is not of equidistribution but of monotonicity and growth.

Proposition 8. *Let $(x_n)_{n\ge1}$ be a sequence of real numbers. Let $y_n = x_{n+1} - x_n$. Assume that y_n is decreasing and satisfies $y_n \to 0$ and $ny_n \to \infty$. Then $(x_n)_n$ is equidistributed.*

Proof. ■

5. A quantitative equidistribution theorem

If a probability measure μ on $T = [0, 2\pi)$ has $\hat{\mu}(m) = 0$ for all $m \neq 0$, then μ must be the normalized Lebesgue measure m on T. If the first hundred Fourier coefficients are zero, can we say that μ is close to the Lebesgue measure? One must decide what is the sense of closeness one wants, and we take the *Kolmogorov-Smirnov* distance defined as $d(\mu, \nu) := \sup\{|\mu(I) - \nu(I)| : I \text{ is an arc in } T\}$. When we say arc of the circle T, we mean an interval $[a, b] \subseteq T$ for some $a < b$ or $[a, 2\pi) \cup [0, b]$ for some $b < a$. Then we have the following theorem of Erdös and Turan[18].

[18] Our presentation is taken from some unpublished notes of Mikhail Sodin and the paper *Equidistribution of zeros of polynomials* by Kannan Soundararajan. Both are exceedingly well-written and we have added almost nothing to the presentation. Sodin gives multiple proofs of the main step in the equidistribution result.

Lemma 9 (Erdös-Turan)**.** *For any probability measure μ on T and any $n \geq 1$, we have*

$$d(\mu, m) \leq C \left[\sum_{k=1}^{n} \frac{|\hat{\mu}(k)|}{k} \ + \ \frac{1}{n} \right]$$

for some constant C.

This is a quantitative version of Weyl's criterion. Indeed, if μ_n is a sequence of measures such that $\hat{\mu}_m(k) \to 0$ for all $k \neq 0$, then the above lemma implies that $\limsup d(\mu_m, m) \leq Cn^{-1}$, for any n, which of course shows that μ_n converges to m weakly.

Proof. Let $F_\mu(t) = \mu[0, t]$, $0 \leq t < 2\pi$ denote the distribution function of μ. For example, $F_m(t) = t/2\pi$. Let $V(t) = F_m(t) - F_\mu(t) - a$ where a is chosen so that $\int_T V(t)dt = 0$. Then, for any arc $I \subseteq T$, we have $|\mu(I) - m(I)| \leq 2\|V\|_{\text{sup}}$ (if $I = (a, b]$ then $\mu(I) - m(I) = V(b) - V(a)$, and similar expression if the arc is $(a, 1) \cup [0, b)$).

To bound the sup-norm of V, we smooth it by convolving with the Fejér kernel K_n to get $\sigma_n V(t) = (V \star K_n)(t) = \int_T V(t-s)K_n(s)ds$. Recall that $K_n(u) \leq \frac{1}{n\sin^2(u/2)} \leq \frac{\pi^2}{nu^2}$, from which we get $\int_{[-\delta,\delta]^c} K_n(u)\frac{du}{2\pi} \leq \frac{10}{n\delta}$. We divide into two cases, either $\|V\|_{\text{sup}} = \sup V$ or $\|V\|_{\text{sup}} = -\inf V$

First assume that $\|V\|_{\text{sup}} = \sup V$. Fix $\delta = 40/n$ and find $t \in T$ such that $V(t+\delta) \geq \|V\|_{\text{sup}} - \omega_V(2\delta)$. We have

$$\begin{aligned}
\sigma_n V(t) &= \int_{[-\delta,\delta]} V(t-s)K_n(s)\frac{ds}{2\pi} + \int_{[-\delta,\delta]} V(t-s)K_n(s)\frac{ds}{2\pi} \\
&\geq (V(t+\delta) - \omega_V(2\delta))(1 - \frac{10}{n\delta}) - \|V\|_{\text{sup}}\frac{10}{n\delta} \\
&\geq \frac{1}{2}\|V\|_{\text{sup}} - 2\omega_V(80/n).
\end{aligned}$$

Hence $\|\sigma_n V\|_{\text{sup}} \geq \frac{1}{2}\|V\|_{\text{sup}} - 2\omega_V(80/n)$.

In the other case, $\|V\|_{\text{sup}} = \sup(-V)$. Pick $\delta = 40/n$ and t such that $V(t-\delta) < -\|V\|_{\text{sup}} + \omega_V(2\delta)$. Then

$$\begin{aligned}
\sigma_n V(t) &= \int_{[-\delta,\delta]} V(t-s)K_n(s)\frac{ds}{2\pi} + \int_{[-\delta,\delta]} V(t-s)K_n(s)\frac{ds}{2\pi} \\
&\leq (V(t-\delta) + \omega_V(2\delta))(1 - \frac{10}{n\delta}) + \|V\|_{\text{sup}}\frac{10}{n\delta} \\
&\leq -\frac{1}{2}\|V\|_{\text{sup}} + 2\omega_V(80/n).
\end{aligned}$$

Thus again $\|\sigma_n V\|_{\text{sup}} \geq \frac{1}{2}\|V\|_{\text{sup}} - 2\omega_V(80/n)$.

