7.2 The Fourier-analytic approach

Now we present the Fourier-analytic approach of Halász. It is convenient to use the language of probability theory. For any n-tuple $\mathbf{v} = (v_1, \dots, v_n)$ of steps in an additive group Z, we use the notation $X_\mathbf{v}$ to denote the random variable

$$X_\mathbf{v} := \epsilon_1 v_1 + \cdots + \epsilon_n v_n$$

where $\epsilon_1, \dots, \epsilon_n$ are independent random variables taking values in $\{-1, +1\}$ with probability $1/2$ for each value. Clearly $\mathbf{P}(X_\mathbf{v} = x)$ equals the number of representations of x as $\epsilon_1 v_1 + \cdots + \epsilon_n v_n$ with $\epsilon_1, \dots, \epsilon_n \in \{-1, 1\}$, divided by 2^n. Note that $X_\mathbf{v}$ is invariant under permutations of the n-tuple $\mathbf{v}$. We use $\mathbf{vw}$ to denote the concatenation of $\mathbf{v}$ and $\mathbf{w}$. The Littlewood–Offord problem then asks to control the distribution of $X_\mathbf{v}$ for a given $\mathbf{v}$, while the inverse Littlewood–Offord problem asks for some structural information on $\mathbf{v}$ given some unexpected distributional property of $X_\mathbf{v}$.

It will be useful to consider the more general random variables $X_\mathbf{v}^{(\mu)}$ for any $0 \leq \mu \leq 1$, defined as

$$X_\mathbf{v} := \epsilon_1^{(\mu)} v_1 + \cdots + \epsilon_n^{(\mu)} v_n,$$

where $\epsilon_1^{(\mu)}, \dots, \epsilon_n^{(\mu)}$ are independent random variables which take the values $+1$ and -1 with probability $\mu/2$, and 0 with probability $1 - \mu$. Thus $X_\mathbf{v}^{(\mu)}$ is the same as $X_\mathbf{v}$ when $\mu = 1$, and at the other extreme $\mu = 0$ becomes the constant 0. The intermediate cases correspond to "lazy random walks" with step sizes $v_1, \dots, v_n$. As ϵ_i can be 0 with considerable probability, one expects $X_\mathbf{v}^{(\mu)}$ to be more concentrated than $X_\mathbf{v}$, and this will indeed be the case. In practice, the cases $\mu \leq 1/2$ are more amenable to Fourier analysis than the $\mu = 1$ case due to a certain "positivity" property which we shall come to shortly.

In this section we shall consider the discrete problem of understanding the probabilities $\mathbf{P}(X_\mathbf{v}^{(\mu)} = x)$ that a random variable $X_\mathbf{v}^{(\mu)}$ concentrates at a single point. In the next section we briefly discuss the analogous probability $\mathbf{P}(X_\mathbf{v}^{(\mu)} \in Q)$ for concentration in a cube.

Let us first make some technical reductions to the problem. Firstly, we can reduce to the case when the ambient group Z is finite. This can be achieved by applying a suitable Freiman isomorphism of order n to the steps $v_1, \dots, v_n$ (see Exercise 5.3.3) while noting that this does not affect the distribution of $X_\mathbf{v}$. Secondly, we can reduce further to the case that Z is odd. To see this, observe from Corollary 3.8 that any finite additive group can be written as the product of a 2-torsion group and a group of odd order. The behavior of the random variable $X_\mathbf{v}$, when projected down to the 2-torsion group is trivial (since $+v_j = -v_j$ in this group), so we may, without loss of generality, project onto the other factor. Note that if the original elements $v_1, \dots, v_n$ lived in some torsion-free group such as $\mathbf{Z}^d$, then by Lemma 5.25 we could now place the vectors in a *cyclic* group of odd prime order. (In doing so we may temporarily obscure some of the "dimensional" structure of the elements $v_1, \dots, v_d$, so in some cases it is convenient to revert back to the original ambient group at certain stages of the argument.)

With these reductions we can now express the distribution of $X_{\mathbf{v}}$ in terms of the Fourier transform. As usual we fix a symmetric non-degenerate bilinear form $\xi \cdot x$ on Z.

