1.1. Equidistribution of polynomial sequences in tori

(Linear) *Fourier analysis* can be viewed as a tool to study an arbitrary function f on (say) the integers $\mathbf{Z}$, by looking at how such a function correlates with *linear phases* such as $n \mapsto e(\xi n)$, where $e(x) := e^{2\pi i x}$ is the fundamental character, and $\xi \in \mathbf{R}$ is a frequency. These correlations control a number of expressions relating to f, such as the expected behaviour of f on arithmetic progressions $n, n+r, n+2r$ of length three.

In this text we will be studying higher-order correlations, such as the correlation of f with quadratic phases such as $n \mapsto e(\xi n^2)$, as these will control the expected behaviour of f on more complex patterns, such as arithmetic progressions $n, n+r, n+2r, n+3r$ of length four. In order to do this, we must first understand the behaviour of *exponential sums* such as

$$\sum_{n=1}^{N} e(\alpha n^2).$$

Such sums are closely related to the *distribution* of expressions such as $\alpha n^2 \bmod 1$ in the unit circle $\mathbf{T} := \mathbf{R}/\mathbf{Z}$, as n varies from 1 to N. More generally, one is interested in the distribution of polynomials $P\colon \mathbf{Z}^d \to \mathbf{T}$ of one or more variables taking values in a torus $\mathbf{T}$; for instance, one might be interested in the distribution of the quadruplet $(\alpha n^2, \alpha(n+r)^2, \alpha(n+2r)^2, \alpha(n+3r)^2)$ as n, r both vary from 1 to N. Roughly speaking, once we understand these types of distributions, then the general machinery of quadratic Fourier analysis will then allow us to understand the distribution of the quadruplet $(f(n), f(n+r), f(n+2r), f(n+3r))$ for more general classes of functions f; this can lead for instance to an understanding of the distribution of arithmetic progressions of length 4 in the primes, if f is somehow related to the primes.

More generally, to find arithmetic progressions such as $n, n+r, n+2r, n+3r$ in a set A, it would suffice to understand the equidistribution of the quadruplet[1] $(1_A(n), 1_A(n+r), 1_A(n+2r), 1_A(n+3r))$ in $\{0,1\}^4$ as n and r vary. This is the starting point for the fundamental connection between *combinatorics* (and more specifically, the task of finding patterns inside sets) and *dynamics* (and more specifically, the theory of equidistribution and recurrence in measure-preserving dynamical systems, which is a subfield of *ergodic theory*). This connection was explored in the previous monograph [**Ta2009**]; it will also be important in this text (particularly as a source of motivation), but the primary focus will be on finitary, and Fourier-based, methods.

[1] Here 1_A is the *indicator function* of A, defined by setting $1_A(n)$ equal to 1 when $n \in A$ and equal to zero otherwise.

The theory of equidistribution of polynomial orbits was developed in the linear case by Dirichlet and Kronecker, and in the polynomial case by Weyl. There are two regimes of interest; the (qualitative) *asymptotic regime* in which the scale parameter N is sent to infinity, and the (quantitative) *single-scale regime* in which N is kept fixed (but large). Traditionally, it is the asymptotic regime which is studied, which connects the subject to other asymptotic fields of mathematics, such as dynamical systems and ergodic theory. However, for many applications (such as the study of the primes), it is the single-scale regime which is of greater importance. The two regimes are not directly equivalent, but are closely related: the single-scale theory can be usually used to derive analogous results in the asymptotic regime, and conversely the arguments in the asymptotic regime can serve as a simplified model to show the way to proceed in the single-scale regime. The analogy between the two can be made tighter by introducing the (qualitative) *ultralimit regime*, which is formally equivalent to the single-scale regime (except for the fact that explicitly quantitative bounds are abandoned in the ultralimit), but resembles the asymptotic regime quite closely.

For the finitary portion of the text, we will be using *asymptotic notation*: $X \ll Y$, $Y \gg X$, or $X = O(Y)$ denotes the bound $|X| \leq CY$ for some absolute constant C, and if we need C to depend on additional parameters, then we will indicate this by subscripts, e.g., $X \ll_d Y$ means that $|X| \leq C_d Y$ for some C_d depending only on d. In the ultralimit theory we will use an analogue of asymptotic notation, which we will review later in this section.

1.1.1. Asymptotic equidistribution theory. Before we look at the single-scale equidistribution theory (both in its finitary form, and its ultralimit form), we will first study the slightly simpler, and much more classical, *asymptotic* equidistribution theory.

Suppose we have a sequence of points $x(1), x(2), x(3), \ldots$ in a compact metric space X. For any finite $N > 0$, we can define the probability measure

$$\mu_N := \mathbf{E}_{n \in [N]} \delta_{x(n)}$$

which is the average of the *Dirac point masses* on each of the points $x(1), \ldots, x(N)$, where we use $\mathbf{E}_{n \in [N]}$ as shorthand for $\frac{1}{N} \sum_{n=1}^{N}$ (with $[N] := \{1, \ldots, N\}$). *Asymptotic equidistribution theory* is concerned with the limiting behaviour of these probability measures μ_N in the limit $N \to \infty$, for various sequences $x(1), x(2), \ldots$ of interest. In particular, we say that the sequence $x \colon \mathbf{N} \to X$ is *asymptotically equidistributed* on $\mathbf{N}$ with respect to a reference *Borel probability measure* μ on X if the μ_N converge in the vague topology to μ or, in other words, that

$$\mathbf{E}_{n \in [N]} f(x(n)) = \int_X f \, d\mu_N \to \int_X f \, d\mu \tag{1.1}$$

for all continuous scalar-valued functions $f \in C(X)$. Note (from the *Riesz representation theorem*) that any sequence is asymptotically equidistributed with respect to at most one Borel probability measure μ.

It is also useful to have a slightly stronger notion of equidistribution: we say that a sequence $x\colon \mathbf{N} \to X$ is *totally asymptotically equidistributed* if it is asymptotically equidistributed on every infinite arithmetic progression, i.e. that the sequence $n \mapsto x(qn+r)$ is asymptotically equidistributed for all integers $q \geq 1$ and $r \geq 0$.

A doubly infinite sequence $(x(n))_{n\in\mathbf{Z}}$, indexed by the integers rather than the natural numbers, is said to be asymptotically equidistributed relative to μ if both halves[2] of the sequence $x(1), x(2), x(3), \ldots$ and $x(-1), x(-2), x(-3), \ldots$ are asymptotically equidistributed relative to μ. Similarly, one can define the notion of a doubly infinite sequence being totally asymptotically equidistributed relative to μ.

Example 1.1.1. If $X = \{0,1\}$, and $x(n) := 1$ whenever $2^{2j} \leq n < 2^{2j+1}$ for some natural number j and $x(n) := 0$ otherwise, show that the sequence x is not asymptotically equidistributed with respect to any measure. Thus we see that asymptotic equidistribution requires all scales to behave "the same" in the limit.

Exercise 1.1.1. If $x\colon \mathbf{N} \to X$ is a sequence into a compact metric space X, and μ is a probability measure on X, show that x is asymptotically equidistributed with respect to μ if and only if one has

$$\lim_{N\to\infty} \frac{1}{N}|\{1 \leq n \leq N : x(n) \in U\}| = \mu(U)$$

for all open sets U in X whose boundary ∂U has measure zero. (*Hint:* For the "only if" part, use *Urysohn's lemma*. For the "if" part, reduce (1.1) to functions f taking values between 0 and 1, and observe that almost all of the level sets $\{y \in X : f(y) < t\}$ have a boundary of measure zero.) What happens if the requirement that ∂U have measure zero is omitted?

Exercise 1.1.2. Let x be a sequence in a compact metric space X which is equidistributed relative to some probability measure μ. Show that for any open set U in X with $\mu(U) > 0$, the set $\{n \in \mathbf{N} : x(n) \in U\}$ is infinite, and furthermore has positive lower density in the sense that

$$\liminf_{N\to\infty} \frac{1}{N}|\{1 \leq n \leq N : x(n) \in U\}| > 0.$$

In particular, if the support of μ is equal to X, show that the set $\{x(n) : n \in \mathbf{N}\}$ is dense in X.

[2]This omits $x(0)$ entirely, but it is easy to see that any individual element of the sequence has no impact on the asymptotic equidistribution.

Exercise 1.1.3. Let $x\colon \mathbf{N} \to X$ be a sequence into a compact metric space X which is equidistributed relative to some probability measure μ. Let $\varphi\colon \mathbf{R} \to \mathbf{R}$ be a compactly supported, piecewise continuous function with only finitely many pieces. Show that for any $f \in C(X)$ one has

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n\in\mathbf{N}} \varphi(n/N) f(x(n)) = \left(\int_X f\ d\mu\right)\left(\int_0^\infty \varphi(t)\ dt\right)$$

and for any open U whose boundary has measure zero, one has

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n\in\mathbf{N}:x(n)\in U} \varphi(n/N) = \mu(U)\left(\int_0^\infty \varphi(t)\ dt\right).$$

In this section, X will be a torus (i.e., a compact connected abelian Lie group), which from the theory of Lie groups is isomorphic to the standard torus $\mathbf{T}^d$, where d is the dimension of the torus. This torus is then equipped with *Haar measure*, which is the unique Borel probability measure on the torus which is translation-invariant. One can identify the standard torus $\mathbf{T}^d$ with the standard fundamental domain $[0,1)^d$, in which case the Haar measure is equated with the usual Lebesgue measure. We shall call a sequence $x_1, x_2, \ldots$ in $\mathbf{T}^d$ (asymptotically) *equidistributed* if it is (asymptotically) equidistributed with respect to Haar measure.