Observe that $\|\sigma_n V\|_{\text{sup}} \leq |\hat{V}(0)| + 2\sum_{k=1}^{n-1} |\hat{V}(k)|$, and hence

$$\begin{aligned}
\|V\|_{\text{sup}} &\leq 2\|\sigma_n V\|_{\text{sup}} + 4\omega_V(80/n) \\
&\leq 4\sum_{k=0}^{n-1} |\hat{V}(k)| + 4\omega_V(10/n).
\end{aligned}$$

This conclusion holds for any $V \in L^1(T)$. For the particular V that we defined above, $\hat{V}(0) = 0$ and $\hat{V}(k) = \pm ik\hat{\mu}(k)$ for $k \neq 0$. Further, $\omega_V(\delta) \leq 2\delta$, hence the inequality $|\mu(I) - m(I)| \leq 2\|V\|_{\text{sup}}$ leads to

$$d(\mu, m) \leq C\left[\sum_{k=1}^{n-1} \frac{|\hat{\mu}(k)|}{k} + \frac{1}{n}\right]$$

This completes the proof. ■

Remark 10. One may care about the constants. As written, the proof gives $C = 80$, but that is because of the second summand. One can do better by writing the inequality in the form

$$d(\mu, m) \leq 4\left[\sum_{k=1}^{n-1} \frac{|\hat{\mu}(k)|}{k} + \frac{20}{n}\right]$$

Even these are not optimal, but we do not bother to do better.

Exercise 11. For simplicity, assume that $\|V\|_{\text{sup}}$ is attained, and choose t such that $t \pm \delta$ is such a point (depending on whether $\|V\|$ is equal to $\sup V$ or $-\inf V$). Then choose $\delta = c/n$ for a c as small as you can and get a better bound with explicit constants.

Remove the assumption that $\|V\|_{\text{sup}}$ is attained by introducing $\varepsilon > 0$ and choosing t such that either $V(t + \delta) \geq \|V\| - \varepsilon$ or $V(t - \delta) \leq -\|V\| + \varepsilon$, and finally letting $\varepsilon \to 0$.

6. Distribution of roots of polynomials

A polynomial of degree n can have any n complex numbers as its roots. But if one picks coefficients at random, often it turns out that the zeros are very close to the unit circle, and uniformly distributed around it. See Figure 6. Can one prove a theorem to this effect? What we see in the pictures can be captured in two statements:

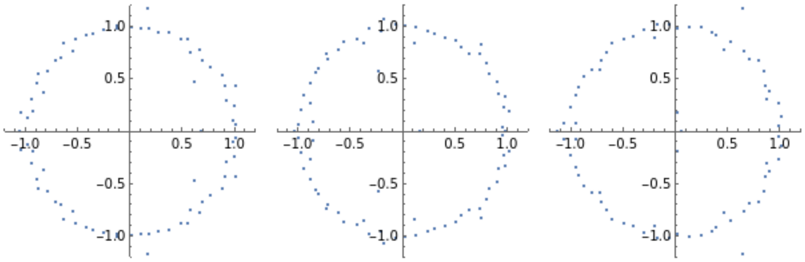

Figure 1. Roots of polynomials of degree 80. The coefficients are independent random variables with different distributions: Left: Random ±1. Middle: Gaussian. Right: Cauchy. There is no qualitative difference in the pictures!

(1) The absolute values of the roots are close to 1.

(2) The projections of the roots to the unit circle are approximately uniformly distributed on the circle.

It turns out that the first statement is relatively easy, and we give it at the end. The second one, on angular equidistribution, is covered by an amazing result of Erdös and Turan. Observe that there is no randomness in the statements!

Before stating the results, some notation. Both the radial and angular distributions will be controlled by the size of the polynomial on the unit circle. We can measure this in several ways. For a polynomial $P(z) = a_n z^n + \ldots + a_1 z + a_0$ with $a_0 a_n \neq 0$, define

(1) $h(P) := \int_0^{2\pi} \log_+ \frac{|P(e^{is})|}{\sqrt{|a_0 a_n|}} \frac{ds}{2\pi}$,

(2) $h_*(P) := \log \frac{\|P\|}{\sqrt{|a_0 a_n|}}$, where $\|P\| = \|P\|_{\sup(T)}$,

(3) $h_\#(P) := \log \frac{|a_0| + \ldots + |a_n|}{\sqrt{|a_0 a_n|}}$.

It is easy to see that

$$h(P) \leq h_*(P) \leq h_\#(P). \tag{3}$$

In many cases, it is easier to control $h_\#$ than h_* which in turn is easier to control than h. Hence, although the best inequalities are stated in terms of h, in using them we often replace h by $h_\#$.

6.1. **Angular distribution of roots.**

Theorem 12 (Erdös-Turan)**.** *Let $P(z) = a_n z^n + \ldots + a_1 z + a_0$ where $a_k \in \mathbb{C}$ and $a_0 a_n \neq 0$. Let the roots of P be $\zeta_k = r_k e^{i\theta_k}$, $1 \leq k \leq n$ (repeated according to multiplicity). Let $\mu = \frac{1}{n}\sum_{k=1}^n \delta_{e^{i\theta_k}}$. Then, for any arc $I \subseteq T$, we have*[19]

$$|\mu(I) - m(I)| \leq \frac{C}{\sqrt{n}}\sqrt{h(P)}.$$

For the right hand side to be a good bound, $h(P)$ must be small, or equivalently $\log_+ |P|$ must be small on the unit circle, on average. Using the inequality $h(P) \leq h_\#(P)$

Corollary 13. *Suppose P_n is a sequence of polynomials of degree n with coefficients having absolute values between B_n and $\frac{1}{B_n}$. Assume that $1 \leq B_n = e^{o(n)}$. Let μ_n be the probability measure on T that puts mass $\frac{1}{n}$ at $\zeta/|\zeta|$ for each root ζ of P (counted with multiplicity). Then, as $n \to \infty$,*

$$\sup\{|\mu_n(I) - m(I)| : I \text{ is an arc in } T\} \to 0.$$

Not all polynomials have equidistribution of the angular parts of the roots. What fails then?

[19] In Soundararajan's paper he gives the inequality with the explicit constant $C = \frac{8}{\pi}$.

Example 14. Let $P(z) = (z-1)^n$. All roots are at 1. What about $h(P)$? Observe that $|e^{it} - 1| > 1$ if and only if $\frac{\pi}{3} < t < \frac{5\pi}{3}$. Hence

$$h(P) = n \int_{\pi/3}^{5\pi/3} \log |e^{it} - 1| \, \frac{dt}{2\pi} \; = \; cn$$

for some $c > 0$. Thus the bound on $|\mu(I) - m(I)|$ does not go to zero as $n \to \infty$.

