smoother (smoothed out at scale εp) and then compute the Fourier expansion. Apply Plancherel's theorem (4.3) with this smoothed out function and $1_A - \mathbf{P}(A)$.) This inequality is a crude form of the famous *Erdős–Turán inequality* in discrepancy theory, and is related to the Weyl criterion for uniform distribution modulo one.

4.3.14 Let $A = \mathbf{Z}_p^2$ be the set of squares in a cyclic group of prime order. Show that for any arithmetic progression P in $\mathbf{Z}_p$, we have

$$|A \cap P| = \frac{1}{2}|P| + O(\sqrt{p}\log p).$$

(Hint: use Lemma 4.14 and the preceding exercise.) This is a special case of the *Polya–Vinogradov inequality* from analytic number theory.

4.3.15 Let F be a finite field, let Z be a vector space over F, and let $M : Z \to Z$ be a linear transformation. Show that if $\dim_F(Z) \geq 3$, then there exists a non-zero $x \in Z$ such that $Mx \cdot x = 0$. (Hint: reduce to the case when M has full rank, and then modify Lemma 4.14. One can also solve this problem by purely algebraic methods.)

4.3.16 [160] Let W be a vector space over a finite field F of odd order, and let $M : W \to W$ be a linear transformation. Show that there exists a subspace U of W with dimension $\dim_F(U) \geq \frac{1}{2}\dim_F(W) - \frac{3}{2}$ such that M is null on U, i.e. $Mx \cdot y = 0$ for all $x, y \in U$. (Hint: take a maximal space U which is null with respect to M. If the orthogonal complement $U^\perp := \{y \in W : Mx \cdot y = 0 \text{ for all } x \in U\}$ is at least three dimensions larger than U, then use the previous lemma.) For a purely algebraic proof of this fact, see Exercise 9.4.11.

4.4 Bohr sets

In many applications of the Fourier-analytic method, one starts with some additive set A and concludes some information about the Fourier transform $\hat{1}_A$ of A (for instance, one may obtain some bound on the Fourier bias $\|A\|_u$). One would then like to pass from this back to some new combinatorial information on the original set A. For some special groups (e.g. finite field geometries F_p^n) one can do this quite directly (see for instance Lemma 10.15). However, to convert Fourier information on general groups to combinatorialinformation we need the notion of a *Bohr set* (also known as *Bohr neighborhoods* in the literature). We first define a "norm" $\|\theta\|_{\mathbf{R}/\mathbf{Z}}$ on the circle group by defining $\|\theta + \mathbf{Z}\|_{\mathbf{R}/\mathbf{Z}} = |\theta|$ whenever $-1/2 < \theta \leq 1/2$; in other words, $\|\theta\|_{\mathbf{R}/\mathbf{Z}}$ is the distance from θ (or more precisely, any representative of the coset θ) to the integers. We observe the elementary bounds

$$4\|\theta\|_{\mathbf{R}/\mathbf{Z}} \leq |e(\theta) - 1| \leq 2\pi\|\theta\|_{\mathbf{R}/\mathbf{Z}} \tag{4.24}$$

which follow from elementary trigonometry and the observation that the sinc function $\sin(x)/x$ varies between 1 and $2/\pi$ when $|x| \leq \pi/2$.

Definition 4.17 (Bohr set) Let $S \subset Z$ be a set of frequencies, and let $\rho > 0$. We define the *Bohr set* $\mathrm{Bohr}(S, \rho) = \mathrm{Bohr}_Z(S, \rho)$ as

$$\mathrm{Bohr}(S, \rho) := \left\{x \in Z : \sup_{\xi \in S} \|\xi \cdot x\|_{\mathbf{R}/\mathbf{Z}} < \rho\right\}.$$

We refer to S as the *frequency set* of the Bohr set, and ρ as the *radius*. The quantity $|S|$ is known as the *rank of the Bohr set*.