We have a simple criterion for when a sequence is asymptotically equidistributed, that reduces the problem to that of estimating exponential sums:

Proposition 1.1.2 (Weyl equidistribution criterion)**.** *Let $x\colon \mathbf{N} \to \mathbf{T}^d$. Then x is asymptotically equidistributed if and only if*

$$\lim_{N\to\infty} \mathbf{E}_{n\in[N]} e(k\cdot x(n)) = 0 \tag{1.2}$$

for all $k \in \mathbf{Z}^d\backslash\{0\}$, where $e(y) := e^{2\pi i y}$. Here we use the dot product

$$(k_1,\ldots,k_d)\cdot(x_1,\ldots,x_d) := k_1x_1 + \cdots + k_dx_d$$

which maps $\mathbf{Z}^d \times \mathbf{T}^d$ to $\mathbf{T}$.

Proof. The "only if" part is immediate from (1.1). For the "if" part, we see from (1.2) that (1.1) holds whenever f is a plane wave $f(y) := e(k\cdot y)$ for some $k \in \mathbf{Z}^d$ (checking the $k = 0$ case separately), and thus by linearity whenever f is a trigonometric polynomial. But by Fourier analysis (or from the *Stone-Weierstrass theorem*), the trigonometric polynomials are dense in $C(\mathbf{T}^d)$ in the uniform topology. The claim now follows from a standard limiting argument. □

As one consequence of this proposition, one can reduce multidimensional equidistribution to single-dimensional equidistribution:

Corollary 1.1.3. *Let* $x\colon \mathbf{N} \to \mathbf{T}^d$*. Then* x *is asymptotically equidistributed in* $\mathbf{T}^d$ *if and only if, for each* $k \in \mathbf{Z}^d \backslash \{0\}$*, the sequence* $n \mapsto k \cdot x(n)$ *is asymptotically equidistributed in* $\mathbf{T}$*.*

Exercise 1.1.4. Show that a sequence $x : \mathbf{N} \to \mathbf{T}^d$ is totally asymptotically equidistributed if and only if one has

$$\lim_{N\to\infty} \mathbf{E}_{n\in[N]} e(k \cdot x(n))e(\alpha n) = 0 \tag{1.3}$$

for all $k \in \mathbf{Z}^d \backslash \{0\}$ and all rational α.

This quickly gives a test for equidistribution for linear sequences, sometimes known as the *equidistribution theorem*:

Exercise 1.1.5. Let $\alpha, \beta \in \mathbf{T}^d$. By using the geometric series formula, show that the following are equivalent:

(i) The sequence $n \mapsto n\alpha + \beta$ is asymptotically equidistributed on $\mathbf{N}$.

(ii) The sequence $n \mapsto n\alpha + \beta$ is totally asymptotically equidistributed on $\mathbf{N}$.

(iii) The sequence $n \mapsto n\alpha + \beta$ is totally asymptotically equidistributed on $\mathbf{Z}$.

(iv) α is *irrational*, in the sense that $k \cdot \alpha \neq 0$ for any non-zero $k \in \mathbf{Z}^d$.

Remark 1.1.4. One can view Exercise 1.1.5 as an assertion that a linear sequence x_n will equidistribute itself unless there is an "obvious" algebraic obstruction to it doing so, such as $k \cdot x_n$ being constant for some non-zero k. This theme of algebraic obstructions being the "only" obstructions to uniform distribution will be present throughout the text.

Exercise 1.1.5 shows that linear sequences with irrational shift α are equidistributed. At the other extreme, if α is *rational* in the sense that $m\alpha = 0$ for some positive integer m, then the sequence $n \mapsto n\alpha + \beta$ is clearly periodic of period m, and definitely not equidistributed.

In the one-dimensional case $d = 1$, these are the only two possibilities. But in higher dimensions, one can have a mixture of the two extremes, that exhibits irrational behaviour in some directions and periodic behaviour in others. Consider for instance the two-dimensional sequence $n \mapsto (\sqrt{2}n, \frac{1}{2}n) \bmod \mathbf{Z}^2$. The first coordinate is totally asymptotically equidistributed in $\mathbf{T}$, while the second coordinate is periodic; the shift $(\sqrt{2}, \frac{1}{2})$ is neither irrational nor rational, but is a mixture of both. As such, we see that the two-dimensional sequence is equidistributed with respect to Haar measure on the group $\mathbf{T} \times (\frac{1}{2}\mathbf{Z}/\mathbf{Z})$.

This phenomenon generalises:

Proposition 1.1.5 (Equidistribution for abelian linear sequences)**.** *Let T be a torus, and let $x(n) := n\alpha + \beta$ for some $\alpha, \beta \in T$. Then there exists a decomposition $x = x' + x''$, where $x'(n) := n\alpha'$ is totally asymptotically equidistributed on $\mathbf{Z}$ in a subtorus T' of T (with $\alpha' \in T'$, of course), and $x''(n) = n\alpha'' + \beta$ is periodic (or equivalently, that $\alpha'' \in T$ is rational).*

Proof. We induct on the dimension d of the torus T. The claim is vacuous for $d = 0$, so suppose that $d \geq 1$ and that the claim has already been proven for tori of smaller dimension. Without loss of generality we may identify T with $\mathbf{T}^d$.

If α is irrational, then we are done by Exercise 1.1.5, so we may assume that α is not irrational; thus $k \cdot \alpha = 0$ for some non-zero $k \in \mathbf{Z}^d$. We then write $k = mk'$, where m is a positive integer and $k' \in \mathbf{Z}^d$ is *irreducible* (i.e., k' is not a proper multiple of any other element of $\mathbf{Z}^d$); thus $k' \cdot \alpha$ is rational. We may thus write $\alpha = \alpha_1 + \alpha_2$, where α_2 is rational, and $k' \cdot \alpha_1 = 0$. Thus, we can split $x = x_1 + x_2$, where $x_1(n) := n\alpha_1$ and $x_2(n) := n\alpha_2 + \beta$. Clearly x_2 is periodic, while x_1 takes values in the subtorus $T_1 := \{y \in T : k' \cdot y = 0\}$ of T. The claim now follows by applying the induction hypothesis to T_1 (and noting that the sum of two periodic sequences is again periodic). $\square$

As a corollary of the above proposition, we see that any linear sequence $n \mapsto n\alpha + \beta$ in a torus T is equidistributed in some union of finite cosets of a subtorus T'. It is easy to see that this torus T is uniquely determined by α, although there is a slight ambiguity in the decomposition $x = x' + x''$ because one can add or subtract a periodic linear sequence taking values in T from x' and add it to x'' (or vice versa).

Having discussed the linear case, we now consider the more general situation of *polynomial* sequences in tori. To get from the linear case to the polynomial case, the fundamental tool is

Lemma 1.1.6 (van der Corput inequality)**.** *Let $a_1, a_2, \ldots$ be a sequence of complex numbers of magnitude at most 1. Then for every $1 \leq H \leq N$, we have*

$$|\mathbf{E}_{n\in[N]} a_n| \ll \left(\mathbf{E}_{h\in[H]} |\mathbf{E}_{n\in[N]} a_{n+h} \overline{a_n}|\right)^{1/2} + \frac{1}{H^{1/2}} + \frac{H^{1/2}}{N^{1/2}}.$$

Proof. For each $h \in [H]$, we have

$$\mathbf{E}_{n\in[N]} a_n = \mathbf{E}_{n\in[N]} a_{n+h} + O\left(\frac{H}{N}\right)$$

and hence, on averaging,

$$\mathbf{E}_{n\in[N]} a_n = \mathbf{E}_{n\in[N]} \mathbf{E}_{h\in[H]} a_{n+h} + O\left(\frac{H}{N}\right).$$

Applying Cauchy-Schwarz, we conclude

$$\mathbf{E}_{n\in[N]}a_n \ll (\mathbf{E}_{n\in[N]}|\mathbf{E}_{h\in[H]}a_{n+h}|^2)^{1/2} + \frac{H}{N}.$$

We expand out the left-hand side as

$$\mathbf{E}_{n\in[N]}a_n \ll (\mathbf{E}_{h,h'\in[H]}\mathbf{E}_{n\in[N]}a_{n+h}\overline{a_{n+h'}})^{1/2} + \frac{H}{N}.$$

The diagonal contribution $h = h'$ is $O(1/H)$. By symmetry, the off-diagonal contribution can be dominated by the contribution when $h > h'$. Making the change of variables $n \mapsto n - h'$, $h \mapsto h + h'$ (accepting a further error of $O(H^{1/2}/N^{1/2})$), we obtain the claim. □

Corollary 1.1.7 (van der Corput lemma)**.** *Let $x\colon \mathbf{N} \to \mathbf{T}^d$ be such that the derivative sequence $\partial_h x\colon n \mapsto x(n+h)-x(n)$ is asymptotically equidistributed on $\mathbf{N}$ for all positive integers h. Then x_n is asymptotically equidistributed on $\mathbf{N}$. Similarly with $\mathbf{N}$ replaced by $\mathbf{Z}$.*

Proof. We just prove the claim for $\mathbf{N}$, as the claim for $\mathbf{Z}$ is analogous (and can in any case be deduced from the $\mathbf{N}$ case).