Exercise 15. Let $P(z) = z^n - 1$. Compute $h(P)$ explicitly, and compare $d(\mu, m)$ with the bound given by the Erdos-Turan theorem.

Proof of Erdös-Turan theorem. By the Erdös-Turan lemma on equidistribution, we know that for any $N \geq 1$,

$$|\mu(I) - m(I)| \leq C \left[\sum_{k=1}^{N-1} \frac{|\hat{\mu}(k)|}{k} + \frac{1}{N} \right].$$

Hence the issue is to get a control on $\hat{\mu}(k)$.

Case when roots are on T: Assume that $P(z) = a_n \prod_{j=1}^{n} (z - e^{it_j})$. Then $\mu = \frac{1}{n} \sum_{j=1}^{n} \delta_{e^{it_j}}$ and $\hat{\mu}(k) = \frac{1}{n} \sum_{j=1}^{n} e^{-ikt_j}$. Being a symmetric polynomial of the roots, these are expressible as polynomials of the coefficients of P (these are called Newton's identities), but the way we do it is as follows:

Claim: $\int_0^{2\pi} e^{iks} \log |e^{it} - e^{is}| \frac{ds}{2\pi} = \frac{e^{ikt}}{2|k|}$.

To see this, observe that if $r < 1$, then by the power series expansion of logarithm,

$$\log |e^{it} - re^{is}| = \mathrm{Re} \log(1 - re^{i(s-t)}) = -\,\mathrm{Re} \sum_{k=1}^{\infty} \frac{1}{k} r^k e^{ik(s-t)} = -\sum_{k \neq 0} \cdot \frac{1}{2|k|} r^{|k|} e^{ik(s-t)}.$$

Consequently, $\int_0^{2\pi} e^{iks} \log |e^{it} - re^{is}| \frac{ds}{2\pi} = \frac{r^{|k|} e^{ikt}}{2|k|}$. Now let $r \uparrow 1$ and argue that the integral on the left converges to $\int_0^{2\pi} e^{iks} \log |e^{it} - e^{is}| \frac{ds}{2\pi}$. This completes the proof of the claim.

Setting $t = t_j$ in the claim and summing over j gives us

$$\frac{1}{2|k|} \hat{\mu}(k) = \frac{1}{n} \int_0^{2\pi} e^{-iks} \log \frac{|P(e^{is})|}{|a_n|} \frac{ds}{2\pi}$$

which implies that

$$\frac{|\hat{\mu}(k)|}{|k|} \leq \frac{2}{n} \log \frac{\|P\|}{|a_n|}. \tag{4}$$

Hence the Erdös-Turan bound gives

$$|\mu(I) - m(I)| \leq C \left[\frac{2m}{n} \log \frac{\|P\|}{|a_n|} + \frac{1}{m} \right].$$

If we set $m = \sqrt{n} / \sqrt{2 \log \frac{\|P\|}{|a_n|}}$ (more precisely an integer close to this number), the two summands have about the same contribution and

$$|\mu(I) - m(I)| \leq \frac{C}{\sqrt{n}} \sqrt{\log \frac{\|P\|}{|a_n|}}.$$

Observe that our assumption that roots are on the unit circle forces $|a_0| = |a_n|$, hence the above expression can also be written as

$$|\mu(I) - m(I)| \leq \frac{C}{\sqrt{2n}}\sqrt{\log \frac{\|P\|}{\sqrt{|a_0 a_n|}}}.$$

On the right is the middle quantity in (3). How to get $h(P)$ is indicated later (it is simple, but we postpone it to avoid distractions from the main point).

General case when roots are anywhere in the plane: If $P(z) = a_n(z - \zeta_1)\dots(z - \zeta_n)$, set $\xi_k = \zeta_k/|\zeta_k|$ and $Q(z) = a_n(z - \xi_1)\dots(z - \xi_n)$. The angular distribution measures μ_P and μ_Q are equal by construction. Hence (to get the weaker bound as before) it suffices to show that

$$\frac{\|Q\|}{|a_n|} \leq \frac{\|P\|}{\sqrt{|a_0 a_n|}}$$

because we already know the theorem for Q. This trick of replacing P by Q is attributed to Schur. Writing $\zeta_j = r_j e^{i\theta_j}$ (then $\xi_j = e^{i\theta_j}$) and taking any $z = e^{i\alpha} \in T$, the desired inequality can be written as

$$\prod_{j=1}^{n} |e^{i\alpha} - e^{i\theta_j}| \leq \frac{1}{\sqrt{|a_0/a_n|}} \prod_{j=1}^{n} |e^{i\alpha} - r_j e^{i\theta_j}| \ = \prod_{j=1}^{n} \frac{|e^{i\alpha} - r_j e^{i\theta_j}|}{\sqrt{r_j}}$$

since a_0/a_n is the products of the roots, up to a sign. We show that the jth factor on the left is bounded by the jth factor on the right. This is easy, because for any α, θ,

$$\begin{aligned} r|e^{i\alpha} - e^{i\theta}|^2 - |e^{i\alpha} - re^{i\theta}|^2 &= r(2 - 2\cos(\alpha - \theta)) - (1 + r^2 - 2r\cos(\alpha - \theta)) \\ &= -(r-1)^2 \end{aligned}$$

which is negative. ■

Remark 16. How to improve the bound to $\sqrt{h(P)}$? We gave up too much in (4) by moving from the integral to the supremum. Instead, as $|x| = 2x_+ - x$, we can write

