for each $1 \leq j \leq n$. We can estimate

$$|1 - \mu + \mu \cos(2\pi\xi \cdot v_j)| \leq \exp\big(-\Omega\big(\mu \|2\xi \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2\big)\big)$$

(cf. (7.1)) and then make the change of variables $t = 2\xi \cdot v_j$ (using (3.8) to estimate the volume of the $d-1$-dimensional balls that are integrated out) to reduce to showing the one-dimensional estimate

$$\frac{1}{|v_j|} \int_{|t| \leq 2|v_j|/k} \exp\big(-\Omega\big(\mu n \|t\|_{\mathbf{R}/\mathbf{Z}}^2\big)\big)\, dt = O\left(\frac{1}{\sqrt{\mu n}}\right).$$

Subdividing the t variable into unit intervals and using the periodicity of $\|t\|_{\mathbf{R}/\mathbf{Z}}$ and the hypothesis $|v_j| \geq 1$, the claim then follows from the easily verified estimate

$$\int_{-\infty}^{\infty} \exp(-\Omega(\mu n |t|^2))\, dt = O\left(\frac{1}{\sqrt{\mu n}}\right).$$

□

One can similarly develop analogs of many of the results of the preceding section, though the analysis is a little more technical as the analogs of Corollary 7.12 are somewhat messier. See [167] for further development of this theory.

Exercises

7.3.1 Prove (7.4).

7.3.2 Establish the following dimension-independent analog of the Esséen concentration inequality:

$$\sup_{x_0 \in \mathbf{R}^d} \mathbf{P}\big(e^{-\pi|X-x_0|^2}\big) \leq \int_{\xi \in \mathbf{R}^d} |\mathbf{E}(e(\xi \cdot X))| e^{-\pi|\xi|^2}\, d\xi.$$

7.3.3 [367] Obtain an analog of Exercise 7.2.3 for the probability $\mathbf{P}(X_{\mathbf{v}}^{\mu} \in B)$ for some unit ball B, assuming that, for every proper subspace of $\mathbf{R}^n$, at most $n-k$ of the vectors lie within a unit distance of this subspace.

7.3.4 Use the previous exercise to develop an analog of Erdős, results in any dimension [108, 367].

7.4 Inverse Littlewood–Offord results

In the preceding sections we considered *direct* Littlewood–Offord results, in which some assumptions were made on the steps $\mathbf{v} = (v_1, \ldots, v_n)$, and as a conclusion some upper bounds were obtained for concentration probabilities such as $\mathbf{P}(X_{\mathbf{v}}^{(\mu)} = x)$. In many applications it is of more interest to establish *inverse* Littlewood–Offord results, in which a lower bound on a concentration probability

is assumed, and some structural property of $\mathbf{v}$ is deduced as a consequence. Of course, every direct Littlewood–Offord result can be converted into an inverse by taking contrapositives. For instance, from Corollary 7.13 we know that if $v_1, \ldots, v_n$ live in a torsion-free group Z and

$$\mathbf{P}\big(X_{\mathbf{v}}^{(\mu)} = x\big) \geq \frac{1}{\sqrt{\mu k}}$$

for some $0 < \mu \leq 1$ and some $x \in Z$, then at most $O(k)$ of the steps $v_1, \ldots, v_n$ are non-zero. Similarly, from Corollary 7.16, we see that if $v_1, \ldots, v_n$ are positive integers and $\mathbf{P}(X_{\mathbf{v}}^{(\mu)} = x)$ is much larger than $\mu^{-1/2} n^{-3/2}$ for some $0 < \mu \leq 1$ and $x \in \mathbf{Z}$, then at least two of the v_j are equal (in fact one can easily establish that a large number of pairs (v_i, v_j) must be equal).

Now we consider inverse Littlewood–Offord theorems that give more structure on the steps $v_1, \ldots, v_n$. The results in this section can be viewed in analogy with inverse sum set estimates, in which one assumes that a certain set A has small doubling constant and concludes some structural information on A, for instance containing A inside a progression. For simplicity we shall focus on the case $\mu = 1$ (though one can use results such as Corollary 7.12 or Lemma 7.14 to then extend to more general μ).