Lemma 7.11 (Fourier representation of $X_{\mathbf{v}}$) *Let Z be a finite group of odd order. If $\mathbf{v} = (v_1 \dots v_n)$ is an n-tuple of elements of Z, then for any $0 \leq \mu \leq 1$ and $x \in Z$ we have*

$$\mathbf{P}\big(X_{\mathbf{v}}^{(\mu)} = x\big) = \mathbf{E}_{\xi \in Z} \cos(2\pi \xi \cdot x) \prod_{j=1}^{n} (1 - \mu + \mu \cos(2\pi \xi \cdot v_j)).$$

Proof Since the quantity $\prod_{j=1}^{n}(1 - \mu + \mu \cos(2\pi\xi \cdot v_j))$ is an even function of ξ, we can write the right-hand side as

$$\mathbf{E}_{\xi \in Z} e(-\xi \cdot x) \prod_{j=1}^{n} (1 - \mu + \mu \cos(2\pi \xi \cdot v_j)).$$

Observing that $1 - \mu + \mu \cos(2\pi\xi \cdot v_j) = \mathbf{E}(e(\xi \cdot \epsilon_j^{(\mu)} v_j))$ and using the independence of the $\epsilon_j^{(\mu)}$, we can rewrite this as

$$\mathbf{E}\mathbf{E}_{\xi \in Z} e\big(\xi \cdot \big(X_{\mathbf{v}}^{(\mu)} - x\big)\big).$$

But the claim now follows from Lemma 4.5. □

This lemma already highlights the special role of the case $0 \leq \mu \leq \frac{1}{2}$, as in this case $1 - \mu + \mu \cos(2\pi\xi \cdot v_j)$ becomes non-negative. In the further case $0 \leq \mu \leq \frac{1}{4}$, we have the elementary but very useful estimate

$$1 - \mu + \mu \cos(2\pi \xi \cdot v_j) = \exp\big(- \Theta\big(\mu \|\xi \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2\big)\big) \tag{7.1}$$

where we recall that $\|x\|_{\mathbf{R}/\mathbf{Z}}$ denotes the distance to the nearest integer.

From Lemma 7.11 we can immediately establish a number of useful bounds on how one distribution $X_{\mathbf{v}}^{(\mu)}$ controls another.

Corollary 7.12 *Let $\mathbf{v} = (v_1, \dots, v_n)$, $\mathbf{w} = (w_1, \dots, w_m)$ be tuples in an additive group Z which is torsion-free or is finite of odd order. Let $x \in Z$.*

- *(Domination) If $0 \leq \mu \leq \mu' \leq 1$, and at least one of $\mu' \leq 1/2$ or $\mu \leq \mu'/4$ hold, then*

$$\mathbf{P}\big(X_{\mathbf{vw}}^{(\mu')} = x\big) \leq \mathbf{P}\big(X_{\mathbf{v}}^{(\mu)} = 0\big) = \mathbf{E}_{\xi \in Z} \prod_{j=1}^{n} (1 - \mu + \mu \cos(2\pi \xi \cdot v_j)).$$

In particular, if $\mu \leq 1/2$*, then* $X_{\mathbf{v}}^{(\mu)}$ *concentrates more at the origin than anywhere else.*

- *(Duplication) If* $0 \leq \mu \leq 1/2$*, then*

$$\mathbf{P}\big(X_{\mathbf{vw}}^{(\mu)} = x\big) \leq \mathbf{P}\big(X_{\mathbf{v}^k}^{(\mu/k)} = 0\big)$$

for all integers $k \geq 1$*, where we use* $\mathbf{v}^k$ *to denote the concatenation of* k *copies of* $\mathbf{v}$.

- *(Hölder) If* $\mathbf{w}_1, \ldots, \mathbf{w}_k$ *are tuples in* Z *(possibly of different length) and* $0 \leq \mu \leq 1/2$*, then*

$$\mathbf{P}\big(X_{\mathbf{vw}_1\ldots\mathbf{w}_k}^{(\mu)} = x\big) \leq \prod_{i=1}^{k} \mathbf{P}\big(X_{\mathbf{vw}_i^k}^{(\mu)} = 0\big)^{1/k}.$$

Proof As discussed earlier we may take Z to be finite of odd order. In all cases we rewrite the probabilities using Lemma 7.11. The Hölder formula is clear, as is the domination formula when $\mu' \leq 1/2$. In the case $\mu \leq \mu'/4$, one observes the elementary inequality

$$|\cos(\pi\theta)| \leq \frac{3}{4} + \frac{1}{4}\cos(2\pi\theta)$$

and hence (by the triangle inequality)

$$|(1-\mu') + \mu'\cos(\pi\theta)| \leq \left(1 - \frac{\mu'}{4}\right) + \frac{\mu'}{4}\cos(2\pi\theta).$$

The claim then follows from the change of variables $\xi \to 2\xi$ (which is invertible when Z has odd order).