$$\begin{aligned} \frac{1}{2|k|}|\hat{\mu}(k)| &\leq \frac{2}{n}\int_0^{2\pi} \Big|\log \frac{|P(e^{is})|}{|a_n|}\Big| \frac{ds}{2\pi} \\ &= \frac{2}{n}\int_0^{2\pi} \log_+ \frac{|P(e^{is})|}{|a_n|}\frac{ds}{2\pi} - \frac{1}{n}\int_0^{2\pi} \log \frac{|P(e^{is})|}{|a_n|}\frac{ds}{2\pi} \\ &= \frac{2}{n}h(P) \end{aligned}$$

because the second integral vanishes. This is because the integral is the average of $\sum_{j=1}^{n} \log|z - e^{it_j}|$ which is harmonic inside the disk and 0 at $z = 0$. If this is not clear, show directly that $\int_0^{2\pi} \log|e^{is} - re^{it_j}|ds = 0$ if $r < 1$ (use Taylor expansion as we did earlier) and let $r \uparrow 1$. Once we get the $h(P)$ bound for polynomials with roots on the unit circle, for the general case it follows from the inequality $h(Q) \leq h(P)$. That in turn is true because we showed that $|Q(z)| \leq |P(z)|$ for $z \in T$.

6.2. **Radial distribution of roots.**

Theorem 17. *Let $P(z) = a_n z^n + \ldots + a_1 z + a_0$ with $a_0 a_n \neq 0$. Let $\zeta_j = r_j e^{it_j}$ for $1 \leq j \leq n$, be the roots, counted with multiplicity. Then*

$$\sum_{j=1}^{n} \log(r_j \vee \frac{1}{r_j}) \leq 2h(P).$$

In particular, if $\nu = \frac{1}{n}\sum_{j=1}^{n} \delta_{r_j}$, then for any $r < 1$,

$$\nu_P\left([0, r] \cup [r^{-1}, \infty)\right) \leq \frac{2h(P)}{n \log \frac{1}{r}}.$$

Corollary 18. *Suppose P_n is a sequence of polynomials of degree n with coefficients having absolute values between B_n and $\frac{1}{B_n}$. Assume that $1 \leq B_n = e^{o(n)}$. Then $\nu_{P_n}(1-\delta, 1+\delta) \to 0$ for any $\delta > 0$.*

Proof of Theorem 17. We claim that

$$\int_0^{2\pi} \log|e^{i\theta} - re^{it}| \frac{d\theta}{2\pi} = \begin{cases} 0 & \text{if } r \leq 1, \\ \log r & \text{if } r \geq 1. \end{cases}$$

For $r < 1$, this is the 0th Fourier coefficient of $\theta \mapsto \log|e^{i\theta} - re^{it}|$ that we saw earlier. For $r > 1$, rewrite the integrand as $\log r + \log|\frac{1}{r}e^{-it} - e^{-i\theta}|$ to reduce it to the previous case. The case $r = 1$ can be taken as a limitingg case from either direction.

Apply this with $r = r_j$, $t = t_j$ and sum up over $j \leq n$ to get

$$\int_0^{2\pi} \log \frac{|P(e^{i\theta})|}{|a_n|} = \sum_{j=1}^{n} \log_+ r_j.$$

Observe that $\frac{|a_0|}{|a_n|} = \prod_j r_j$ to rewrite this as

$$\int_0^{2\pi} \log \frac{|P(e^{i\theta})|}{\sqrt{|a_0 a_n|}} = -\frac{1}{2}\log\frac{|a_0|}{|a_n|} + \sum_{j=1}^{n} \log_+ r_j$$
$$= \frac{1}{2}\sum_{j=1}^{n} |\log r_j|.$$

Now $|\log r|$ is the same as $\log(r \vee \frac{1}{r})$ and the proof of the first statement is complete.

The second one follows from the first by observing that all the terms $\log r_j \vee \frac{1}{r_j}$ are positive, and each zero with absolute value in $(0, r] \cup [\frac{1}{r}, \infty)$ contributes at least $\log \frac{1}{r}$ to the sum. Hence their number is at most

$$\frac{2}{\log \frac{1}{r}} \int_0^{2\pi} \log \frac{|P(e^{i\theta})|}{\sqrt{|a_0 a_n|}}$$

Dividing by n gives the second statement. ■

CHAPTER 9

Expander graphs

1. Expansion in graphs

Let $G = (V, E)$ be a finite graph. This means that V (vertex set) is a finite set and E (edge set) is a subset of $\{\{i, j\} : i, j \in V\}$. Such graphs are called simple. A more general notion is to allow E to be a multiset, and also allow it to contain elements of the form $\{x, x\}$ for $x \in V$. An edge of the form $\{x, x\}$ is called a loop and if $\{x, y\}$ occurs k times, we say there are k edges connecting x to y. We shall use the notation as if our graphs are simple, although much of it extends to graphs with multi-edges and loops. Even more generally, one can consider weighted graphs, where $w_{i,j} = w_{j,i}$ is the weight of an edge between i and j (zero weight means no edge)[20].

The *adjacency matrix* $A_G = (a_{i,j})_{i,j\in V}$ where $a_{i,j} = 1$ if $i \sim j$ (i.e., $\{i, j\}$ is an edge) and 0 otherwise. For multiple edges, $a_{i,j}$ is the multiplicity of the edge. For weighted graphs $a_{i,j}$ is the weight. In any case, A_G is a symmetric matrix. The *Laplacian matrix* $L_G = D - A_G$ where $D = \mathrm{diag}(d_i)_{i\in V}$ is the diagonal matrix of degrees of the vertices (d_i is the number of edges connected to i, or more generally the row sum of the ith row of A_G). Then for $f \in \mathbb{R}^V$,

$$\langle Lf, f\rangle = \sum_{i\in V} d_i f(i)^2 - \sum_{\{i,j\}\in E} 2a_{i,j} f(i)f(j) = \sum_{\{i,j\}\in E} a_{i,j}(f(i) - f(j))^2$$

showing that L is positive semi-definite. It always has a zero eigenvalue, as $L\mathbf{1} = 0$. The normalized Laplacian is defined as $\mathcal{L} = D^{-\frac{1}{2}} L D^{-\frac{1}{2}}$.