By Proposition 1.1.2, we need to show that for each non-zero $k \in \mathbf{Z}^d$, the exponential sum

$$|\mathbf{E}_{n\in[N]}e(k \cdot x(n))|$$

goes to zero as $N \to \infty$. Fix an $H > 0$. By Lemma 1.1.6, this expression is bounded by

$$\ll (\mathbf{E}_{h\in[H]}|\mathbf{E}_{n\in[N]}e(k\cdot(x(n+h)-x(n)))|)^{1/2} + \frac{1}{H^{1/2}} + \frac{H^{1/2}}{N^{1/2}}.$$

On the other hand, for each fixed positive integer h, we have from hypothesis and Proposition 1.1.2 that $|\mathbf{E}_{n\in[N]}e(k\cdot(x(n+h)-x(n)))|$ goes to zero as $N \to \infty$. Taking limit superior as $N \to \infty$, we conclude that

$$\limsup_{N\to\infty} |\mathbf{E}_{n\in[N]}e(k \cdot x(n))| \ll \frac{1}{H^{1/2}}.$$

Since H is arbitrary, the claim follows. □

Remark 1.1.8. There is another famous lemma by van der Corput concerning oscillatory integrals, but it is not directly related to the material discussed here.

Corollary 1.1.7 has the following immediate corollary:

Corollary 1.1.9 (Weyl equidistribution theorem for polynomials)**.** *Let $s \geq 1$ be an integer, and let $P(n) = \alpha_s n^s + \cdots + \alpha_0$ be a polynomial of degree s with $\alpha_0, \ldots, \alpha_s \in \mathbf{T}^d$. If α_s is irrational, then $n \mapsto P(n)$ is asymptotically equidistributed on $\mathbf{Z}$.*

Proof. We induct on s. For $s = 1$ this follows from Exercise 1.1.5. Now suppose that $s > 1$, and that the claim has already been proven for smaller values of s. For any positive integer h, we observe that $P(n + h) - P(n)$ is a polynomial of degree $s - 1$ in n with leading coefficient $sh\alpha_s n^{s-1}$. As α_s is irrational, $sh\alpha_s$ is irrational also, and so by the induction hypothesis, $P(n + h) - P(n)$ is asymptotically equidistributed. The claim now follows from Corollary 1.1.7. □

Exercise 1.1.6. Let $P(n) = \alpha_s n^s + \cdots + \alpha_0$ be a polynomial of degree s in $\mathbf{T}^d$. Show that the following are equivalent:

(i) P is asymptotically equidistributed on $\mathbf{N}$.

(ii) P is totally asymptotically equidistributed on $\mathbf{N}$.

(iii) P is totally asymptotically equidistributed on $\mathbf{Z}$.

(iv) There does not exist a non-zero $k \in \mathbf{Z}^d$ such that $k \cdot \alpha_1 = \cdots = k \cdot \alpha_s = 0$.

(*Hint:* It is convenient to first use Corollary 1.1.3 to reduce to the one-dimensional case.)

This gives a polynomial variant of the equidistribution theorem:

Exercise 1.1.7 (Equidistribution theorem for abelian polynomial sequences)**.** Let T be a torus, and let P be a polynomial map from $\mathbf{Z}$ to T of some degree $s \geq 0$. Show that there exists a decomposition $P = P' + P''$, where P', P'' are polynomials of degree s, P' is totally asymptotically equidistributed in a subtorus T' of T on $\mathbf{Z}$, and P'' is periodic (or equivalently, that all non-constant coefficients of P'' are rational).

In particular, we see that polynomial sequences in a torus are equidistributed with respect to a finite combination of Haar measures of cosets of a subtorus. Note that this finite combination can have multiplicity; for instance, when considering the polynomial map $n \mapsto (\sqrt{2}n, \frac{1}{3}n^2) \bmod \mathbf{Z}^2$, it is not hard to see that this map is equidistributed with respect to 1/3 times the Haar probability measure on $(\mathbf{T}) \times \{0 \bmod \mathbf{Z}\}$, plus 2/3 times the Haar probability measure on $(\mathbf{T}) \times \{\frac{1}{3} \bmod \mathbf{Z}\}$.

Exercise 1.1.7 gives a satisfactory description of the asymptotic equidistribution of arbitrary polynomial sequences in tori. We give just one example of how such a description can be useful:

Exercise 1.1.8 (Recurrence)**.** Let T be a torus, let P be a polynomial map from $\mathbf{Z}$ to T, and let n_0 be an integer. Show that there exists a sequence n_j of positive integers going to infinity such that $P(n_j) \to P(n_0)$.

We discussed recurrence for one-dimensional sequences $x\colon n \mapsto x(n)$. It is also of interest to establish an analogous theory for multi-dimensional sequences, as follows.

Definition 1.1.10. A multidimensional sequence $x\colon \mathbf{Z}^m \to X$ is *asymptotically equidistributed* relative to a probability measure μ if, for every continuous, compactly supported function $\varphi\colon \mathbf{R}^m \to \mathbf{R}$ and every function $f \in C(X)$, one has

$$\frac{1}{N^m} \sum_{n \in \mathbf{Z}^m} \varphi(n/N) f(x(n)) \to \left(\int_{\mathbf{R}^m} \varphi\right)\left(\int_X f\ d\mu\right)$$

as $N \to \infty$. The sequence is *totally asymptotically equidistributed* relative to μ if the sequence $n \mapsto x(qn+r)$ is asymptotically equidistributed relative to μ for all positive integers q and all $r \in \mathbf{Z}^m$.

Exercise 1.1.9. Show that this definition of equidistribution on $\mathbf{Z}^m$ coincides with the preceding definition of equidistribution on $\mathbf{Z}$ in the one-dimensional case $m = 1$.

Exercise 1.1.10 (Multidimensional Weyl equidistribution criterion)**.** Let $x\colon \mathbf{Z}^m \to \mathbf{T}^d$ be a multidimensional sequence. Show that x is asymptotically equidistributed if and only if

$$\lim_{N\to\infty} \frac{1}{N^m} \sum_{n \in \mathbf{Z}^m : n/N \in B} e(k \cdot x(n)) = 0 \tag{1.4}$$

for all $k \in \mathbf{Z}^d \backslash \{0\}$ and all rectangular boxes B in $\mathbf{R}^m$. Then show that x is totally asymptotically equidistributed if and only if

$$\lim_{N\to\infty} \frac{1}{N^m} \sum_{n \in \mathbf{Z}^m : n/N \in B} e(k \cdot x(n)) e(\alpha \cdot n) = 0 \tag{1.5}$$

for all $k \in \mathbf{Z}^d \backslash \{0\}$, all rectangular boxes B in $\mathbf{R}^m$, and all rational $\alpha \in \mathbf{Q}^m$.

Exercise 1.1.11. Let $\alpha_1, \ldots, \alpha_m, \beta \in \mathbf{T}^d$, and let $x\colon \mathbf{Z}^m \to \mathbf{T}^d$ be the linear sequence $x(n_1, \ldots, n_m) := n_1\alpha_1 + \cdots + n_m\alpha_m + \beta$. Show that the following are equivalent:

(i) The sequence x is asymptotically equidistributed on $\mathbf{Z}^m$.

(ii) The sequence x is totally asymptotically equidistributed on $\mathbf{Z}^m$.

(iii) We have $(k \cdot \alpha_1, \ldots, k \cdot \alpha_m) \neq 0$ for any non-zero $k \in \mathbf{Z}^d$.

Exercise 1.1.12 (Multidimensional van der Corput lemma)**.** Let $x\colon \mathbf{Z}^m \to \mathbf{T}^d$ be such that the sequence $\partial_h x\colon n \mapsto x(n+h) - x(n)$ is asymptotically equidistributed on $\mathbf{Z}^m$ for all h outside of a hyperplane in $\mathbf{R}^m$. Show that x is asymptotically equidistributed on $\mathbf{Z}^m$.

Exercise 1.1.13. Let

$$P(n_1, \dots, n_m) := \sum_{i_1,\dots,i_m \geq 0: i_1+\dots+i_m \leq s} \alpha_{i_1,\dots,i_m} n_1^{i_1} \dots n_m^{i_m}$$

be a polynomial map from $\mathbf{Z}^m$ to $\mathbf{T}^d$ of degree s, where $\alpha_{i_1,\dots,i_m} \in \mathbf{T}^d$ are coefficients. Show that the following are equivalent:

(i) P is asymptotically equidistributed on $\mathbf{Z}^m$.

(ii) P is totally asymptotically equidistributed on $\mathbf{Z}^m$.

(iii) There does not exist a non-zero $k \in \mathbf{Z}^d$ such that $k \cdot \alpha_{i_1,\dots,i_m} = 0$ for all $(i_1, \dots, i_m) \neq 0$.

Exercise 1.1.14 (Equidistribution for abelian multidimensional polynomial sequences)**.** Let T be a torus, and let P be a polynomial map from $\mathbf{Z}^m$ to T of some degree $s \geq 0$. Show that there exists a decomposition $P = P' + P''$, where P', P'' are polynomials of degree s, P' is totally asymptotically equidistributed in a subtorus T' of T on $\mathbf{Z}^m$, and P'' is periodic with respect to some finite index sublattice of $\mathbf{Z}^m$ (or equivalently, that all non-constant coefficients of P'' are rational).