Remark 4.18 Note that if Z is a vector space over a finite field F, then every subspace of Z can be viewed as a Bohr set (with radius $O(1/|F|)$, and rank equal to the codimension). Thus Bohr sets can be viewed as a generalization of subspaces. Note that most finite groups Z tend to have very few actual subgroups (the extreme case being the cyclic groups $\mathbf{Z}_p$ of prime order), so it is convenient to be able to rely on the much larger class of Bohr sets as a substitute.

Remark 4.19 One way to think of Bohr sets is to consider the embedding of Z into the complex vector space $\mathbf{C}^S$ (and in particular to the standard unit torus inside $\mathbf{C}^S$) by the multiplicative map $x \mapsto (e(\xi \cdot x))_{\xi \in S}$. A Bohr set is thus the inverse image of a cube.

Observe that the $\| \|_{\mathbf{R}/\mathbf{Z}}$ norm is symmetric and subadditive; $\|-x\|_{\mathbf{R}/\mathbf{Z}} = \|x\|_{\mathbf{R}/\mathbf{Z}}$ and $\|x + y\|_{\mathbf{R}/\mathbf{Z}} \leq \|x\|_{\mathbf{R}/\mathbf{Z}} + \|y\|_{\mathbf{R}/\mathbf{Z}}$. From this we see that the Bohr sets $\mathrm{Bohr}(S, \rho)$ are symmetric, decreasing in S, and increasing in ρ (and fill out the whole space Z once $\rho > 1/2$); they are always unions of cosets of $S^\perp$, and if ρ is sufficiently small they consist entirely of $S^\perp$. One can also easily verify the intersection property

$$\mathrm{Bohr}(S, \rho) \cap \mathrm{Bohr}(S', \rho) = \mathrm{Bohr}(S \cup S', \rho)$$

and the addition property

$$\mathrm{Bohr}(S, \rho) + \mathrm{Bohr}(S, \rho') \subseteq \mathrm{Bohr}(S, \rho + \rho').$$

In particular we have

$$k\mathrm{Bohr}(S, \rho) \subseteq \mathrm{Bohr}(S, k\rho)$$

for any $k \geq 1$.

Next, we establish some bounds for the size of Bohr sets.

Lemma 4.20 (Size bounds) *If $S \subset Z$ and $\rho > 0$, then we have the lower bound*

$$\mathbf{P}_Z(\mathrm{Bohr}(S, \rho)) \geq \rho^{|S|} \tag{4.25}$$

and we have the doubling estimate

$$\mathbf{P}_Z(\mathrm{Bohr}(S, 2\rho)) \leq 4^{|S|}\mathbf{P}_Z(\mathrm{Bohr}(S, \rho)). \tag{4.26}$$

This lemma should be compared with the Kronecker approximation theorem (Corollary 3.25); indeed the two results are very closely related.

Proof For each $\xi \in S$ let θ_ξ be an element of $\mathbf{R}/\mathbf{Z}$ chosen independently and uniformly at random. For any $x \in Z$, one can easily verify that

$$\mathbf{P}_Z(\|\xi \cdot x - \theta_\xi\|_{\mathbf{R}/\mathbf{Z}} < \rho/2 \text{ for all } \xi \in S) = \rho^{|S|}.$$

Summing this over all $x \in Z$ using linearity of expectation (1.4), we conclude

$$\mathbf{E}|\{x \in Z : \|\xi \cdot x - \theta_\xi\|_{\mathbf{R}/\mathbf{Z}} < \rho/2 \text{ for all } \xi \in S\}| \geq \rho^{|S|}|Z|$$

and thus there exists a choice of θ_ξ such that

$$|\{x \in Z : \|\xi \cdot x - \theta_\xi\|_{\mathbf{R}/\mathbf{Z}} < \rho/2 \text{ for all } \xi \in S\}| \geq \rho^{|S|}|Z|. \tag{4.27}$$

Now observe from the triangle inequality that if x, x' lie in the above set, then $x - x'$ lies in $\mathrm{Bohr}(S, \rho)$. The claim (4.25) follows.