For $S, T \subseteq V$, let $E(S, T)$ denote the set of edges with one end in S and the other in T. Then define the *expansion coefficient* of G as

$$h_G = \min_{S:|S|\le\frac{1}{2}n} \frac{|E(S, S^c)|}{|S|}.$$

Example 1. If $G = K_n$, the complete graph on n vertices, then $|E(S, S^c)| = |S| \times (n - |S|)$ and hence $h_G = \lceil \frac{1}{2}n \rceil$. If G is the discrete cycle on n vertices (edges are $\{1, 2\}, \dots, \{n-1, n\}, \{n, 1\}$), then $h_G \asymp \frac{1}{n}$.

Definition 2. A sequence of graphs $G_n = (V_n, E_n)$ is called an expander family if $|V_n| \to \infty$ and $\max_{i\in V_n} \deg(i) \le d$ for some $d < \infty$, and $h_{G_n} \ge h_0$ for some $h_0 > 0$.

[20] There are many good references for what we cover in this chapter. The survey article *Expander graphs and their applications* by Hoory, Linial and Wigderson; the book Expander graphs by E. Kowalski; the book/lecture notes on expanders by Luca Trevisan; the book *Spectral graph theory* by Fan Chung. These are some exceptionally well-written ones.

To understand the notion of expansion, consider a graph in which the degrees are bounded above by d. For $x \in V$ and $r \geq 1$, let $B(x, r)$ denote the set of vertices within graph distance r of x. If $|B(x, r)| \leq \frac{1}{2}n$, then $|B(x, r + 1)| - |B(x, r)| \geq \frac{h_G}{d}|B(x, r)|$, since there are at least $h_G|B(x, r)|$ that connect $B(x, r)$ to its complement, and at most d of them can have a common end-point on the other side (all these end-points are in $B(x, r+1) \setminus B(x, r)$). Hence, $|B(x, r+1)| \geq (1+\frac{h}{d})|B(x, r)|$, showing that $|B(x, r)| \geq (1 + \frac{h}{d})^r$ for $r \leq r(x) := \min\{r : |B(x, r)| \geq n/2\}$. The balls increase exponentially in size till half of the graph is covered. This shows that

$$r(x) \leq \frac{\log n}{\log(1 + \frac{h}{r})} \text{ and diameter}(G) \leq \frac{2 \log n}{\log(1 + \frac{h}{r})}$$

since $B(x, r(x)) \cap B(y, r(y)) \neq \emptyset$ for any x, y. Since we also have the trivial bound $|B(x, r)| \leq 1 + d + \ldots + d^r \asymp d^r$, no connected graph can have diameter more than $c_d \log r$. Thus in an expander sequence, the graphs G_n have diameters growing at the order of $\log |V_n|$.

2. Connection between expansion and spectrum

Henceforth, we assume that the graph $G = (V, E)$ is d-regular. Let the eigenvalues of the Laplacian L_G be denoted $0 = \lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_n$. Then the eigenvalues of A_G are $d = \tilde{\lambda}_1 \geq \tilde{\lambda_2} \geq \ldots \geq \tilde{\lambda}_n \geq -d$ (the last inequality is because the row sums of A_G are at most d).Of these, two are important enough that we give separate notations: $\lambda_G := \lambda_2$ and $\tilde{\lambda}^*_G = \max\{\tilde{\lambda}_2, |\tilde{\lambda}_n|\}$. Two easy observations:

(1) $\lambda_G = 0$ if and only if G is disconnected. In fact, since $\langle Lv, v\rangle = \sum_{i\sim j}(v_i - v_j)^2$, we see that $Lv = 0$ if and only if v is constant on each connected component of G. Therefore, the dimension of the null space of L_G is equal to the number of connected components of G.

(2) $\tilde{\lambda}_n = -d$ if and only if G is bipartite. If G is bipartite, then $V = V_1 \sqcup V_2$ and all edges of G connect a vertex in V_1 to a vertex in V_2. If $v_i = 1$ for $i \in V_1$ and $v_i = -1$ for $i \in V_2$, check that $Lv = -dv$. We leave the converse as an exercise.

Now we go on to more quantitative way in which λ_G measures how well-connected G is. This is a fundamental result in spectral graph theory.

Theorem 3 (Cheeger's inequality+Buser's inequality)**.** *Let G be a finite d-regular graph. Then*

$$\frac{h_G^2}{2d} \leq \lambda_G \leq 2h_G.$$

The key point for us is that expansion can be captured in terms of λ_G. In particular, $\{G_n\}$ is an expander sequence if and only if λ_{G_n} is bounded below. Incidentally, the inequality holds as stated for general finite graphs, if d is interpreted as the maximum degree among all vertices - this follows from the proof given below.

Before going to the proof, recall the Rayleigh-Ritz formulas (variational principles for eigenvalues of Hermitian martices) which imply that

$$\lambda_G = \min_{f \in \mathbb{R}^V, f \perp 1} \frac{\langle L_G f, f\rangle}{\langle f, f\rangle}.$$

Any f that attains the minimum is an eigenvector with eigenvalue λ_G.