We give just one application of this multidimensional theory, that gives a hint as to why the theory of equidistribution of polynomials may be relevant:

Exercise 1.1.15. Let T be a torus, let P be a polynomial map from $\mathbf{Z}$ to T, let $\varepsilon > 0$, and let $k \geq 1$. Show that there exists positive integers $a, r \geq 1$ such that $P(a), P(a+r), \dots, P(a+(k-1)r)$ all lie within ε of each other. (*Hint:* Consider the polynomial map from $\mathbf{Z}^2$ to T^k that maps (a, r) to $(P(a), \dots, P(a+(k-1)r))$. One can also use the one-dimensional theory by freezing a and only looking at the equidistribution in r.)

1.1.2. Single-scale equidistribution theory. We now turn from the asymptotic equidistribution theory to the equidistribution theory at a single scale N. Thus, instead of analysing the qualitative distribution of infinite sequence $x\colon \mathbf{N} \to X$, we consider instead the quantitative distribution of a finite sequence $x\colon [N] \to X$, where N is a (large) natural number and $[N] := \{1, \dots, N\}$. To make everything quantitative, we will replace the notion of a continuous function by that of a *Lipschitz function.* Recall that the (inhomogeneous) Lipschitz norm $\|f\|_{\mathrm{Lip}}$ of a function $f\colon X \to \mathbf{R}$ on a metric space $X = (X, d)$ is defined by the formula

$$\|f\|_{\mathrm{Lip}} := \sup_{x \in X} |f(x)| + \sup_{x,y \in X: x \neq y} \frac{|f(x) - f(y)|}{d(x, y)}.$$

We also define the homogeneous Lipschitz semi-norm

$$\|f\|_{\mathrm{Lip}} := \sup_{x,y \in X: x \neq y} \frac{|f(x) - f(y)|}{d(x, y)}.$$

Definition 1.1.11. Let $X = (X, d)$ be a compact metric space, let $\delta > 0$, let μ be a probability measure on X. A finite sequence $x\colon [N] \to X$ is said to be *δ-equidistributed* relative to μ if one has

$$|\mathbf{E}_{n\in[N]} f(x(n)) - \int_X f\ d\mu| \le \delta \|f\|_{\mathrm{Lip}} \tag{1.6}$$

for all Lipschitz functions $f\colon X \to \mathbf{R}$.

We say that the sequence $x_1, \ldots, x_N \in X$ is *totally δ-equidistributed* relative to μ if one has

$$|\mathbf{E}_{n\in P} f(x(n)) - \int_X f\ d\mu| \le \delta \|f\|_{\mathrm{Lip}}$$

for all Lipschitz functions $f\colon X \to \mathbf{R}$ and all arithmetic progressions P in $[N]$ of length at least δN.

In this section, we will only apply this concept to the torus $\mathbf{T}^d$ with the Haar measure μ and the metric inherited from the Euclidean metric. However, in subsequent sections we will also consider equidistribution in other spaces, most notably on *nilmanifolds*.

Exercise 1.1.16. Let $x(1), x(2), x(3), \ldots$ be a sequence in a metric space $X = (X, d)$, and let μ be a probability measure on X. Show that the sequence $x(1), x(2), \ldots$ is asymptotically equidistributed relative to μ if and only if, for every $\delta > 0$, $x(1), \ldots, x(N)$ is δ-equidistributed relative to μ whenever N is sufficiently large depending on δ, or equivalently if $x(1), \ldots, x(N)$ is $\delta(N)$-equidistributed relative to μ for all $N > 0$, where $\delta(N) \to 0$ as $N \to \infty$. (*Hint:* You will need the *Arzelá-Ascoli theorem.*)

Similarly, show that $x(1), x(2), \ldots$ is totally asymptotically equidistributed relative to μ if and only if, for every $\delta > 0$, $x(1), \ldots, x(N)$ is totally δ-equidistributed relative to μ whenever N is sufficiently large depending on δ, or equivalently if $x(1), \ldots, x(N)$ is totally $\delta(N)$-equidistributed relative to μ for all $N > 0$, where $\delta(N) \to 0$ as $N \to \infty$.

Remark 1.1.12. More succinctly, (total) asymptotic equidistribution of $x(1), x(2), \ldots$ is equivalent to (total) $o_{N\to\infty}(1)$-equidistribution of $x(1), \ldots,$ $x(N)$ as $N \to \infty$, where $o_{n\to\infty}(1)$ denotes a quantity that goes to zero as $N \to \infty$. Thus we see that asymptotic notation such as $o_{n\to\infty}(1)$ can efficiently conceal a surprisingly large number of quantifiers.

Exercise 1.1.17. Let N_0 be a large integer, and let $x(n) := n/N_0 \bmod 1$ be a sequence in the standard torus $\mathbf{T} = \mathbf{R}/\mathbf{Z}$ with Haar measure. Show that whenever N is a positive multiple of N_0, then the sequence $x(1), \ldots, x(N)$ is $O(1/N_0)$-equidistributed. What happens if N is not a multiple of N_0?

If, furthermore, $N \ge N_0^2$, show that $x(1), \ldots, x(N)$ is $O(1/\sqrt{N_0})$-equidistributed. Why is a condition such as $N \ge N_0^2$ necessary?

Note that the above exercise does not specify the exact relationship between δ and N when one is given an asymptotically equidistributed sequence $x(1), x(2), \ldots$; this relationship is the additional piece of information provided by single-scale equidistribution that is not present in asymptotic equidistribution.

It turns out that much of the asymptotic equidistribution theory has a counterpart for single-scale equidistribution. We begin with the Weyl criterion.

Proposition 1.1.13 (Single-scale Weyl equidistribution criterion)**.** *Let* $x_1, x_2, \ldots, x_N$ *be a sequence in* $\mathbf{T}^d$*, and let* $0 < \delta < 1$.

(i) *If* $x_1, \ldots, x_N$ *is* δ*-equidistributed, and* $k \in \mathbf{Z}^d \backslash \{0\}$ *has magnitude* $|k| \leq \delta^{-c}$*, then one has*

$$|\mathbf{E}_{n\in[N]} e(k \cdot x_n)| \ll_d \delta^c$$

if $c > 0$ *is a small enough absolute constant.*

(ii) *Conversely, if* $x_1, \ldots, x_N$ *is* not δ*-equidistributed, then there exists* $k \in \mathbf{Z}^d \backslash \{0\}$ *with magnitude* $|k| \ll_d \delta^{-C_d}$*, such that*

$$|\mathbf{E}_{n\in[N]} e(k \cdot x_n)| \gg_d \delta^{C_d}$$

for some C_d *depending on* d*.*

Proof. The first claim is immediate as the function $x \mapsto e(k \cdot x)$ has mean zero and Lipschitz constant $O_d(|k|)$, so we turn to the second claim. By hypothesis, (1.6) fails for some Lipschitz f. We may subtract off the mean and assume that $\int_{\mathbf{T}^d} f = 0$; we can then normalise the Lipschitz norm to be one; thus we now have

$$|\mathbf{E}_{n\in[N]} f(x_n)| > \delta.$$

We introduce a summation parameter $R \in \mathbf{N}$, and consider the *Fejér partial Fourier series*

$$F_R f(x) := \sum_{k\in\mathbf{Z}^d} m_R(k) \hat{f}(k) e(k \cdot x)$$

where $\hat{f}(k)$ are the Fourier coefficients

$$\hat{f}(k) := \int_{\mathbf{T}^d} f(x) e(-k \cdot x)\, dx$$

and m_R is the Fourier multiplier

$$m_R(k_1, \ldots, k_d) := \prod_{j=1}^{d} \left(1 - \frac{|k_j|}{R}\right)_+ .$$

Standard Fourier analysis shows that we have the convolution representation

$$F_R f(x) = \int_{\mathbf{T}^d} f(y) K_R(x-y)$$

where K_R is the Fejér kernel

$$K_R(x_1, \ldots, x_d) := \prod_{j=1}^d \frac{1}{R} \left(\frac{\sin(\pi R x_j)}{\sin(\pi x_j)} \right)^2 .$$

Using the kernel bounds

$$\int_{\mathbf{T}^d} K_R = 1$$

and

$$|K_R(x)| \ll_d \prod_{j=1}^d R(1 + R\|x_j\|_{\mathbf{T}})^{-2},$$

where $\|x\|_{\mathbf{T}}$ is the distance from x to the nearest integer, and the Lipschitz nature of f, we see that

$$F_R f(x) = f(x) + O_d(1/R).$$

Thus, if we choose R to be a sufficiently small multiple of $1/\delta$ (depending on d), one has

$$|\mathbf{E}_{n \in [N]} F_R f(x_n)| \gg \delta$$

and thus by the pigeonhole principle (and the trivial bound $\hat{f}(k) = O(1)$ and $\hat{f}(0) = 0$) we have

$$|\mathbf{E}_{n \in [N]} e(k \cdot x_n)| \gg_d \delta^{O_d(1)}$$

for some non-zero k of magnitude $|k| \ll_d \delta^{-O_d(1)}$, and the claim follows. $\square$

There is an analogue for total equidistribution:

Exercise 1.1.18. Let $x_1, x_2, \ldots, x_N$ be a sequence in $\mathbf{T}^d$, and let $0 < \delta < 1$.

(i) If $x_1, \ldots, x_N$ is totally δ-equidistributed, $k \in \mathbf{Z}^d \backslash \{0\}$ has magnitude $|k| \leq \delta^{-c_d}$, and a is a rational of height at most δ^{-c_d}, then one has

$$|\mathbf{E}_{n \in [N]} e(k \cdot x_n) e(an)| \ll_d \delta^{c_d}$$

if $c_d > 0$ is a small enough constant depending only on d.