Let us start with an example when $\max_x \mathbf{P}(X_{\mathbf{v}}^1 = x)$ is large. This example has been the main motivation of our results.

Example 7.19 Let P be a symmetric generalized arithmetic progression of (constant) rank d and volume V in Z. Let $v_1, \ldots, v_n$ be (not necessarily different) elements of V. Then the sum $\sum_{i=1}^n \epsilon_i v_i$ takes values in the generalized arithmetic progression nP which have volume $n^d V$. From the pigeonhole principle it follows that

$$\max_x \mathbf{P}\big(X_{\mathbf{v}}^1 = x\big) \geq n^{-d} V^{-1}. \tag{7.6}$$

The above example shows that if the elements of $\mathbf{v}$ belong to a generalized arithmetic progression with small rank and small volume then $\mathbf{P}_\mu(\mathbf{v})$ is large. One might hope that the inverse of this also holds, namely,

If $P_\mu(\mathbf{v})$ is large, then the elements of $\mathbf{v}$ belong to a generalized arithmetic progression with small rank and small volume.

We are going to present a few results which support this statement. Let us first give a simple, but rather weak, result.

Proposition 7.20 *Let $\mathbf{v} = (v_1, \ldots, v_n)$ be a tuple in an additive group Z which is either torsion-free or finite of odd order, such that $\mathbf{P}(X_{\mathbf{v}}^{(1)} = x) > 2^{-d-1}$ for some $x \in Z$ and $d \geq 0$. Then all the steps $v_1, \ldots, v_n$ are contained in a cube $[-1, 1]^d \cdot (w_1, \ldots, w_d)$ of dimension d.*

Proof Suppose the conclusion failed. Then from Lemma 4.35 we see that $\mathbf{v}$ must contain a dissociated subword $\mathbf{w} = (w_1, \ldots, w_{d+1})$ of length $d+1$. By conditioning on the variables not associated to $\mathbf{w}$, we observe that

$$2^{-d-1} < \mathbf{P}\big(X_{\mathbf{v}}^{(1)} = x\big) \leq \sup_{y \in Z} \mathbf{P}\big(X_{\mathbf{w}}^{(1)} = y\big).$$

On the other hand, since $\mathbf{w}$ is dissociated, and Z has no 2-torsion, all the sums in $X_{\mathbf{w}}^{(1)}$ are distinct and so $\mathbf{P}(X_{\mathbf{w}}^{(1)} = y) \leq 2^{-d-1}$, thus yielding the desired contradiction. □

In practice, this proposition is not very useful because the dimension d of the cube can be rather large (typically it is like $\log n$). However, one can lower dimension its by increasing the side lengths, and allowing some exceptional steps v_j to lie outside of the resulting progression.

Proposition 7.21 *Let Z be either torsion-free or finite of odd order. For any integer $d \geq 1$, there is a positive constant δ_d such that the following holds. Let $k \geq 2$ be an integer, let $x \in Z$, and let $\mathbf{v} = (v_1, \ldots, v_n)$ be a tuple in Z. Then either*

$$\mathbf{P}\big(X_{\mathbf{v}}^{(1)} = x\big) \leq \delta_d k^{-d}$$

or there exists a progression $P = [-k, k]^{d-1} \cdot (w_1, \ldots, w_{d-1})$ in Z such that for all but at most k^2 exceptional values of $j \in [1, n]$, there exists $a_0 \in [1, k]$ such that $a_0 v_j \in P$.

Note that Corollary 7.13 (with $\mu = 1$) can be thought of as the $d = 1$ case of this proposition, while Proposition 7.20 can be viewed as the limiting case $k = 1$. Of course one should take $k < \sqrt{n}$ to avoid the claim being vacuous.