Now we prove (4.26). By a limiting argument we may replace 2ρ by $2\rho - \varepsilon$ on the left-hand side for some small $\varepsilon > 0$. Observe that we can cover the interval $\{\theta \in \mathbf{R}/\mathbf{Z} : \|\theta\|_{\mathbf{R}/\mathbf{Z}} < 2\rho - \varepsilon\}$ by four intervals of the form $\{\theta \in \mathbf{R}/\mathbf{Z} : \|\theta - \theta_0\|_{\mathbf{R}/\mathbf{Z}} < \rho/2\}$. We can thus can cover $\mathrm{Bohr}(S, 2\rho)$ by $4^{|S|}$ sets of the type appearing in the left-hand side of (4.27). The claim follows by arguing as before. □

We have already mentioned that subspaces of a vector space are one example of a Bohr set. Progressions can form another example; for instance intervals such as $(-N, N)$ in a cyclic group $\mathbf{Z}_M$ can easily be seen to be a Bohr set of rank 1. We can combine these two examples by introducing the concept of a *coset progression*.

Definition 4.21 (Coset progressions) [157] A *coset progression* in an additive group Z is any set of the form $P + H$ where P is a progression and H is a finite subgroup of Z. We say that the coset progression $P + H$ is *proper* if P is proper and $|P + H| = |P||H|$ (i.e. all the sums in $P + H$ are distinct). We say that a coset progression $P + H$ has rank d if the component P has rank d. We say that $P + H$ is symmetric if P has the form $P = (-N, N) \cdot v$.

Of course, Corollary 3.8 shows that every coset progression can also be viewed as an ordinary progression, but possibly of much larger rank. If however Z is a cyclic group of prime order, then H will either be trivial or equal to the whole space, and will thus increase the rank by at most 1. Indeed we can view vector spaces over small finite fields on the one hand, and cyclic groups of prime order on the other, as the two extremes of additive behavior for finite groups Z.

Now we relate Bohr sets of rank d with coset progressions of rank d.

Lemma 4.22 (Bohr sets contain large coset progressions) *[160] Let* $\mathrm{Bohr}(S, \rho)$ *be a Bohr set of rank d in Z with* $0 < \rho < \frac{1}{2}$. *Then there exists a proper symmetric coset progression* $P + H$ *of rank* $0 \leq d' \leq d$*, obeying the inclusions*

$$\mathrm{Bohr}(S, d'^{-2d'}\rho) \subseteq P + H \subseteq \mathrm{Bohr}(S, \rho). \tag{4.28}$$

In particular, from Lemma 4.20 we have

$$\mathbf{P}_Z(P + H) \geq \rho^d d^{-4d^2}. \tag{4.29}$$

Furthermore we have $H = S^{\perp}$.

Proof Let $\phi : Z \to (\mathbf{R}/\mathbf{Z})^S$ be the group homomorphism $\phi(x) := (\xi \cdot x)_{\xi \in S}$. Observe that $\phi(Z)$ is a finite subgroup of the torus $(\mathbf{R}/\mathbf{Z})^S$, and that $\mathrm{Bohr}(S, \rho)$ contains the inverse image of the cube $Q := \{(y_\xi)_{\xi \in S} \in \mathbf{R}^S : |y_\xi| \leq \rho\} \subset \mathbf{R}^S$ (which we identify with its projection in $(\mathbf{R}/\mathbf{Z})^S$) under ϕ.

Let $\Gamma \subseteq \mathbf{R}^S$ be the lattice $\phi(Z) + \mathbf{Z}^S$. Though it is a slight abuse of notation, we consider $\phi(Z) \cap Q$ to be the same as $\Gamma \cap Q$. Applying Lemma 3.36, we can find a progression $\tilde{P} := (-L, L) \cdot w$ for some linearly independent $w_1, \ldots, w_{d'} \subseteq \Gamma$ with $0 \leq d' \leq d$ such that

$$\Gamma \cap d'^{-2d'} \cdot Q \subseteq \tilde{P} \subseteq \Gamma \cap Q.$$

Since the w_j are independent, $\tilde{P}$ is necessarily proper. The claim now follows by setting v_j to be an arbitrary element of $\phi^{-1}(w_j)$ for each $1 \leq j \leq d'$, and setting H equal to the kernel of ϕ, which is of course just $S^{\perp}$. □