(ii) Conversely, if $x_1, \ldots, x_N$ is *not* totally δ-equidistributed, then there exists $k \in \mathbf{Z}^d \backslash \{0\}$ with magnitude $|k| \ll_d \delta^{-C_d}$, and a rational a of height $O_d(\delta^{-C_d})$, such that

$$|\mathbf{E}_{n \in [N]} e(k \cdot x_n) e(an)| \gg_d \delta^{C_d}$$

for some C_d depending on d.

This gives a version of Exercise 1.1.5:

Exercise 1.1.19. Let $\alpha, \beta \in \mathbf{T}^d$, let $N \geq 1$, and let $0 < \delta < 1$. Suppose that the linear sequence $(\alpha n + \beta)_{n=1}^N$ is not totally δ-equidistributed. Show that there exists a non-zero $k \in \mathbf{Z}^d$ with $|k| \ll_d \delta^{-O_d(1)}$ such that $\|k \cdot \alpha\|_{\mathbf{T}} \ll_d \delta^{-O_d(1)}/N$.

Next, we give an analogue of Corollary 1.1.7:

Exercise 1.1.20 (Single-scale van der Corput lemma)**.** Let $x_1, x_2, \ldots, x_N \in \mathbf{T}^d$ be a sequence which is not totally δ-equidistributed for some $0 < \delta \leq 1/2$. Let $1 \leq H \leq \delta^{-C_d} N$ for some sufficiently large C_d depending only on d. Then there exists at least $\delta^{C_d} H$ integers $h \in [-H, H]$ such that the sequence $(x_{n+h} - x_n)_{n=1}^N$ is not totally δ^{C_d}-equidistributed (where we extend x_n by zero outside of $\{1, \ldots, N\}$). (*Hint:* Apply Lemma 1.1.6.)

Just as in the asymptotic setting, we can use the van der Corput lemma to extend the linear equidistribution theory to polynomial sequences. To get satisfactory results, though, we will need an additional input, namely the following classical lemma, essentially due to Vinogradov:

Lemma 1.1.14. *Let $\alpha \in \mathbf{T}$, $0 < \varepsilon < 1/100$, $100\varepsilon < \delta < 1$, and $N \geq 100/\delta$. Suppose that $\|n\alpha\|_{\mathbf{T}} \leq \varepsilon$ for at least δN values of $n \in [-N, N]$. Then there exists a positive integer $q = O(1/\delta)$ such that $\|\alpha q\|_{\mathbf{T}} \ll \frac{\varepsilon q}{\delta N}$.*

The key point here is that one starts with many multiples of α being somewhat close ($O(\varepsilon)$) to an integer, but concludes that there is a single multiple of α which is *very* close ($O(\varepsilon/N)$, ignoring factors of δ) to an integer.

Proof. By the pigeonhole principle, we can find two distinct integers $n, n' \in [-N, N]$ with $|n - n'| \ll 1/\delta$ such that $\|n\alpha\|_{\mathbf{T}}, \|n'\alpha\|_{\mathbf{T}} \leq \varepsilon$. Setting $q := |n' - n|$, we thus have $\|q\alpha\|_{\mathbf{T}} \leq 2\varepsilon$. We may assume that $q\alpha \neq 0$ since we are done otherwise. Since $N \geq 100/\delta$, we have $N/q \geq 10$ (say).

Now partition $[-N, N]$ into q arithmetic progressions $\{nq + r : -N/q + O(1) \leq n \leq N/q + O(1)\}$ for some $r = 0, \ldots, q-1$. By the pigeonhole principle, there must exist an r for which the set

$$\{-N/q + O(1) \leq n \leq N/q + O(1) : \|\alpha(nq + r)\|_{\mathbf{T}} \leq \varepsilon\}$$

has cardinality at least $\delta N/q$. On the other hand, since $\|q\alpha\|_{\mathbf{T}} \leq 2\varepsilon \leq 0.02$, we see that this set consists of intervals of length at most $2\varepsilon/\|q\alpha\|_{\mathbf{T}}$, punctuated by gaps of length at least $0.9/\|q\alpha\|_{\mathbf{T}}$ (say). Since the gaps are at least $0.45/\varepsilon$ times as large as the intervals, we see that if two or more of these intervals appear in the set, then the cardinality of the set is at most $100\varepsilon N/q < \delta N/q$, a contradiction. Thus at most one interval appears in the set, which implies that $2\varepsilon/\|q\alpha\|_{\mathbf{T}} \geq \delta N/q$, and the claim follows. $\square$

Remark 1.1.15. The numerical constants can of course be improved, but this is not our focus here.

Exercise 1.1.21. Let $P\colon \mathbf{Z} \to \mathbf{T}^d$ be a polynomial sequence $P(n) := \alpha_s n^s + \cdots + \alpha_0$, let $N \geq 1$, and let $0 < \delta < 1$. Suppose that the polynomial sequence P is not totally δ-equidistributed on $[N]$. Show that there exists a non-zero $k \in \mathbf{Z}^d$ with $|k| \ll_{d,s} \delta^{-O_{d,s}(1)}$ such that $\|k \cdot \alpha_s\|_{\mathbf{T}} \ll_{d,s} \delta^{-O_{d,s}(1)}/N^s$. (*Hint:* Induct on s starting with Exercise 1.1.19 for the base case, and then using Exercise 1.1.20 and Lemma 1.1.14 to continue the induction.)

Note the N^s denominator; the higher-degree coefficients of a polynomial need to be *very* rational in order not to cause equidistribution.

The above exercise only controls the top degree coefficient, but we can in fact control all coefficients this way:

Lemma 1.1.16. *With the hypotheses of Exercise* 1.1.21*, we can in fact find a non-zero* $k \in \mathbf{Z}^d$ *with* $|k| \ll_{d,s} \delta^{-O_{d,s}(1)}$ *such that* $\|k \cdot \alpha_i\|_{\mathbf{T}} \ll_{d,s} \delta^{-O_{d,s}(1)}/N^i$ *for all* $i = 0, \ldots, s$.

Proof. We shall just establish the one-dimensional case $d = 1$, as the general dimensional case then follows from Exercise 1.1.18.

The case $s \leq 1$ follows from Exercise 1.1.19, so assume inductively that $s > 1$ and that the claim has already been proven for smaller values of s. We allow all implied constants to depend on s. From Exercise 1.1.21, we already can find a positive k with $k = O(\delta^{-O(1)})$ such that $\|k\alpha_s\|_{\mathbf{T}} \ll \delta^{-O(1)}/N^s$. We now partition $[N]$ into arithmetic progressions of spacing k and length $N' \sim \delta^C N$ for some sufficiently large C; then by the pigeonhole principle, we see that P fails to be totally $\gg \delta^{O(1)}$-equidistributed on one of these progressions. But on one such progression (which can be identified with $[N']$) the degree s component of P is essentially constant (up to errors much smaller than δ) if C is large enough; if one then applies the induction hypothesis to the remaining portion of P on this progression, we can obtain the claim. □

This gives us the following analogue of Exercise 1.1.7. We say that a subtorus T of some dimension d' of a standard torus $\mathbf{T}^d$ has *complexity* at most M if there exists an invertible linear transformation $L \in SL_d(\mathbf{Z})$ with integer coefficients (which can thus be viewed as a homeomorphism of $\mathbf{T}^d$ that maps T to the standard torus $\mathbf{T}^{d'} \times \{0\}^{d-d'}$), and such that all coefficients have magnitude at most M.

Exercise 1.1.22. Show that every subtorus (i.e., compact connected Lie subgroup) T of $\mathbf{T}^d$ has finite complexity. (*Hint:* Let V be the Lie algebra of T, then identify V with a subspace of $\mathbf{R}^d$ and T with $V/(V \cap \mathbf{Z}^d)$. Show that $V \cap \mathbf{Z}^d$ is a full rank sublattice of V, and is thus generated by $\dim(V)$ independent generators.)

Proposition 1.1.17 (Single-scale equidistribution theorem for abelian polynomial sequences). *Let P be a polynomial map from $\mathbf{Z}$ to $\mathbf{T}^d$ of some degree $s \geq 0$, and let $F\colon \mathbf{R}^+ \to \mathbf{R}^+$ be an increasing function. Then there exists an integer $1 \leq M \leq O_{F,s,d}(1)$ and a decomposition*

$$P = P_{\mathrm{smth}} + P_{\mathrm{equi}} + P_{\mathrm{rat}}$$

into polynomials of degree s, where

(i) *(P_{smth} is smooth) The i^{th} coefficient $\alpha_{i,\mathrm{smth}}$ of P_{smth} has size $O(M/N^i)$. In particular, on the interval $[N]$, P_{smth} is Lipschitz with homogeneous norm $O_{s,d}(M/N)$.*

(ii) *(P_{equi} is equidistributed) There exists a subtorus T of $\mathbf{T}^d$ of complexity at most M and some dimension d', such that P_{equi} takes values in T and is totally $1/F(M)$-equidistributed on $[N]$ in this torus (after identifying this torus with $\mathbf{T}^{d'}$ using an invertible linear transformation of complexity at most M).*

(iii) *(P_{rat} is rational) The coefficients $\alpha_{i,\mathrm{rat}}$ of P_{rat} are such that $q\alpha_{i,\mathrm{rat}} = 0$ for some $1 \leq q \leq M$ and all $0 \leq i \leq s$. In particular, $qP_{\mathrm{rat}} = 0$ and P_{rat} is periodic with period q.*

If, furthermore, F is of polynomial growth, and more precisely $F(M) \leq KM^A$ for some $A, K \geq 1$, then one can take $M \ll_{A,s,d} K^{O_{A,s,d}(1)}$.