The duplication formula similarly follows from the elementary inequality

$$(1-\mu) + \mu\cos(2\pi\theta) \leq \left(\left(1 - \frac{\mu}{k}\right) + \frac{\mu}{k}\cos(2\pi\theta)\right)^k,$$

which can be seen by taking logarithms and exploiting the concavity of $\log(1-t)$ in the region $0 < t < 1$. □

The above corollary allows one to show that the quantity $\mathbf{P}(X_{\mathbf{v}}^{(\mu)} = 0)$ is fairly stable when one tinkers with the tuple $\mathbf{v}$ (for instance, by adding or removing duplicates) and the parameter μ, at least when $\mu \leq 1/2$. As an application, let us give a Fourier-analytic analog of Corollary 7.4.

Corollary 7.13 *Let* $\mathbf{v} = (v_1, \ldots, v_n)$ *be an* n*-tuple in a torsion-free group* Z *such that at least* k *of the* v_j *are non-zero. Then for all* $0 < \mu \leq 1$ *and* $x \in Z$ *we have*

$$\mathbf{P}\big(X_{\mathbf{v}}^{(\mu)} = x\big) = O\left(\frac{1}{\sqrt{\mu k}}\right).$$

Proof Using the domination property we may take $\mu \leq 1/2$. Without loss of generality we may take $v_1, \ldots, v_k$ to be non-zero. Applying Corollary 7.12 repeatedly we have

$$\begin{aligned}\mathbf{P}\big(X_{\mathbf{v}}^{(\mu)} = x\big) &\leq \mathbf{P}\big(X_{\mathbf{v}\mathbf{v}}^{(\mu/2)} = 0\big)\\ &\leq \mathbf{P}\big(X_{\mathbf{v}v_j^k}^{(\mu/2)} = 0\big)\\ &\leq \mathbf{P}\big(X_{v_j^k}^{(\mu/2)} = 0\big)\end{aligned}$$

for some $1 \leq j \leq k$. The latter quantity is a standard quantity in the theory of random walks[1] and can be computed combinatorially using Stirling's formula (1.52), but we present here a Fourier-analytic approach. We can map v_j^k via a Freiman isomorphism to the identity 1 in a large cyclic group $\mathbf{Z}_N$, and use Lemma 7.11 to conclude

$$\mathbf{P}\big(X_{v_j^k}^{(\mu/2)} = 0\big) = \mathbf{E}_{\xi \in \mathbf{Z}_N} \left(1 - \frac{\mu}{2} + \frac{\mu}{2}\cos(2\pi\xi/N)\right)^k$$

and thus, on taking limits as $N \to \infty$,

$$\mathbf{P}\big(X_{v_j^k}^{(\mu/2)} = 0\big) = \int_0^1 \left(1 - \frac{\mu}{2} + \frac{\mu}{2}\cos(2\pi\xi)\right)^k \, d\xi.$$

Using (7.1), it suffices to bound $\int_0^1 \exp(-\Theta(k\mu^2\xi))d\xi$. It is easy to show that most of the weight of this integral is in the interval $(0, C/\sqrt{\mu k})$ for some large constant C. The claim follows. □

We remark that in the case $\mu = 1$, Corollary 7.4 gives the sharp bound

$$\mathbf{P}\big(X_{\mathbf{v}}^{(1)} = x\big) \leq \frac{\binom{k}{\lfloor k/2 \rfloor}}{2^k} = \Theta\left(\frac{1}{\sqrt{k}}\right)$$

thanks to Stirling's formula (1.52). This shows that the Fourier-analytic method can give bounds which are sharp up to absolute constants.

If the steps $v_1, \ldots, v_n$ are sufficiently "high-dimensional" one can do better than this $O(1/\sqrt{k})$ type bound; see Exercise 7.2.3.

Now let us give a deeper distributional inequality which relies in particular on the Cauchy–Davenport inequality (Theorem 5.4).