Proof. The second inequality (Buser's) is the easier one. Given any $S \subseteq V$ with cardinality $s \le \frac{1}{2}n$, define $f_i = 1 - \frac{s}{n}$ for $i \in S$ and $f_i = -\frac{s}{n}$ for $i \in S^c$. This is just the indicator of S, shifted to have zero mean. Then

$$\langle f, f\rangle = (1 - \frac{s}{n})^2 s + \left(\frac{s}{n}\right)^2 (n-s) \; = \; \frac{s(n-s)}{n} \; \ge \; \frac{s}{2},$$
$$\langle Lf, f\rangle = \sum_{i \sim j} (f(i) - f(j))^2 \; = \; \sum_{i \sim j : i \in S, j \in S^c} 1 \; = \; |E(S, S^c)|.$$

By the variational formula,

$$\frac{2|E(S,S^c)|}{|S|} \ge \frac{\langle Lf, f\rangle}{\langle f, f\rangle} \ge \lambda_G.$$

To prove the first inequality (Cheeger's), we let f be an eigenvector of L_G with eigenvalue λ_G. Label the vertices so that $f(1) \ge \ldots \ge f(n)$ and fix $k = \lfloor \frac{n}{2} \rfloor$ (the reason for this choice becomes clear later), and define two other vectors

$$g(i) = \begin{cases} f(i) - f(k) & \text{if } i < k, \\ 0 & \text{if } i \ge k, \end{cases} \qquad \text{and} \qquad h(i) = \begin{cases} 0 & \text{if } i \le k, \\ f(k) - f(i) & \text{if } i > k. \end{cases}$$

Then we claim that

$$\|g\|^2 + \|h\|^2 \ge \|f\|^2 \qquad \text{and } \langle Lg, g\rangle + \langle Lh, h\rangle \le \langle Lf, f\rangle. \tag{1}$$

The first inequality is easy as the left side is $\sum_i (f(i) - f(k))^2$ which is equal to $\sum_i f(i)^2 + nf(k)^2$, as $f(1) + \ldots + f(n) = 0$. To prove the second, consider any edge $i \sim j$. If $i < j < k$, then $|g(i) - g(j)| = |f(i) - f(j)|$ and $|h(i) - h(j)| = 0$ while if $k \le i < j$ then $|g(i) - g(j)| = 0$ and $|h(i) - h(j)| = |f(i) - f(j)|$. If $i < k \le j$, then

$$(g(i) - g(j))^2 + (h(i) - h(j))^2 = (f(i) - f(k))^2 + (f(k) - f(j))^2 \le (f(i) - f(j))^2.$$

The ordering of vertices was used in the second inequality because then $(f(i) - f(k))(f(k) - f(j)) \ge 0$. Summing over all edges justifies the second inequality in (1). Consequently,

$$\lambda_G \ge \frac{\langle Lf, f\rangle}{\langle f, f\rangle} \ge \frac{\langle Lg, g\rangle + \langle Lh, h\rangle}{\|g\|^2 + \|h\|^2} \ge \frac{\langle Lg, g\rangle}{\|g\|^2} \wedge \frac{\langle Lh, h\rangle}{\|h\|^2}.$$

Without loss of generality assume that the first one is the smaller of the two. To analyse it, observe that $\sum_{i\sim j}(g(i)+g(j))^2 \le 2\sum_{i\sim j}(g(i)^2+g(j)^2) = 2d\sum_i g(i)^2$. Hence

$$\begin{aligned}\frac{\langle Lg,g\rangle}{\|g\|^2} &\ge \frac{\sum_{i\sim j}(g(i)-g(j))^2}{\|g\|^2} \times \frac{\sum_{i\sim j}(g(i)+g(j))^2}{2d\|g\|^2} \\ &\ge \frac{\left(\sum_{i\sim j}(g(i)^2-g(j)^2)\right)^2}{2d\|g\|^4}\end{aligned}$$

by Cauchy-Schwarz inequality. Now,

$$\begin{aligned}\sum_{i\sim j}(g(i)^2-g(j)^2) &= \sum_{i\sim j,\ i<j}\sum_{\ell=i}^{j-1}(g(\ell)^2-g(\ell+1))^2 \\ &= \sum_{\ell}(g(\ell)^2-g(\ell+1))^2|\{i\sim j : i\le \ell<j\}|.\end{aligned}$$

The second factor in the summand is $|E([\ell],[\ell]^c)|$ which is at least $h_G\ell$ if $\ell \le \frac{1}{2}n$. Of course, $g(\ell)-g(\ell+1)=0$ if $\ell\ge k$, hence we see that

$$\begin{aligned}\sum_{i\sim j}(g(i)^2-g(j)^2) &= h_G\sum_{\ell=1}^{k-1}(g(\ell)^2-g(\ell+1)^2)\ell \\ &= h_G\sum_{\ell=1}^{n}g(\ell)^2.\end{aligned}$$

Thus, we have arrived at

$$\lambda_G \ge \frac{\langle Lg,g\rangle}{\|g\|^2} \ge \frac{(h_G\|g\|^2)^2}{2d\|g\|^4} = \frac{h_G^2}{2d}$$

which is Cheeger's inequality. ■

Later we shall need an inequality like Cheeger's for infinite graphs. While one can try to define the Laplacian and make sense of it as a self-adoint operator, study its spectrum, etc., for our purposes we can cut short all that and directly prove an inequality between quadratic forms and expansion. Follow the above proof to complete the following exercise.

Exercise 4. Let $G=(V,E)$ be a d-regular graph with a countable vertex set V. Then for any $g\in L^2(V)$, show that

$$\frac{\sum_{i\sim j}(g(i)-g(j))^2}{\sum_{i\in V}g(i)^2} \ge \frac{1}{2d}\left(\inf_{S\subseteq V,\ |S|<\infty}\frac{|E(S,S^c)|}{|S|}\right)^2$$

Observe that the condition that $g\perp \mathbf{1}$ and the condition that $|S|\ge\frac{1}{2}|V|$ are neither needed nor meaningful.

Returning to finite graphs, as discussed earlier, these inequalities allow us to define expanders algebraically in terms of the second eigenvalue. If we take this definition, the following result places a limitation on how good the expansion can be.

Theorem 5 (Alon–Boppana)**.** *Let G be a d-regular random graph. Then $\lambda_G \leq d - 2\sqrt{d-1} + \delta_{n,d}$ where $\delta_{n,d} \to 0$ as $n \to \infty$ for fixed d.*

The usual way to state this is that $\tilde{\lambda}_2 \geq 2\sqrt{d-1} - o(1)$. This inequality explains why the following definition is meaningful.