In the case of a cyclic group, we can dispense with the group H and sharpen the constants somewhat (though at the cost of losing the first inclusion in (4.28)):

Proposition 4.23 *Let* $Z = \mathbf{Z}_N$ *be a cyclic group, and let* $\mathrm{Bohr}(S, \rho)$ *be a Bohr set of rank d with* $0 < \rho < \frac{1}{2}$. *Then* $\mathrm{Bohr}(S, \rho)$ *contains a symmetric proper progression P of rank d and cardinality*

$$|P| \geq \frac{\rho^d}{d^d} N.$$

Furthermore we may choose P to be symmetric (i.e. $P = -P$*).*

Proof The main tool here will be Minkowski's second theorem. We use the standard bilinear form $\xi \cdot x = \xi x / N$, and write $S = (\xi_1, \ldots, \xi_d)$. Let $\alpha \in \mathbf{R}^d$ be the vector $\alpha := (\frac{\xi_1}{N}, \ldots, \frac{\xi_d}{N})$, and let Γ be the lattice $\mathbf{Z} \cdot \alpha + \mathbf{Z}^d$; this clearly has full rank, and by (3.12)

$$\mathrm{mes}(\mathbf{R}^d/\Gamma) = \mathrm{mes}(\mathbf{R}^d/\mathbf{Z}^d)/|\Gamma/\mathbf{Z}^d| \geq 1/N.$$

Let Q be the cube

$$Q := \{(x_1, \ldots, x_d) \in \mathbf{R}^d : |x_j| < \rho \text{ for all } 1 \leq j \leq n\},$$

and let $0 < \lambda_1 \leq \cdots \leq \lambda_d$ be the succesive minima of Q with respect to Γ, with a corresponding directional basis $v_1, \ldots, v_d \in \Gamma$ as given by Theorem 3.30. In particular we see that every coordinate of v_j has magnitude at most $\lambda_j \rho$.

Let $1 \leq j \leq d$ be arbitrary. Since $v_j \in \Gamma$, we see from the definition of Γ that there exists $w_j \in \mathbf{Z}_N$ such that $v_j \in \alpha w_j + \mathbf{Z}^d$. In particular we see that $\|\xi_i \cdot w_j\|_{\mathbf{R}/\mathbf{Z}} \leq \lambda_j \rho$ for all $1 \leq i, j \leq d$. Set $w := (w_1, \ldots, w_d)$. Now we let $M_j := \lfloor \frac{1}{d\lambda_j} \rfloor$, and let $M := (M_1, \ldots, M_d)$; we now claim that the progression $P := (-M, M) \cdot w$ is proper and lies in $\mathrm{Bohr}(S, \rho)$ (it is clearly symmetric). Let us first verify that $P \subseteq \mathrm{Bohr}(S, \rho)$. If $n = (n_1, \ldots, n_d) \in (-M, M)$, then for any $1 \leq j \leq d$ we have

$$\|\xi_j \cdot (n \cdot w)\|_{\mathbf{R}/\mathbf{Z}} \leq \sum_{j=1}^{d} |n_j| \|\xi_j \cdot w_j\|_{\mathbf{R}/\mathbf{Z}} < \sum_{j=1}^{d} \frac{1}{d\lambda_j} \lambda_j \rho = \rho$$

and hence $n_1 w_1 + \cdots + n_d w_d \in \mathrm{Bohr}(S, \rho)$. This proves the inclusion $P \subseteq \mathrm{Bohr}(S, \rho)$.

Now we show that P is proper. Suppose for contradiction that there exist distinct $n, n' \in (-M, M)$ such that $n \cdot w = n' \cdot w$; setting $\tilde{n} := n - n' \in (-2M, 2M)$, we thus see that $\tilde{n} \cdot w = 0$. In particular, $(\tilde{n} \cdot v)_i$ is an integer for each i. On the other hand, by arguing as before, we see that

$$|(\tilde{n} \cdot v)_i| \leq \sum_{j=1}^{d} |\tilde{n}_j| |\xi_i w_j / N| < \sum_{j=1}^{d} \frac{2}{d\lambda_j} \lambda_j \rho = 2\rho.$$

Since $\rho < 1/2$, we conclude that $(\tilde{n} \cdot v)_i = 0$ for all i, and thus $\sum_j \tilde{n}_j v_j = 0$. But this contradicts the linear independence of the directional basis $v_1, \ldots, v_d$. Thus P is proper.