Lemma 7.14 (Halász relative concentration inequality) *[195] Let Z be either torsion-free or cyclic of odd prime order. Let $\mathbf{v}$ be a tuple in Z. Then for any*

---

[1] Indeed, a useful heuristic is to think of $X_{v^k}^{(\mu)}$ as behaving (up to constants) similarly to the uniform distribution on the progression $[-\sqrt{\mu k}, \sqrt{\mu k}] \cdot v$; note that this heuristic is supported by the Chernoff inequality.

$0 < \mu \leq \mu' \leq 1$ with $\mu \leq 1/4$, we have

$$\mathbf{P}\big(X_{\mathbf{v}}^{(\mu')} = x\big) \leq O\left(\sqrt{\frac{\mu}{\mu'}}\mathbf{P}\big(X_{\mathbf{v}}^{(\mu)} = 0\big)\right) + O\big(\mathbf{P}\big(X_{\mathbf{v}}^{(\mu)} = 0\big)^{\Theta(\mu'/\mu)}\big)$$

for all $x \in Z$.

Note that the domination inequality only gives $\mathbf{P}(X_{\mathbf{v}}^{(\mu')} = x) \leq \mathbf{P}(X_{\mathbf{v}}^{(\mu)} = 0)$. Thus Halász's inequality becomes superior when μ is significantly smaller than μ', in which case it asserts that $X_{\mathbf{v}}^{(\mu)}$ concentrates at the origin substantially more often than $X_{\mathbf{v}}^{(\mu')}$ does. For some further discussion and more quantitative versions of this inequality, see [195], [364], [365].

Proof Using the domination inequality we may assume that $\mu' \leq 1/2$ and $x = 0$. We may also take μ'/μ to be large. By Corollary 5.25 we may take $Z = \mathbf{Z}_p$ for some odd prime p. Introduce the functions $F, G : Z \to \mathbf{R}^+$ by

$$F(\xi) := \prod_{j=1}^{n}(1 - \mu' + \mu' \cos(2\pi\xi \cdot v_j)); \quad G(\xi) := \prod_{j=1}^{n}(1 - \mu + \mu \cos(2\pi\xi \cdot v_j));$$

then by Lemma 7.11 our task is to show that

$$\mathbf{E}_{\mathbf{Z}_p}(F) = O\left(\sqrt{\frac{\mu}{\mu'}}\mathbf{E}_{\mathbf{Z}_p}(G)\right) + O\big(\mathbf{E}_{\mathbf{Z}_p}(G)^{\Omega(\mu'/\mu)}\big).$$

Now let $0 < \alpha \leq 1$ be arbitrary. Observe from (7.1) that if $\xi \in \mathbf{Z}_p$ is such that $F(\xi) > \alpha$, then

$$\left(\sum_{j=1}^{n} \|\xi \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2\right)^{1/2} = O\left(\frac{\sqrt{\log\frac{1}{\alpha}}}{\sqrt{\mu'}}\right).$$

By the triangle inequality, we thus conclude that if $\xi_1, \ldots, \xi_m$ are arbitrary elements of the set $\{\xi \in \mathbf{Z}_p : F(\xi) \geq \alpha\}$, then

$$\left(\sum_{j=1}^{n} \|(\xi_1 + \cdots + \xi_m) \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2\right)^{1/2} = O\left(m\frac{\sqrt{\log\frac{1}{\alpha}}}{\sqrt{\mu'}}\right).$$

If we take m to be $\lfloor c\sqrt{\frac{\mu'}{\mu}}\rfloor$ for some small absolute constant $c > 0$, another application of (7.1) then gives

$$G(\xi_1 + \cdots + \xi_m) > \alpha.$$

In other words we have established the sum set inclusion

$$m\{\xi \in \mathbf{Z}_p : F(\xi) > \alpha\} \subseteq \{\xi \in \mathbf{Z}_p : G(\xi) > \alpha\}.$$

Applying the Cauchy–Davenport inequality repeatedly, we have[1]

$$\mathbf{P}_{\mathbf{Z}_p}(m\{\xi \in \mathbf{Z}_p : G(\xi) > \alpha\}) \geq \max(m\mathbf{P}_{\mathbf{Z}_p}(\{\xi \in \mathbf{Z}_p : F(\xi) > \alpha\}), 1).$$