Definition 6. A sequence of d-regular graphs $G_n = (V_n, E_n)$ with $|V_n| \to \infty$ is said to be *Ramanujan* if $\lambda_{G_n} \geq d - 2\sqrt{d-1}$ for all n.

Thus, Ramanujan graphs are the most extreme possible expanders. Constructing them is much harder than constructing general expanders. It was done first by Lubotsky–Phillip–Sarnak using some number theory results related to Ramanujan conjectures, and hence they gave the name. Their construction was for specific d (of the form 1+prime). Marcus–Spielman–Srivastava recently constructed Ramanujan graphs of all degrees. In this chapter we shall only talk about general expanders.

While the proof of the Alon-Boppana bound requires some work, it is not hard to get some bounds of this nature. For example, writing the trace of A^2 in terms of eigenvalues we see that $\text{tr}(A_G^2) \leq d^2 + (n-1)\tilde{\lambda}_*^2$. On the other hand, the trace is also the sum of squares of all the entries of A_G, hence $\text{tr}(A_G^2) = nd$. Thus, we see that $\tilde{\lambda}_* \geq \sqrt{d}$. The loss of the factor of 2 on the right (at least if d is large enough for us to ignore the difference between d and $d-1$) can be fixed by considering higher powers $\text{tr}(A_G^{2p})$. On one side we can bound it using eigenvalues and on the other side one can relate it to the number of closed paths of length $2p$ on the graph. Some analysis (mainly the idea that for a given starting point on the graph G, there are at least as many closed paths of a given length $2p$ as there are for a given starting point on a d-regular tree) leads to a weaker form of Alon-Boppana bound that says that

$$\tilde{\lambda}_* \geq 2\sqrt{d-1} - o(1).$$

The weakness is because $\tilde{\lambda}_2 \leq \tilde{\lambda}_*$.

3. Construction of expanders

The original Margulis construction with important improvements and simplications by Gabber and Galil, and then many others is presented in many places. We just give an outline and refer the reader to these sources[21].

The graph G_n is defined as follows: Let $V_n = \mathbb{Z}_n \times \mathbb{Z}_n = \{0, 1 \ldots, n-1\}^2$. The edges adjacent to $u = (k, \ell) \in V_n$ connect it to $u \pm e_1 = (k \pm 1, \ell)$, $u \pm e_2 = (k, \ell \pm 1)$, $S(u) = (k + \ell, \ell)$, $S^{-1}(u) = (k - \ell, \ell)$, $T(u) = (k, k + \ell)$, $T^{-1}(u) = (k, \ell - k)$ (all addition is modulo n).This allows multiple edges and loops, but is a d-regular graph.

Theorem 7. *G_n is an expander family. In fact $\lambda_{G_n} \geq$??.*

[21] The presentation in Luca Trevisan's notes is superb and we follow it here closely.

If f is the second eigenvector of L_{G_n}, then we know that $f \perp \mathbf{1}$ and

$$\lambda_{G_n} = \frac{\sum_{u\in V_n}(f(u)-f(Su))^2+(f(u)-f(Tu))^2+(f(u)-f(u+e_1))^2+(f(u)-f(u+e_2))^2}{\sum_{u\in V_n} f(u)^2}.$$

Correspnding to f, we define a function $F : [0,n)^2 \mapsto \mathbb{R}$ by $F(x) = f(u)$ if $x \in u + [0,1)^2$, $u \in V_n$. To convert it to our standard convention for the torus, also define $G : [0,2\pi)^2 \mapsto \mathbb{R}$ by $G(x) = F(nx/2\pi)$. Then

(1) $\int_{T^2} G(x)dx = \frac{1}{n}\int_{[0,n)^2} F(x)dx = \frac{1}{n}\sum_{u\in G_n} f(u) = 0$.
(2) $\int_{T^2} G(x)^2 dx = \frac{1}{n^2}\int_{[0,n)^2} F(x)^2 = \frac{1}{n^2}\sum_{u\in V_n} f(u)^2$.
(3) Further, if we define $S, T : [0,n)^2 \mapsto [0,n)^2$ by $S(x_1,x_2) = (x_1+x_2, x_2)$ and $T(x_1,x_2) = (x_1, x_1+x_2)$, then if $x \in u+[0,1)^2$, then $S(x) \in S(u)+[0,1)^2$ or $S(x) \in S(u)+e_1+[0,1)^2$. Hence, if we define $\bar{S} : T^2 \mapsto T^2$ by $\bar{S}(x) = \frac{2\pi}{n}S(nx/2\pi) = (x_1+x_2, x_2)$ (the addition is in T^2, i.e., modulo 2π), then we can work out that

$$\begin{aligned}\int_{T^2}(G(x)-G(\bar{S}x))^2 &= \frac{1}{n^2}\int_{[0,n)^2}(F(x)-F(S(x)))^2\\ &= \frac{1}{n^2}\sum_{u\in V_n}(f(u)-f(S(u))^2+(f(u)-f(Su+e_1))^2\\ &\le \frac{1}{n^2}\sum_{u\in V_n}(f(u)-f(S(u))^2+2[(f(u)-f(Su))^2+(f(Su)-f(Su+e_1))^2]\\ &\le \frac{3}{n^2}\sum_{u\in V_n}(f(u)-f(S(u))^2+(f(u)-f(u+e_1))^2.\end{aligned}$$

Adding it to the analogous identity for T in place of S, we find that

$$\begin{aligned}&\int_{T^2}(G(x)-G(\bar{S}x))^2+(G(x)-G(\bar{T}x))^2\\ &=\sum_{u\in V_n}(f(u)-f(Su))^2+(f(u)-f(Tu))^2+(f(u)-f(u+e_1))^2+(f(u)-f(u+e_2))^2.\end{aligned}$$