Example 1.1.18. Consider the linear flow $P(n) := (\sqrt{2}n, (\frac{1}{2}+\frac{1}{N})n) \bmod \mathbf{Z}^2$ in $\mathbf{T}^2$ on $[N]$. This flow can be decomposed into a smooth flow $P_{\mathrm{smth}}(n) := (0, \frac{1}{N}n) \bmod \mathbf{Z}^2$ with a homogeneous Lipschitz norm of $O(1/N)$, an equidistributed flow $P_{\mathrm{equi}}(n) := (\sqrt{2}n, 0) \bmod \mathbf{Z}^2$ which will be δ-equidistributed on the subtorus $\mathbf{T}^1 \times \{0\}$ for a reasonably small δ (in fact one can take δ as small as N^{-c} for some small absolute constant $c > 0$), and a rational flow $P_{\mathrm{rat}}(n) := (0, \frac{1}{2}n) \bmod \mathbf{Z}^2$, which is periodic with period 2. This example illustrates how all three components of this decomposition arise naturally in the single-scale case.

Remark 1.1.19. Comparing this result with the asymptotically equidistributed analogue in Example 1.1.7, we notice several differences. Firstly, we now have the smooth component P_{smth}, which did not previously make an appearance (except implicitly, as the constant term in P'). Secondly, the equidistribution of the component P_{equi} is not infinite, but is the next best thing, namely it is given by an arbitrary function F of the quantity M, which controls the other components of the decomposition.

Proof. The case $s = 0$ is trivial, so suppose inductively that $s \geq 1$, and that the claim has already been proven for lower degrees. Then for fixed degree, the case $d = 0$ is vacuously true, so we make a further inductive assumption $d \geq 1$ and the claim has already been proven for smaller dimensions (keeping s fixed).

If P is already totally $1/F(1)$-equidistributed then we are done (setting $P_{\mathrm{equi}} = P$ and $P_{\mathrm{smth}} = P_{\mathrm{rat}} = 0$ and $M = 1$), so suppose that this is not the case. Applying Exercise 1.1.21, we conclude that there is some non-zero $k \in \mathbf{Z}^d$ with $|k| \ll_{d,s} F(1)^{O_{d,s}(1)}$ such that

$$\|k \cdot \alpha_i\|_{\mathbf{T}} \ll_{d,s} F(1)^{O_{d,s}(1)}/N^i$$

for all $i = 0, \ldots, s$. We split $k = mk'$ where k' is irreducible and m is a positive integer. We can therefore split $\alpha_i = \alpha_{i,\mathrm{smth}} + \alpha_{i,\mathrm{rat}} + \alpha'_i$ where $\alpha_{i,\mathrm{smth}} = O(F(1)^{O_{d,s}(1)}/N^i)$, $q\alpha_i = 0$ for some positive integer $q = O_{d,s}(F(1)^{O_{d,s}(1)})$, and $k' \cdot \alpha'_i = 0$. This then gives a decomposition $P = P_{\mathrm{smth}} + P' + P_{\mathrm{rat}}$, with P' taking values in the subtorus $\{x \in \mathbf{T}^d : k' \cdot x = 0\}$, which can be identified with $\mathbf{T}^{d-1}$ after an invertible linear transformation with integer coefficients of size $O_{d,s}(F(1)^{O_{d,s}(1)})$. If one applies the induction hypothesis to P' (with F replaced by a suitably larger function F') one then obtains the claim.

The final claim about polynomial bounds can be verified by a closer inspection of the argument (noting that all intermediate steps are polynomially quantitative, and that the length of the induction is bounded by $O_{d,s}(1)$). □

Remark 1.1.20. It is instructive to see how this smooth-equidistributed-rational decomposition evolves as N increases. Roughly speaking, the torus T that the P_{equi} component is equidistributed on is stable at most scales, but there will be a finite number of times in which a "growth spurt" occurs and T jumps up in dimension. For instance, consider the linear flow $P(n) := (n/N_0, n/N_0^2) \bmod \mathbf{Z}^2$ on the two-dimensional torus. At scales $N \ll N_0$ (and with F fixed, and N_0 assumed to be sufficiently large depending on F), P consists entirely of the smooth component. But as N increases past N_0, the first component of P no longer qualifies as smooth, and becomes equidistributed instead; thus in the range $N_0 \ll N \ll N_0^2$, we have $P_{\mathrm{smth}}(n) = (0, n/N_0^2) \bmod \mathbf{Z}^2$ and $P_{\mathrm{equi}}(n) = (n/N_0, 0) \bmod \mathbf{Z}^2$ (with P_{rat} remaining trivial), with the torus T increasing from the trivial torus $\{0\}^2$ to $\mathbf{T}^1 \times \{0\}$. A second transition occurs when N exceeds N_0^2, at which point P_{equi} encompasses all of P. Evolving things in a somewhat different direction, if one then increases F so that $F(1)$ is much larger than N_0^2, then P will now entirely consist of a rational component P_{rat}. These sorts of dynamics are not directly seen if one only looks at the asymptotic theory, which roughly speaking is concerned with the limit after taking $N \to \infty$, and *then* taking a second limit by making the growth function F go to infinity.

There is a multidimensional version of Proposition 1.1.17, but we will not describe it here; see [**GrTa2011**] for a statement (and also see the next section for the ultralimit counterpart of this statement).

Remark 1.1.21. These single-scale abelian equidistribution theorems are a special case of a more general single-scale *nilpotent* equidistribution theorem, which will play an important role in later aspects of the theory, and which was the main result of the aforementioned paper of Ben Green and myself.

As an example of this theorem in action, we give a single-scale strengthening of Exercise 1.1.8 (and Exercise 1.1.15):

Exercise 1.1.23 (Recurrence)**.** Let P be a polynomial map from $\mathbf{Z}$ to $\mathbf{T}^d$ of degree s, and let $N \geq 1$ be an integer. Show that for every $\varepsilon > 0$ and $N > 1$, and every integer $n_0 \in [N]$, we have

$$|\{n \in [N] : \|P(n) - P(n_0)\| \leq \varepsilon\}| \gg_{d,s} \varepsilon^{O_{d,s}(1)} N.$$

Exercise 1.1.24 (Multiple recurrence)**.** With the notation of Exercise 1.1.23, establish that

$$\begin{aligned}|\{r \in&[-N, N] : \|P(n_0 + jr) - P(n_0)\| \leq \varepsilon \text{ for } j = 0, 1, \ldots, k-1\}| \\ &\gg_{d,s,k} \varepsilon^{O_{d,s,k}(1)} N\end{aligned}$$

for any $k \geq 1$.

Exercise 1.1.25 (Syndeticity)**.** A set of integers is *syndetic* if it has bounded gaps (or equivalently, if a finite number of translates of this set can cover all of $\mathbf{Z}$). Let $P \colon \mathbf{Z} \to \mathbf{T}^d$ be a polynomial and let $\varepsilon > 0$. Show that the set $\{n \in \mathbf{Z} \colon \|P(n) - P(n_0)\| \leq \varepsilon\}$ is syndetic. (*Hint:* First reduce to the case when P is (totally) asymptotically equidistributed. Then, if N is large enough, show (by inspection of the proof of Exercise 1.1.21) that the translates $P(\cdot + n_0)$ are ε-equidistributed on $[N]$ uniformly for all $n \in \mathbf{Z}$, for any fixed $\varepsilon > 0$. Note how the asymptotic theory and the single-scale theory need to work together to obtain this result.)

1.1.3. Ultralimit equidistribution theory. The single-scale theory was somewhat more complicated than the asymptotic theory, in part because one had to juggle parameters such as N, δ, and (for the equidistribution theorems) F as well. However, one can clean up this theory somewhat (especially if one does not wish to quantify the dependence of bounds on the equidistribution parameter δ) by using an ultralimit, which causes the δ

and F parameters to disappear, at the cost of converting the finitary theory to an infinitary one. Ultralimit analysis is discussed in Section 2.1; we give a quick review here.

We first fix a *non-principal ultrafilter* $\alpha_\infty \in \beta\mathbf{N}\backslash\mathbf{N}$ (see Section 2.1 for a definition of a non-principal ultrafilter). A property P_α pertaining to a natural number α is said to hold *for all α sufficiently close to α_∞* if the set of α for which P_α holds lies in the ultrafilter α_∞. Two sequences $(x_\alpha)_{\alpha\in\mathbf{N}}, (y_\alpha)_{\alpha\in\mathbf{N}}$ of objects are *equivalent* if one has $x_\alpha = y_\alpha$ for all α sufficiently close to α_∞, and we define the *ultralimit* $\lim_{\alpha\to\alpha_\infty} x_\alpha$ to be the equivalence class of all sequences equivalent to $(x_\alpha)_{\alpha\in\mathbf{N}}$, with the convention that x is identified with its own ultralimit $\lim_{\alpha\to\alpha_\infty} x_\alpha$. Given any sequence X_α of sets, the *ultraproduct* $\prod_{\alpha\to\alpha_\infty} X_\alpha$ is the space of all ultralimits $\lim_{\alpha\to\alpha_\infty} x_\alpha$, where $x_\alpha \in X_\alpha$ for all α sufficiently close to α_∞. The ultraproduct $\prod_{\alpha\to\alpha_\infty} X$ of a single set X is the *ultrapower* of X and is denoted ${}^*X$.