If $\alpha \geq \mathbf{E}_{\mathbf{Z}_p}(G)$, then $\mathbf{P}_{\mathbf{Z}_p}(\{\xi \in \mathbf{Z}_p : G(\xi) > \alpha\}) < 1$ by Markov's inequality, and hence

$$\mathbf{P}_{\mathbf{Z}_p}(\{\xi \in \mathbf{Z}_p : F(\xi) > \alpha\}) \leq \frac{1}{m}\mathbf{P}_{\mathbf{Z}_p}(\{\xi \in \mathbf{Z}_p : G(\xi) > \alpha\}).$$

Integrating this in α, we conclude

$$\mathbf{E}_{\mathbf{Z}_p}(F\mathbf{I}(F \geq \mathbf{E}_{\mathbf{Z}_p}(G))) \leq \frac{1}{m}\mathbf{E}_{\mathbf{Z}_p}(G) = O\left(\sqrt{\frac{\mu}{\mu'}}\mathbf{E}_{\mathbf{Z}_p}(G)\right).$$

On the other hand, from (7.1) we have the pointwise bound

$$F(\xi) \leq G^{\Theta(\mu'/\mu)}(\xi)$$

and hence

$$\mathbf{E}_{\mathbf{Z}_p}(F\mathbf{I}(F < \mathbf{E}_{\mathbf{Z}_p}(G))) \leq \mathbf{E}_{\mathbf{Z}_p}(G)^{\Theta(\mu'/\mu)}.$$

Adding this to the preceding inequality, we obtain the claim. □

A modification of the above argument gives a more direct bound on $\mathbf{P}(X_{\mathbf{v}}^{(\mu)} = x)$.

Lemma 7.15 (Halász concentration inequality) *[167] Let Z be a cyclic group of prime odd order, and let* $\mathbf{v} = (v_1, \ldots, v_n)$ *be a tuple in Z with all the* v_j *non-zero. Then for any* $0 < \mu \leq 1$ *and* $x \in Z$ *we have*

$$\mathbf{P}\big(X_{\mathbf{v}}^{(\mu)} = x\big) \leq O\left(\frac{1}{\sqrt{\mu n}}\mathbf{P}_{\xi \in Z}\left(\sum_{j=1}^{n} \cos(\xi \cdot v_j) \geq \frac{n}{2}\right)\right) + \exp(-\Omega(\mu n)). \tag{7.2}$$

Proof Using the domination property we may take $\mu \leq 1/2$. By Lemma 7.11 and (7.1) we have

$$\mathbf{P}\big(X_{\mathbf{v}}^{(\mu)} = x\big) \leq \mathbf{E}_Z F \leq \mathbf{E}_{\xi \in Z} \exp\left(-\Theta\left(\mu \sum_{j=1}^{n} \|\xi \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2\right)\right).$$

---

[1] To be absolutely precise here, we should have written

$$\mathbf{P}_{\mathbf{Z}_p}(m\{\xi \in \mathbf{Z}_p : G(\xi) > \alpha\}) \geq \max(m\mathbf{P}_{\mathbf{Z}_p}(\{\xi \in \mathbf{Z}_p : F(\xi) > \alpha\}) - (m-1)/p, 1),$$

since Cauchy–Davenport inequality only implies $|A + B| \geq \min\{|A| + |B| - 1, p\}$, for any two subsets A, B of Z_p. However, the term $(m-1)/p$ is negligible as we can take p arbitrarily large.

We can subdivide the right-hand side based on the size of $(\sum_{j=1}^n \|\xi \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2)^{1/2}$, and bound the above expression by

$$O\left(\sum_{1 \le m \le c\mu n} \exp(-\Theta(m)) \mathbf{P}_{\xi \in Z}\left(\left(\sum_{j=1}^n \|\xi \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2\right)^{1/2} \le \sqrt{m/\mu}\right)\right) + \exp(-\Omega(c\mu n))$$

where $c > 0$ is a small absolute constant. Now observe that

$$\|\xi \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2 = \Theta(1 - \cos(2\pi\xi \cdot v_j)) \tag{7.3}$$

which in conjunction with Lemma 4.5 gives

$$\mathbf{E}_{\xi \in Z} \|\xi \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2 = \Theta(1).$$