Finally, the cardinality of the proper progression P is

$$|P| = \prod_{j=1}^{d} (2M_j - 1) \geq \prod_{j=1}^{d} \frac{1}{d\lambda_j}$$

and the claim follows from Minkowski's second theorem. □

One undesirable feature of Bohr sets of large rank d is that they have large doubling constant: (4.26) suggests that $\mathrm{Bohr}(S, \rho) + \mathrm{Bohr}(S, \rho)$ can be 4^d times larger than $\mathrm{Bohr}(S, \rho)$. A useful observation of Bourgain [39] is that if one considers an imbalanced sum $\mathrm{Bohr}(S, \rho) + \mathrm{Bohr}(S, \rho')$, with ρ' much smaller than ρ, then it is still possible for $\mathrm{Bohr}(S, \rho) + \mathrm{Bohr}(S, \rho')$ to be close to $\mathrm{Bohr}(S, \rho)$. This intuition is formalized by the notion of a *regular Bohr set*.

Definition 4.24 (Regular Bohr sets) A Bohr set $\mathrm{Bohr}(S, \rho)$ of rank d is said to be *regular* if one has the estimate

$$\begin{aligned}(1 - 100d|\kappa|)\mathbf{P}_Z(\mathrm{Bohr}(S, \rho)) &\leq \mathbf{P}_Z(\mathrm{Bohr}(S, (1+\kappa)\rho)) \\ &\leq (1 + 100d|\kappa|)\mathbf{P}_Z(\mathrm{Bohr}(S, \rho))\end{aligned}$$

whenever $|\kappa| \leq \frac{1}{100d}$.

Not all Bohr sets are regular. However, it turns out that every Bohr set is "close" to a regular one:

Lemma 4.25 (Regular Bohr sets are ubiquitous) *[39] Let S be a non-empty additive set and let $0 < \varepsilon < 1$. Then there exists $\rho \in [\varepsilon, 2\varepsilon]$ such that $\mathrm{Bohr}(S, \rho)$ is regular.*

Proof Let $f : [0, 1] \to \mathbf{R}$ be the function $f(a) := \frac{1}{d} \log_2 \mathbf{P}_Z(\mathrm{Bohr}(S, 2^a\varepsilon))$. Observe that f is non-decreasing in a, and from Lemma 4.20 we have $f(1) - f(0) \leq \log_2 5$.

Suppose we could find $0.1 \leq a \leq 0.9$ such that $|f(a') - f(a)| \leq 20|a - a'|$ for all $|a| \leq 0.1$. Then it is easy to see that $\mathrm{Bohr}(S, 2^a\varepsilon)$ is regular. Thus, it suffices to obtain an a with this property. This can be done directly from the Hardy–Littlewood maximal inequality (applied to the Lebesgue–Stieltjes measure df), or as follows. If no such a exists, then for every $0.1 \leq a \leq 0.9$ there exists a real interval I of length at most 0.1 and with one endpoint equal to a, such that $\int_I df > \int_I 20\, dx$. These intervals cover $\{a : 0.1 \leq a \leq 0.9\}$, which has measure 0.8. By the Vitali covering lemma (see exercises), one can find thus find a finite subcollection of disjoint intervals $I_1, \dots, I_n$ of total length $|I_1| + \cdots + |I_n| \geq 0.8/5$ (say). But then we have

$$\log_2 5 \geq \int_0^1 df \geq \sum_{i=1}^n \int_{I_i} df \geq \sum_{i=1}^n \int_{I_i} 20dx \geq \frac{0.8}{5} \times 20,$$

a contradiction. □

We shall make a crucial use of this lemma in proving Bourgain's quantitative version of Roth's theorem in Section 10.4.