Consequently,

$$3\lambda_{G_n} \ge \frac{\int_{T^2}(G(x)-G(\bar{S}x))^2+(G(x)-G(\bar{T}x))^2}{\int_{T^2} G(x)^2}.$$

To estimate the right hand side, we use Fourier analysis on T^2. Write

$$G(x) \overset{L^2(T^2)}{=} \sum_{p\in\mathbb{Z}^2}\hat{G}(p)e^{i(p_1x_1+p_2x_2)}$$

Note that $\hat{G}(0,0) = \int_{T^2} G = 0$. We now compute the Fourier expansions of $G \circ \overline{S}$ and $G \circ \overline{T}$.

$$\begin{aligned} G(\overline{S}x) = G(x_1 + x_2, x_2) \ &= \sum_{p \in \mathbb{Z}^2} \hat{G}(p) e^{i(p_1(x_1+x_2)+p_2x_2)} \\ &= \sum_{p \in \mathbb{Z}^2} \hat{G}(p_1, p_2) e^{i(p_1x_1+(p_1+p_2)x_2)} \\ &= \sum_{p \in \mathbb{Z}^2} \hat{G}(q_1, q_2 - q_1) e^{i(q_1x_1+q_2x_2)}. \end{aligned}$$

This means that $(\widehat{G \circ \overline{S}})(p_1, p_2) = \hat{G}(p_1, p_2 - p_1)$. Similary, check that $(\widehat{G \circ \overline{T}})(p_1, p_2) = \hat{G}(p_1 - p_2, p_2)$. To state this clearly, let us introduce $\tilde{S}, \tilde{T} : \mathbb{Z}^2 \mapsto \mathbb{Z}^2$ by $\tilde{S}(p_1, p_2) = (p_1 + p_2, p_2)$ and $\tilde{T}(p_1, p_2) = (p_1, p_2 + p_1)$ where the addition is in $\mathbb{Z}^2$ (in retrospect, it would have been simpler to have just defined S, T on any group and understand from the context which one is being used). Then,

$$\widehat{G \circ \overline{S}} = \hat{G} \circ \tilde{T}^{-1}, \qquad \widehat{G \circ \overline{T}} = \hat{G} \circ \tilde{S}^{-1}.$$

This is the key observation which explains the choice of the two maps S, T in defining G_n. Using the Fourier expansions of $G, G \circ \overline{S}, G \circ \overline{T}$ and Plancherel's theorem, we see that

$$\begin{aligned} \int_{T^2} G(x)^2 &= (2\pi)^2 \sum_{p \in \mathbb{Z}^2} |\hat{G}(p)|^2 \\ \int_{T^2} |G(x) - G(\overline{S}x)|^2 &= (2\pi)^2 \sum_{p \in \mathbb{Z}^2} |\hat{G}(\tilde{T}p) - \hat{G}(p)|^2 \\ \int_{T^2} |G(x) - G(\overline{T}x)|^2 &= (2\pi)^2 \sum_{p \in \mathbb{Z}^2} |\hat{G}(\tilde{S}p) - \hat{G}(p)|^2 \end{aligned}$$

and hence

$$2\lambda_{G_n} \geq \frac{\sum_{p \in \mathbb{Z}^2} |\hat{G}(p) - \hat{G}(\tilde{T}p)|^2 + \sum_{p \in \mathbb{Z}^2} |\hat{G}(p) - \hat{G}(\tilde{S}p)|^2}{\sum_{p \in \mathbb{Z}^2} |\hat{G}(p)|^2}.$$

Observe that all the sums are over $\mathbb{Z}^2 \setminus \{(0,0)\}$ (and $\tilde{S}, \tilde{T}$ map $\mathbb{Z}^2 \setminus \{(0,0)\}$ into itself). If we define an graph G with vertex set $\mathbb{Z}^2 \setminus \{(0,0)\}$ and edges $p \sim \tilde{T}p$ and $p \sim \tilde{S}p$, then the right hand side above is precisely the Rayleigh-Ritz quotient for the Laplacian on this graph. From Cheeger's inequality for infinite graphs as given in Exercise 4, we deduce that

$$3\lambda_{G_n} \geq \frac{1}{8} \left(\inf_{S \subseteq \mathbb{Z}^2,\ |S| < \infty} \frac{|E(S, S^c)|}{|S|} \right)^2.$$

It looks like we are back to where we started. Instead of the graph G_n, we now have the infinite graph and we must show that the expansion coefficient is strictly positive. Turns out, this can be done directly by an elementary argument! It is tempting to think that may be a variant of this argument can be directly carried out for the original graph G_n, but I have not seen such a proof anywhere and I am even unable to visualize the graph G_n well.

Take any finite $S \subseteq \mathbb{Z}^2 \setminus \{(0,0)\}$. In Trevisan's notes, he partitions it into $S_0, S_1, \dots, S_4$, where $S_1, \dots, S_4$ are those points of S that lie in the (strict) first to fourth quadrants and S_0 consists of all the other vertices (those that have at least one zero co-ordinate). We quote (and leave it to you to work it out or refer to Trevisan's notes)

(1) S_1 has at least S_1 edges that go out of S and connect it to vertices in the first quadrant. Hence deduce that $|E(S_1 \cup S_2 \cup S_3 \cup S_4, S^c)| \geq |S_1 \cup S_2 \cup S_3 \cup S_4|$.

(2) The $4|S_0$ edges with one vertex in S_0 have the other in S_0^c, but only $3/4$ of these can land in S (from the first step). Hence deduce that $|E(S_0, S^c)| \geq 7|S_0| - 3|S|$.

Use the two inequalities to deduce that $|E(S, S^c)| \geq \frac{1}{7}|S|$.

Putting everything together, we have

$$\lambda_{G_n} \geq \frac{1}{3} \times \frac{1}{8} \times \frac{1}{49} = \frac{1}{1176}.$$

This completes the proof. ■