By linearity of expectation we thus have

$$\mathbf{E}_{\xi \in Z} \sum_{j=1}^n \|\xi \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2 = \Theta(n);$$

in particular, we see that $\mathbf{P}_{\xi \in Z}(((\sum_{j=1}^n \|\xi \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2)^{1/2} \le c\sqrt{n}))$ is strictly less than one if c is small enough. Applying the Cauchy–Davenport inequality as in the preceding proof, we conclude

$$\begin{aligned} &\mathbf{P}_{\xi \in Z}\left(\left(\sum_{j=1}^n \|\xi \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2\right)^{1/2} \le \sqrt{m/\mu}\right) \\ &\quad \le O\left(\sqrt{\frac{m}{\mu n}}\right) \mathbf{P}_{\xi \in Z}\left(\left(\sum_{j=1}^n \|\xi \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2\right)^{1/2} \le c\sqrt{n}\right). \end{aligned}$$

Using (7.3) again, we conclude

$$\mathbf{P}_{\xi \in Z}\left(\left(\sum_{j=1}^n \|\xi \cdot v_j\|_{\mathbf{R}/\mathbf{Z}}^2\right)^{1/2} \le \sqrt{m/\mu}\right) \le O\left(\sqrt{\frac{m}{\mu n}}\right) \mathbf{P}_{\xi \in Z}\left(\sum_{j=1}^n \cos(\xi \cdot v_j) \ge \frac{n}{2}\right)$$

if c is sufficiently small. The claim then follows from the observation that

$$\sum_{1 \le m \le \sqrt{\mu n}} \exp(-\Theta(m)) \sqrt{\frac{m}{\mu n}} = O\left(\frac{1}{\sqrt{\mu n}}\right)$$

(the geometric decay of $\exp(-\Theta(m))$ being more than sufficient to counteract the polynomial growth of $\sqrt{m}$). □

This bound easily implies Corollary 7.13, and is in fact significantly stronger. For instance, we have

Corollary 7.16 *[167] Let* $0 < \mu \leq 1$*, and let* n *be sufficiently large depending on* μ*. Let* $\mathbf{v} = (v_1, \ldots, v_n)$ *be a tuple of positive integers. For each integer* $j > 0$*, let* m_j *denote the number of times* j *occurs in* $\mathbf{v}$*, thus* $m_j := \{1 \leq i \leq n : v_i = j\}$*. Then for any* $x \in Z$ *we have*

$$\mathbf{P}\big(X_{\mathbf{v}}^{(\mu)} = x\big) \leq O\left(\mu^{-1/2} n^{-5/2} \sum_{j>0} m_j^2\right).$$

In particular, if all the v_i *are distinct, then*

$$\mathbf{P}\big(X_{\mathbf{v}}^{(\mu)} = x\big) \leq O\big(\mu^{-1/2} n^{-3/2}\big).$$

We remark that in the $\mu = 1$ case, the second half of this Corollary was first established by combinatorial means in [310] (with the precise threshold given in [330]).

Proof We may use a Freiman isomorphism to place $v_1, \ldots, v_n$ inside $\mathbf{Z}_p$ for some very large prime p. A direct application of Parseval's theorem 4.2 gives

$$\mathbf{E}_{\xi \in \mathbf{Z}_p} \left| \sum_{j=1}^{n} \cos(\xi \cdot v_j) \right|^2 = O\left(\sum_{j>0} m_j^2\right)$$

and hence by Markov's inequality

$$\mathbf{P}_{\xi \in \mathbf{Z}_p} \left(\sum_{j=1}^{n} \cos(\xi \cdot v_j) \geq \frac{n}{2} \right) = O\left(\frac{1}{n^2} \sum_{j>0} m_j^2 \right).$$

The claim then follows from Lemma 7.15 (observing that $\exp(-\Theta(\mu n)) = O(\mu^{-1/2} n^{-5/2})$ when n is large). □

Exercises

7.2.1 Show that in the condition $\mu \leq \mu'/4$ in the domination inequality of Corollary 7.12, the constant 4 cannot be replaced by any smaller constant, even in the most important case $\mu = 1$.

7.2.2 If $\mathbf{v} = (v_1, \ldots, v_n)$ are a tuple of integers, show that

$$\mathbf{P}\big(X_{\mathbf{v}}^{(\mu)} = m\big) = \int_0^1 \cos(2\pi m \xi) \prod_{j=1}^{n} (1 - \mu + \mu \cos(2\pi v_j \xi))\, d\xi$$

for all integers m.