Exercises

4.4.1 Show that if $0 < \rho \leq 1/6$ and $|S| \geq 1$, then $|\widehat{1_{\mathrm{Bohr}(S,\rho)}}(\xi)| \geq \frac{1}{2}\mathbf{P}_Z(\mathrm{Bohr}(S, \rho))$ for all $\xi \in S$. In particular Bohr sets are extremely non-uniform: $\|\mathrm{Bohr}(S, \rho)\|_u \geq \frac{1}{2}\mathbf{P}_Z(\mathrm{Bohr}(S, \rho))$. By applying Plancherel's theorem, conclude the additional bound $\mathbf{P}_Z(\mathrm{Bohr}(S, \rho)) \leq \frac{4}{|S|}$.

4.4.2 Give examples to show that the density $\mathbf{P}_Z(\mathrm{Bohr}(S, \rho))$ of a Bohr set can be as low as $\Theta(\rho)^{|S|}$, and as large as $\Theta(1/|S|)$, even when ρ is small and $|S|$ is large. Thus the bounds in (4.25) and the preceding exercise cannot be significantly improved.

4.4.3 Establish the bound $\mathbf{P}_Z(\mathrm{Bohr}(S, k\rho)) \leq O(k)^{|S|}\mathbf{P}_Z(\mathrm{Bohr}(S, \rho))$ for any $k \geq 1$. Using the Ruzsa covering lemma (Lemma 2.14), conclude that one can cover $\mathrm{Bohr}(S, k\rho)$ by $O(k)^{|S|}$ translates of $\mathrm{Bohr}(S, \rho)$. In particular, in the notation of Definition 2.25, $\mathrm{Bohr}(S, \rho)$ is a $O(1)^{|S|}$-approximate group.

4.4.4 In the setting of Lemma 4.22, show that $\mathrm{Bohr}(S, \rho)$ can be covered by $O(d)^{d^2}$ translates of $P + H$.

4.4.5 Show that a Bohr set $\mathrm{Bohr}(S, \rho)$ of rank d always contains an arithmetic progression of length $\Theta(|\mathrm{Bohr}(S, \rho)|^{1/d})$ and non-zero step size. (Hint: if $|\mathrm{Bohr}(S, \rho)|^{1/d}$ is large, use the preceding exercise to show that $\mathrm{Bohr}(S, \rho/k)$ contains a non-zero element for some integer $k = \Theta(|\mathrm{Bohr}(S, \rho)|^{1/d})$.)

4.4.6 [160] Let A be an additive set in Z that contains 0. Show that there exists a set S of frequencies with $|S| \leq 1 + \log_2 |A|$ such that $A \cap \mathrm{Bohr}(S, \sqrt{2}) = \{0\}$. (Hint: choose $1 + \lfloor \log_2 |A| \rfloor$ frequencies randomly and independently (allowing for collisions) and use the first moment method.)

4.4.7 (Vitali covering lemma) Let $\mathcal{I}$ be a finite collection of intervals in the real line. Show that there exist a subcollection $I_1, \ldots, I_n$ of these intervals whose interiors are disjoint, and such that $\sum_{i=1}^n |I_i| \geq \frac{1}{5}\mathrm{mes}(\bigcup_{I\in\mathcal{I}} I)$. (Hint: use a greedy algorithm, picking the largest intervals first.) By being more sophisticated in the argument, lower $\frac{1}{5}$ to $\frac{1}{2}$. (Hint: eliminate nested intervals, and then move greedily from left to right to cover $\bigcup_{I\in\mathcal{I}} I$ by two families of interior-disjoint intervals.)

4.4.8 (Hardy–Littlewood maximal inequality) Let μ be a non-negative finite measure on the real line, and let $M\mu$ denote the *Hardy–Littlewood maximal function* $M\mu(x) := \sup_{r>0} \frac{1}{2r}\mu\{y : x - r < y < x + r\}$. (It can be verified that $M\mu$ is a measurable function.) Using the Vitali covering lemma, establish the distributional inequality

$$\mathrm{mes}(\{x : M\mu(x) \geq \lambda\}) \leq \frac{2}{\lambda}\mu(\mathbf{R}).$$