Proof Call a tuple $(w_1, \ldots, w_r)$ *k-dissociated* if the progression $[-k, k]^r \cdot (w_1, \ldots, w_r)$ is proper. We now construct an k-dissociated tuple $(w_1, \ldots, w_r)$ for some $0 \leq r \leq d$ by the following algorithm.

- Step 0. Initialize $r = 0$. In particular, $(w_1, \ldots, w_r)$ is trivially k-dissociated, and from Corollary 7.12 we have

$$\mathbf{P}\left(X^{(1/4d)}_{\mathbf{v}^{d-r} w_1^{k^2} \ldots w_r^{k^2}} = 0\right) \geq \mathbf{P}\big(X_{\mathbf{v}}^{(1)} = x\big). \tag{7.7}$$

- Step 1. Count how many $1 \leq j \leq n$ there are such that $(w_1, \ldots, w_r, v_j)$ is k-dissociated. If this number is less than k^2, halt the algorithm. Otherwise, move on to Step 2.
- Step 2. Applying Corollary 7.12, we can locate a v_j such that $(w_1, \ldots, w_r, v_j)$ is k-dissociated, and

$$\mathbf{P}\left(X^{(1/4d)}_{\mathbf{v}^{d-r} w_1^{k^2} \ldots w_r^{k^2}} = 0\right) \leq \mathbf{P}\left(X^{(1/4d)}_{\mathbf{v}^{d-r-1} w_1^{k^2} \ldots w_r^{k^2} v_j^{k^2}} = 0\right).$$

We then set $w_{r+1} := v_j$ and increase r to $r+1$. Return to Step 1. Note that $(w_1, \dots, w_r)$ remains k-dissociated, and (7.7) remains true, when doing so.

Suppose that we terminate at some step $r \leq d-1$. Then we have an r-tuple $(w_1, \dots, w_r)$ which is k-dissociated, but such that $(w_1, \dots, w_r, v_j)$ is k-dissociated for at most k^2 values of v_j. Unwinding the definitions, this shows that for all but at most k^2 values of v_j, there exists $a_0 \in [1,k]$ such that $a_0 v_j \in Q - Q$, where $Q := [0,k]^r \cdot (w_1, \dots, w_r)$ and $r \leq d-1$. The claim then follows by adding some dummy vectors to the w_j.

Now we prove that we must indeed terminate at some step $r \leq d-1$. Assume (for a contradiction) that we have reached step d. Then we have an k-dissociated tuple $(w_1, \dots, w_d)$ such that

$$\mathbf{P}\big(X_{\mathbf{v}}^{(1)} = x\big) \leq \mathbf{P}\left(X^{(1/4d)}_{w_1^{k^2} \dots w_d^{k^2}} = 0\right).$$

Let $\Gamma \subset \mathbf{Z}^d$ be the lattice

$$\Gamma := \{(m_1, \dots, m_d) \in \mathbf{Z}^d : m_1 w_1 + \dots + m_d w_d = 0\},$$

then by using independence we can write

$$\mathbf{P}\big(X_{\mathbf{v}}^{(1)} = x\big) \leq \mathbf{P}\left(X^{(1/4d)}_{w_1^{k^2} \dots w_d^{k^2}} = 0\right) = \sum_{(m_1,\dots,m_d) \in \Gamma} \prod_{j=1}^{d} \mathbf{P}\left(X_{1^{k^2}}^{(1/4d)} = m_j\right) \quad (7.8)$$

where $X_{1^{k^2}}^{(1/4d)} = \eta_1^{(1/4d)} + \dots + \eta_{k^2}^{(1/4d)}$.