Ultralimits of real numbers (i.e., elements of ${}^*\mathbf{R}$) will be called *limit real numbers*; similarly one defines limit natural numbers, limit complex numbers, etc. Ordinary numbers will be called *standard* numbers to distinguish them from limit numbers, thus for instance a limit real number is an ultralimit of standard real numbers. All the usual arithmetic operations and relations on standard numbers are inherited by their limit analogues; for instance, a limit real number $\lim_{\alpha\to\alpha_\infty} x_\alpha$ is larger than another $\lim_{\alpha\to\alpha_\infty} y_\alpha$ if one has $x_\alpha > y_\alpha$ for all α sufficiently close to α_∞. The axioms of a non-principal ultrafilter ensure that these relations and operations on limit numbers obey the same axioms as their standard counterparts[3].

Ultraproducts of sets will be called *limit sets*; they are roughly analogous to "elementary sets" in measure theory. Ultraproducts of finite sets will be called *limit finite sets*. Thus, for instance, if $N = \lim_{\alpha\to\alpha_\infty} N_\alpha$ is a limit natural number, then $[N] = \prod_{\alpha\to\alpha_\infty}[N_\alpha]$ is a limit finite set, and can be identified with the set of limit natural numbers between 1 and N.

Remark 1.1.22. In the language of *non-standard analysis*, limit numbers and limit sets are known as *non-standard numbers* and *internal sets*, respectively. We will, however, use the language of ultralimit analysis rather than non-standard analysis in order to emphasise the fact that limit objects are the ultralimits of standard objects; see Section 2.1 for further discussion of this perspective.

[3]The formalisation of this principle is *Łos's theorem*, which roughly speaking asserts that any first-order sentence which is true for standard objects, is also true for their limit counterparts.

Given a sequence of functions $f_\alpha\colon X_\alpha \to Y_\alpha$, we can form the *ultralimit* $\lim_{\alpha\to\alpha_\infty} f_\alpha\colon \lim_{\alpha\to\alpha_\infty} X_\alpha \to \lim_{\alpha\to\alpha_\infty} Y_\alpha$ by the formula

$$(\lim_{\alpha\to\alpha_\infty} f_\alpha)\left(\lim_{\alpha\to\alpha_\infty} x_\alpha\right) := \lim_{\alpha\to\alpha_\infty} f_\alpha(x_\alpha);$$

one easily verifies that this is a well-defined function between the two ultraproducts. We refer to ultralimits of functions as *limit functions*; they are roughly analogous to "simple functions" in measurable theory. We identify every standard function $f\colon X \to Y$ with its ultralimit $\lim_{\alpha\to\alpha_\infty} f\colon {}^*X \to {}^*Y$, which extends the original function f.

Now we introduce limit asymptotic notation, which is deliberately chosen to be similar (though not identical) to ordinary asymptotic notation. Given two limit numbers X, Y, we write $X \ll Y$, $Y \gg X$, or $X = O(Y)$ if we have $|X| \leq CY$ for some standard $C > 0$. We also write $X = o(Y)$ if we have $|X| \leq cY$ for every standard $c > 0$; thus for any limit numbers X, Y with $Y > 0$, exactly one of $|X| \gg Y$ and $X = o(Y)$ is true. A limit real is said to be *bounded* if it is of the form $O(1)$, and *infinitesimal* if it is of the form $o(1)$; similarly for limit complex numbers. Note that the bounded limit reals are a subring of the limit reals, and the infinitesimal limit reals are an ideal of the bounded limit reals.

Exercise 1.1.26 (Relation between limit asymptotic notation and ordinary asymptotic notation)**.** Let $X = \lim_{\alpha\to\alpha_\infty} X_\alpha$ and $Y = \lim_{\alpha\to\alpha_\infty} Y_\alpha$ be two limit numbers.

(i) Show that $X \ll Y$ if and only if there exists a standard $C > 0$ such that $|X_\alpha| \leq CY_\alpha$ for all α sufficiently close to α_0.

(ii) Show that $X = o(Y)$ if and only if, for every standard $\varepsilon > 0$, one has $|X_\alpha| \leq \varepsilon Y_\alpha$ for all α sufficiently close to α_0.

Exercise 1.1.27. Show that every bounded limit real number x has a unique decomposition $x = \mathrm{st}(x) + (x - \mathrm{st}(x))$, where $\mathrm{st}(x)$ is a standard real (called the *standard part* of x) and $x - \mathrm{st}(x)$ is infinitesimal.

We now give the analogue of single-scale equidistribution in the ultralimit setting.

Definition 1.1.23 (Ultralimit equidistribution)**.** Let $X = (X, d)$ be a standard compact metric space, let N be an unbounded limit natural number, and let $x\colon [N] \to {}^*X$ be a limit function. We say that x is *equidistributed* with respect to a (standard) Borel probability measure μ on X if one has

$$\mathrm{st}\mathbf{E}_{n\in[N]} f(x(n)) = \int_X f\ d\mu$$

for all standard continuous functions $f \in C(X)$. Here, we define the expectation of a limit function in the obvious limit manner, thus

$$\mathbf{E}_{n\in[N]} f(x(n)) = \lim_{\alpha\to\alpha_\infty} \mathbf{E}_{n\in[N_\alpha]} f(x_\alpha(n))$$

if $N = \lim_{\alpha\to\alpha_\infty} N_\alpha$ and $x = \lim_{\alpha\to\alpha_\infty} x_\alpha$.

We say that x is *totally equidistributed* relative to μ if the sequence $n \mapsto x(qn + r)$ is equidistributed on $[N/q]$ for every standard $q > 0$ and $r \in \mathbf{Z}$ (extending x arbitrarily outside $[N]$ if necessary).

Remark 1.1.24. One could just as easily replace the space of continuous functions by any dense subclass in the uniform topology, such as the space of Lipschitz functions.

The ultralimit notion of equidistribution is closely related to that of both asymptotic equidistribution and single-scale equidistribution, as the following exercises indicate:

Exercise 1.1.28 (Asymptotic equidistribution vs. ultralimit equidistribution)**.** Let $x\colon \mathbf{N} \to X$ be a sequence into a standard compact metric space (which can then be extended from a map from ${}^*\mathbf{N}$ to ${}^*X$ as usual), let μ be a Borel probability measure on X. Show that x is asymptotically equidistributed on $\mathbf{N}$ with respect to μ if and only if x is equidistributed on $[N]$ for every unbounded natural number N and every choice of non-principal ultrafilter α_∞.

Exercise 1.1.29 (Single-scale equidistribution vs. ultralimit equidistribution)**.** For every $\alpha \in \mathbf{N}$, let N_α be a natural number that goes to infinity as $\alpha \to \infty$, let $x_\alpha\colon [N_\alpha] \to X$ be a map to a standard compact metric space. Let μ be a Borel probability measure on X. Write $N := \lim_{\alpha\to\alpha_\infty} N_\alpha$ and $x := \lim_{\alpha\to\alpha_\infty} x_\alpha$ for the ultralimits. Show that x is equidistributed with respect to μ if and only if, for every standard $\delta > 0$, x_α is δ-equidistributed with respect to μ for all α sufficiently close to α_∞.

In view of these correspondences, it is thus not surprising that one has ultralimit analogues of the asymptotic and single-scale theory. These analogues tend to be *logically equivalent* to the single-scale counterparts (once one concedes all quantitative bounds), but are *formally similar* (though not identical) to the asymptotic counterparts, thus providing a bridge between the two theories, which we can summarise by the following three statements:

(i) Asymptotic theory is analogous to ultralimit theory (in particular, the statements and proofs are formally similar);

(ii) ultralimit theory is logically equivalent to qualitative finitary theory; and

(iii) quantitative finitary theory is a strengthening of qualitative finitary theory.

For instance, here is the ultralimit version of the Weyl criterion:

Exercise 1.1.30 (Ultralimit Weyl equidistribution criterion)**.** Let $x\colon [N] \to {}^*\mathbf{T}^d$ be a limit function for some unbounded N and standard d. Then x is equidistributed if and only if

$$\mathbf{E}_{n\in[N]} e(k \cdot x(n)) = o(1) \tag{1.7}$$

for all standard $k \in \mathbf{Z}^d\backslash\{0\}$. *Hint:* Mimic the proof of Proposition 1.1.2.

Exercise 1.1.31. Use Exercise 1.1.30 to recover a weak version of Proposition 1.1.13, in which the quantities δ^{c_d}, δ^{C_d} are replaced by (ineffective) functions of δ that decay to zero as $\delta \to 0$. Conversely, use this weak version to recover Exercise 1.1.30. (*Hint:* Similar arguments appear in Section 2.1.)

Exercise 1.1.32. With the notation of Exercise 1.1.30, show that x is totally equidistributed if and only if

$$\mathbf{E}_{n\in[N]} e(k \cdot x(n)) e(\theta n) = o(1)$$

for all standard $k \in \mathbf{Z}^d\backslash\{0\}$ and standard rational θ.

Exercise 1.1.33. With the notation of Exercise 1.1.30, show that x is equidistributed in $\mathbf{T}^d$ on $[N]$ if and only if $k \cdot x$ is equidistributed in $\mathbf{T}$ on $[N]$ for every non-zero standard $k \in \mathbf{Z}^d$.

Now we establish the ultralimit version of the linear equidistribution criterion:

Exercise 1.1.34. Let $\alpha, \beta \in {}^*\mathbf{T}^d$, and let N be an unbounded integer. Show that the following are equivalent:

(i) The sequence $n \mapsto n\alpha + \beta$ is equidistributed on $[N]$.