7.2.3 [167] Let $1 \leq k \leq n$ and $d \geq 1$, and let $\mathbf{v} = (v_1, \ldots, v_n)$, a tuple of vectors in $\mathbf{R}^d$, be "non-degenerate" in the sense that every proper subspace of $\mathbf{R}^d$ contains at most $n - k$ of the $v_1, \ldots, v_n$. Show that

$$\mathbf{P}\big(X_{\mathbf{v}}^{(\mu)} = x\big) = O_d\big((\mu k)^{-d/2}\big)$$

for every $0 < \mu \leq 1$ and $x \in \mathbf{R}^d$. (Hint: argue as Corollary 7.13, starting with an expression such as $\mathbf{P}(X^{(\mu/d)}_{\mathbf{v}^d} = 0)$ and applying Hölder's inequality suitably to arrive at a quantity such as $\mathbf{P}(X^{(\mu/d)}_{w_1^k w_2^k \dots w_d^k})$, where $w_1, \dots, w_d \in \mathbf{R}^d$ are linearly independent.) Give examples that show this bound is best possible up to the implicit constants in the $O_d()$ notation.

7.2.4 [364] With the notation and assumptions of Lemma 7.14, establish the following quantitative special case of the Halász inequality:

$$\mathbf{P}\big(X^{(1)}_{\mathbf{v}} = x\big) \leq \frac{1}{2}\mathbf{P}\big(X^{(1/16)}_{\mathbf{v}} = 0\big) + \mathbf{P}\big(X^{(1/16)}_{\mathbf{v}} = 0\big)^4.$$

7.2.5 Show that Lemma 7.14 can fail when Z is a non-cyclic finite group. In particular, if $Z = F_3^d$, show that $\mathbf{P}(X^{(\mu)}_{\mathbf{v}} = 0)$ can be comparable to $1/3^d$ for a large range of μ if the tuple $\mathbf{v}$ is chosen appropriately. This shows the pivotal role played by the Cauchy–Davenport inequality in the Halász argument.

7.2.6 Show that if the m_j are decreasing in j, then the right-hand side of Corollary 7.16 cannot be improved except for the implicit constant. (Hint: compute the variance of $X^{(\mu)}_{\mathbf{v}}$.)

7.2.7 Let $0 < \mu \leq 1$, and suppose n is sufficiently large depending on μ. Let $\mathbf{v} = (v_1, \dots, v_n)$ take values in an additive set S in $\mathbf{Z}_p$ for some odd prime p. Show that for any even integer $k \geq 2$ and $x \in \mathbf{Z}$ we have

$$\mathbf{P}\big(X^{(\mu)}_{\mathbf{v}} = x\big) \leq O_k\left(\mu^{-1/2} n^{-2k-\frac{1}{2}} \|S\|^{2k}_{\Lambda(2k)} \left(\sum_{j \in S} m_j^2\right)^k\right)$$

where m_j is the number of times j occurs in $\mathbf{v}$, and the $\Lambda(2k)$ constant is defined in Definition 4.26. In particular, if the v_j are all distinct, then

$$\mathbf{P}\big(X^{(\mu)}_{\mathbf{v}} = x\big) \leq O_k\big(\mu^{-1/2} n^{-(2k+1)/2} \|S\|^{2k}_{\Lambda(2k)}\big).$$

Thus $X^{(\mu)}_{\mathbf{v}}$ can only concentrate significantly when the $\Lambda(p)$ constants of the support of $\mathbf{v}$ are large.

7.2.8 [167] Let $0 < \mu \leq 1$, and let n be sufficiently large depending on μ. Let $v_1, \dots, v_n$ be non-zero integers, and let $k \geq 2$ be an even integer. Generalize Corollary 7.16 to show that for any $x \in \mathbf{Z}$ we have

$$\mathbf{P}\big(X^{(\mu)}_{\mathbf{v}} = x\big) \leq O_k\big(\mu^{-1/2} n^{-2k-\frac{1}{2}} R_k\big)$$

where R_k is the number of solutions to the equation

$$\epsilon_1 v_{i_1} + \cdots + \epsilon_{2k} v_{i_{2k}} = 0$$