Now we use a volume-packing argument. A simple computation involving the binomial formula (or induction on the k^2 parameter) shows that the expression $\mathbf{P}(X_{1^{k^2}}^{(1/4d)} = m)$ is even in m, and decreasing for positive m. It is also $\Theta_d(1/k)$ when $|m| \leq k$ (this can be seen either from Stirling's formula (1.52), or from Corollary 7.13 and variance and monotonicity considerations). Thus we have

$$\mathbf{P}\left(X_{1^{k^2}}^{(1/4d)} = m\right) = O_d\left(\frac{1}{k} \sum_{m' \in m + (-k/2, k/2)} \mathbf{P}\left(X_{1^{k^2}}^{(1/4d)} = m\right)\right)$$

and hence from (7.8) we have

$$\mathbf{P}\big(X_{\mathbf{v}}^{(1)} = x\big) \leq O_d\left(k^{-d} \sum_{(m_1,\dots,m_d) \in \Gamma} \sum_{(m_1',\dots,m_d') \in (m_1,\dots,m_d) + (-k/2, k/2)^d} \prod_{j=1}^{d} \mathbf{P}\left(X_{1^{k^2}}^{(1/4d)} = m_j\right)\right).$$

Since $(w_1, \ldots, w_d)$ is k-dissociated, all the $(m'_1, \ldots, m'_d)$ tuples in $\Gamma + (-k/2, k/2)^d$ are different. Thus, we conclude

$$\mathbf{P}(X_{\mathbf{v}}^{(1)} = x) \leq O_d\left(k^{-d} \sum_{(m_1,\ldots,m_d)\in\mathbf{Z}^d} \prod_{j=1}^{d} \mathbf{P}\left(X_{1^{k^2}}^{(1/4d)} = m_j\right)\right).$$

But from the union bound we have

$$\sum_{(m_1,\ldots,m_d)\in\mathbf{Z}^d} \prod_{j=1}^{d} \mathbf{P}\left(X_{1^{k^2}}^{(1/4d)} = m_j\right) = 1.$$

To complete the proof, set the constant δ_d in the proposition to be larger than the hidden constant in $O_d(k^{-d})$. □

The a_0 factor in the above proposition is somewhat undesirable. With some more effort, one can remove this factor, but at the cost of enlarging the progression somewhat.

Theorem 7.22 (Inverse Littlewood–Offord theorem) *[366] Let $0 < \mu < 1$ and let α and A be arbitrary positive constants. Then there is a constant $B = B(\mu, \alpha, A)$ such that the following holds. Assume that $\mathbf{v} = (v_1, \ldots, v_n)$ is a tuple of rational numbers satisfying $\max_x \mathbf{P}(X_{\mathbf{v}}^{\mu} = x) \geq n^{-A}$. Then there is a generalized arithmetic progression P of rational numbers of rank at most B and volume at most n^B which contains all but at most Bn^{α} elements of $\mathbf{v}$.*

The proof of Theorem 7.22 is somewhat lengthy but is a modification of that of Proposition 7.21. For details see [366].

An inverse theorem in a similar spirit for the relative Halász inequality, Lemma 7.14, was also obtained in [365]:

Theorem 7.23 (Inverse Halász inequality) *[365] Let Z be either torsion-free or cyclic of odd prime order. Let $\mathbf{v} = (v_1, \ldots, v_n)$ be a tuple in Z, and suppose that $\varepsilon_0 > \varepsilon_1 > 0$ are such that*

$$\mathbf{P}\left(X_{\mathbf{v}}^{(1)} = 0\right) \geq \varepsilon_1 \mathbf{P}\left(X_{\mathbf{v}}^{(1/4-\varepsilon_0/100)} = 0\right)$$

and

$$\mathbf{P}\left(X_{\mathbf{v}}^{(1)} = 0\right) \geq \left(\frac{3}{4} + 2\varepsilon_0\right)^n.$$

Then there exists a proper progression P of rank $O_{\varepsilon_0,\varepsilon_1}(1)$ and volume $O_{\varepsilon_0,\varepsilon_1}(\frac{1}{\mathbf{P}(X_{\mathbf{v}}^{(1)}=0)})$ which contain the $v_1, \ldots, v_n$.

In fact some additional structural information was obtained, namely that the $v_1, \ldots, v_n$ are mostly contained in the "core" of the progression P, and