(ii) The sequence $n \mapsto n\alpha + \beta$ is totally equidistributed on $[N]$.

(iii) α is *irrational to scale* $1/N$, in the sense that $k \cdot \alpha \neq O(1/N)$ for any non-zero standard $k \in \mathbf{Z}^d$.

Note that in the ultralimit setting, assertions such as $k \cdot \alpha \neq O(1/N)$ make perfectly rigorous sense (it means that $|k\cdot\alpha| \geq C/N$ for every standard C), but when using finitary asymptotic big-O notation

Next, we establish the analogue of the van der Corput lemma:

Exercise 1.1.35 (van der Corput lemma, ultralimit version)**.** Let N be an unbounded integer, and let $x\colon [N] \to {}^*\mathbf{T}^d$ be a limit sequence. Let $H = o(N)$ be unbounded, and suppose that the derivative sequence $\partial_h x\colon n \mapsto$

$x(n+h) - x(n)$ is equidistributed on $[N]$ for $\gg H$ values of $h \in [H]$ (by extending x arbitrarily outside of $[N]$). Show that x is equidistributed on $[N]$. Similarly, "equidistributed" is replaced by "totally equidistributed".

Here is the analogue of the Vinogradov lemma:

Exercise 1.1.36 (Vinogradov lemma, ultralimit version)**.** Let $\alpha \in {}^*\mathbf{T}$, N be unbounded, and $\varepsilon > 0$ be infinitesimal. Suppose that $\|n\alpha\|_{\mathbf{T}} \leq \varepsilon$ for $\gg N$ values of $n \in [-N, N]$. Show that there exists a positive standard integer q such that $\|\alpha q\|_{\mathbf{T}} \ll \varepsilon/N$.

These two lemmas allow us to establish the ultralimit polynomial equidistribution theory:

Exercise 1.1.37. Let $P\colon {}^*\mathbf{Z} \to {}^*\mathbf{T}^d$ be a polynomial sequence $P(n) := \alpha_s n^s + \cdots + \alpha_0$ with s, d standard, and $\alpha_0, \ldots, \alpha_s \in {}^*\mathbf{T}^d$. Let N be an unbounded natural number. Suppose that P is not totally equidistributed on $[N]$. Show that there exists a non-zero standard $k \in \mathbf{Z}^d$ with $\|k \cdot \alpha_s\|_{\mathbf{T}} \ll N^{-s}$.

Exercise 1.1.38. With the hypotheses of Exercise 1.1.37, show in fact that there exists a non-zero standard $k \in \mathbf{Z}^d$ such that $\|k \cdot \alpha_i\|_{\mathbf{T}} \ll N^{-i}$ for all $i = 0, \ldots, s$.

Exercise 1.1.39 (Ultralimit equidistribution theorem for abelian polynomial sequences)**.** Let P be a polynomial map from ${}^*\mathbf{Z}$ to ${}^*\mathbf{T}^d$ of some standard degree $s \geq 0$. Let N be an unbounded natural number. Then there exists a decomposition

$$P = P_{\mathrm{smth}} + P_{\mathrm{equi}} + P_{\mathrm{rat}}$$

into polynomials of degree s, where

(i) (P_{smth} is smooth) The i^{th} coefficient $\alpha_{i,\mathrm{smth}}$ of P_{smth} has size $O(N^{-i})$. In particular, on the interval $[N]$, P_{smth} is Lipschitz with homogeneous norm $O(1/N)$.

(ii) (P_{equi} is equidistributed) There exists a standard subtorus T of $\mathbf{T}^d$, such that P_{equi} takes values in T and is totally equidistributed on $[N]$ in this torus.

(iii) (P_{rat} is rational) The coefficients $\alpha_{i,\mathrm{rat}}$ of P_{rat} are standard rational elements of $\mathbf{T}^d$. In particular, there is a standard positive integer q such that $qP_{\mathrm{rat}} = 0$ and P_{rat} is periodic with period q.

Exercise 1.1.40. Show that the torus T is uniquely determined by P, and decomposition $P = P_{\mathrm{smth}} + P_{\mathrm{equi}} + P_{\mathrm{rat}}$ in Exercise 1.1.39 is unique up to expressions taking values in T (i.e., if one is given another decomposition

$P = P'_{\mathrm{smth}} + P'_{\mathrm{equi}}, P'_{\mathrm{rat}}$, then P_i and P'_i differ by expressions taking values in T).

Exercise 1.1.41 (Recurrence)**.** Let P be a polynomial map from $*\mathbf{Z}$ to $*\mathbf{T}^d$ of some standard degree s, and let N be an unbounded natural number. Show that for every standard $\varepsilon > 0$ and every $n_0 \in N$, we have

$$|\{n \in [N] : \|P(n) - P(n_0)\| \leq \varepsilon\}| \gg N$$

and more generally

$$|\{r \in [-N, N] : \|P(n_0 + jr) - P(n_0)\| \leq \varepsilon \text{ for } j = 0, 1, \ldots, k-1\}| \gg N$$

for any standard k.

As before, there are also multidimensional analogues of this theory. We shall just state the main results without proof:

Definition 1.1.25 (Multidimensional equidistribution)**.** Let X be a standard compact metric space, let N be an unbounded limit natural number, let $m \geq 1$ be standard, and let $x \colon [N]^m \to *X$ be a limit function. We say that x is *equidistributed* with respect to a (standard) Borel probability measure μ on X if one has

$$\mathrm{st}\mathbf{E}_{n \in [N]^m} 1_B(n/N) f(x(n)) = \mathrm{mes}(\Omega) \int_X f \, d\mu$$

for every standard box $B \subset [0,1]^m$ and for all standard continuous functions $f \in C(X)$.

We say that x is *totally equidistributed* relative to μ if the sequence $n \mapsto x(qn + r)$ is equidistributed on $[N/q]^d$ for every standard $q > 0$ and $r \in \mathbf{Z}^m$ (extending x arbitrarily outside $[N]$ if necessary).

Remark 1.1.26. One can replace the indicators 1_B by many other classes, such as indicators of standard convex sets, or standard open sets whose boundary has measure zero, or continuous or Lipschitz functions.

Theorem 1.1.27 (Multidimensional ultralimit equidistribution theorem for abelian polynomial sequences)**.** *Let $m, d, s \geq 0$ be standard integers, and let P be a polynomial map from $*\mathbf{Z}^m$ to $*\mathbf{T}^d$ of degree s. Let N be an unbounded natural number. Then there exists a decomposition*

$$P = P_{\mathrm{smth}} + P_{\mathrm{equi}} + P_{\mathrm{rat}}$$

into polynomials of degree s, where

(i) *(P_{smth} is smooth) The i^{th} coefficient $\alpha_{i,\mathrm{smth}}$ of P_{smth} has size $O(N^{-|i|})$ for every multi-index $i = (i_1, \ldots, i_m)$. In particular, on the interval $[N]$, P_{smth} is Lipschitz with homogeneous norm $O(1/N)$.*

(ii) *(P_{equi} is equidistributed) There exists a standard subtorus T of $\mathbf{T}^d$, such that P_{equi} takes values in T and is totally equidistributed on $[N]^m$ in this torus.*

(iii) *(P_{rat} is rational) The coefficients $\alpha_{i,\mathrm{rat}}$ of P_{rat} are standard rational elements of $\mathbf{T}^d$. In particular, there is a standard positive integer q such that $qP_{\mathrm{rat}} = 0$ and P_{rat} is periodic with period q.*

Proof. This is implicitly in [**GrTa2011**]; the result is phrased using the language of single-scale equidistribution, but this easily implies the ultralimit version. □

1.2. Roth's theorem

We now give a basic application of Fourier analysis to the problem of counting additive patterns in sets, namely the following famous theorem of Roth [**Ro1953**]:

Theorem 1.2.1 (Roth's theorem)**.** *Let A be a subset of the integers $\mathbf{Z}$ whose upper density*

$$\overline{\delta}(A) := \limsup_{N\to\infty} \frac{|A \cap [-N,N]|}{2N+1}$$

is positive. Then A contains infinitely many arithmetic progressions $a, a+r, a+2r$ of length three, with $a \in \mathbf{Z}$ and $r > 0$.

This is the first non-trivial case of *Szemerédi's theorem* [**Sz1975**], which is the same assertion but with length three arithmetic progressions replaced by progressions of length k for any k.

As it turns out, one can prove Roth's theorem by an application of linear Fourier analysis by comparing the set A (or more precisely, the indicator function 1_A of that set, or of pieces of that set) against linear characters $n \mapsto e(\alpha n)$ for various frequencies $\alpha \in \mathbf{R}/\mathbf{Z}$. There are two extreme cases to consider (which are model examples of a more general dichotomy between structure and randomness, as discussed in [**Ta2008**]). One is when A is aligned almost completely with one of these linear characters, for instance, by being a *Bohr set* of the form

$$\{n \in \mathbf{Z} : \|\alpha n - \theta\|_{\mathbf{R}/\mathbf{Z}} < \varepsilon\}$$

or, more generally, of the form

$$\{n \in \mathbf{Z} : \alpha n \in U\}$$

for some multi-dimensional frequency $\alpha \in \mathbf{T}^d$ and some open set U. In this case, arithmetic progressions can be located using the equidistribution theory from Section 1.1. At the other extreme, one has *Fourier-uniform* or *Fourier-pseudorandom sets*, whose correlation with any linear character is