

FURSTENBERG-POISSON BOUNDARY 101
LECTURE NOTES

OMER TAMUZ

CONTENTS

1. Random Walks on \mathbb{Z}^d and Pólya's Theorem	1
2. Random walks on groups	2
3. The tails of random walks on abelian groups: statement	4
4. Harmonic functions and the Furstenberg transform	4
5. The tails of random walks on abelian groups: proof	7
6. A non-trivial example	7
7. The Furstenberg-Poisson boundary	8
8. Lamplighters	10
9. Furstenberg entropy and the Kaimanovich-Vershik Theorem	12

1. RANDOM WALKS ON \mathbb{Z}^d AND PÓLYA'S THEOREM

Let $\{X_n\}_{n \in \mathbb{N}}$ be i.i.d. real random variables, each distributed uniformly on $\{-1, +1\}$. Let

$$Z_n = X_1 + X_2 + \cdots + X_n.$$

The process $\{Z_n\}_{n \in \mathbb{N}}$ is called the *simple random walk on \mathbb{Z}* . A classical question is whether the random walk is *recurrent*: does it return to the origin infinitely many times? Let R be the event that $Z_n = 0$ infinitely often.

Theorem 1.1 (Pólya). *For the simple random walk on \mathbb{Z} , $\mathbb{P}[R] = 1$.*

Some intuition behind this theorem is the following: the distribution of Z_n looks like a normal distribution with variance n . If we think of this distribution as the uniform distribution on $[-\sqrt{n}, \sqrt{n}]$, we see that $\mathbb{P}[Z_n = 0]$ is about $1/\sqrt{n}$. Hence the expected number of visits to the origin is about $\sum_n 1/\sqrt{n} = \infty$.

Now, by the Markov property of this random walk, the number of visits to the origin is distributed geometrically with parameter p , where p is the probability to return to the origin. Since the expectation is infinite then $p = 1$.

The same question can be asked about the simple random walk on \mathbb{Z}^2 , for which the X_n 's have the uniform distribution on $\{(-1, 0), (1, 0), (0, -1), (0, 1)\}$, and in general for \mathbb{Z}^d .

Theorem 1.2 (Pólya). *For the simple random walk on \mathbb{Z}^d , $\mathbb{P}[R] = 1$ if $d \leq 2$ and $\mathbb{P}[R] = 0$ if $d \geq 3$.*

The intuition given above for \mathbb{Z} is also helpful for understanding the general case.

The simple random walk on \mathbb{Z}^3 will eventually leave the origin and never come back. It follows that it will get further and further away: $\mathbb{P}[\lim_n |Z_n| = \infty] = 1$. One can wonder: will this random walk eventually “settle” on a single octant and never leave it? It turns out that the answer is no:

$$\mathbb{P}[Z_n \geq (0, 0, 0) \text{ for all } n \text{ large enough}] = 0.$$

While it is easy to see why this is true, we defer the proof to later, when we will, in fact, prove a much stronger statement.

2. RANDOM WALKS ON GROUPS

We now turn to formally define our probability space and measure, for the general case of random walks on groups. Let G be a countable discrete group, and let μ be a probability measure on G . We assume that the support of μ generates G as a semi-group. We note that this entire theory extends to the (very interesting) case that G is locally compact and second countable, and some power of μ is absolutely continuous with respect to the Haar measure.

Equip $G^{\mathbb{N}}$ with the product topology, and let Σ be the Borel sigma-algebra. Let

$$\begin{aligned} \varphi : G^{\mathbb{N}} &\longrightarrow G^{\mathbb{N}} \\ (h_1, h_2, h_3, \dots) &\longmapsto (h_1, h_1 h_2, h_1 h_2 h_3, \dots) \end{aligned}$$

Our probability space is $(G^{\mathbb{N}}, \Sigma, \mathbb{P})$, where $\mathbb{P} = \varphi_*(\mu^{\mathbb{N}})$. When $G = \mathbb{Z}$ and μ is the uniform distribution on $\{-1, +1\}$ then \mathbb{P} is precisely the distribution of (Z_1, Z_2, Z_3, \dots) , the locations of the simple random walk on \mathbb{Z} . We thus refer to μ as the step distribution and to $G^{\mathbb{N}}$ as the space of random walk paths. We will still, when convenient, use the random variables $\{X_n\}$ and $\{Z_n\}$, which are now formally given by

$$Z_n(g_1, g_2, \dots) = g_n.$$

and

$$X_n(g_1, g_2, \dots) = g_{n-1}^{-1} g_n$$

where g_0 is understood to equal e , the identity of G .

We let G act on $G^{\mathbb{N}}$ by the diagonal action:

$$g(g_1, g_2, \dots) = (gg_1, gg_2, \dots).$$

This action can be extended in the usual way to an action on probability measures on $G^{\mathbb{N}}$. In particular we will act on \mathbb{P} by

$$(g_*\mathbb{P})[A] = \mathbb{P}[g^{-1}A],$$

for any $A \in \Sigma$. The measure $g_*\mathbb{P}$ can be thought of as the measure of a random walk which starts from g rather than from the identity. It is easy to see that $g_*\mathbb{P} = \varphi_*^g(\mu^{\mathbb{N}})$ where

$$\begin{aligned} \varphi^g &: G^{\mathbb{N}} && \longrightarrow G^{\mathbb{N}} \\ &(h_1, h_2, h_3, \dots) && \longmapsto (gh_1, gh_1h_2, gh_1h_2h_3, \dots) \end{aligned}$$

Denote by e the identity of G . The set of recurrent paths is given by

$$R = \{(g_1, g_2, g_3, \dots) : g_n = e \text{ i.o.}\}.$$

The random walk is said to be recurrent if $\mathbb{P}[R] = 1$. Note that the indicator function of R has the property that it is independent of any prefix of the path: for any $n \in \mathbb{N}$, $(h_1, \dots, h_n) \in G^n$ and $(g_1, g_2, \dots) \in G^{\mathbb{N}}$ it holds that

$$\mathbb{1}_{\{A\}}(g_1, g_2, \dots, g_n, g_{n+1}, \dots) = \mathbb{1}_{\{A\}}(h_1, h_2, \dots, h_n, g_{n+1}, \dots).$$

This property of R makes it a *tail event*. Another definition of tail events is the following. Let $\mathcal{T}_n \subset \Sigma$ be the sigma-algebra of sets measurable in the coordinates g_{n+1}, g_{n+2}, \dots ; equivalently, \mathcal{T}_n is the pre-image of Σ under the map σ^n , where σ is the shift given by

$$\begin{aligned} \sigma &: G^{\mathbb{N}} && \longrightarrow G^{\mathbb{N}} \\ &(g_1, g_2, g_3, \dots) && \longmapsto (g_2, g_3, \dots) \end{aligned}$$

The *tail sigma-algebra* is given by $\mathcal{T} = \bigcap_n \mathcal{T}_n$. A tail event is simply an element of \mathcal{T} . Note that \mathcal{T} is G -invariant; if $T \in \mathcal{T}$ then $gT \in \mathcal{T}$.

An example of a tail event, for $G = \mathbb{Z}^d$, is the “settling on the octant” event

$$P = \{(g_1, g_2, \dots) : g_n \geq (0, 0, 0) \text{ for all } n \text{ large enough}\}.$$

Another important sigma-algebra is the *shift-invariant* sigma-algebra \mathcal{I} . This is the sub-sigma-algebra of sets A such that $(g_1, g_2, \dots) \in A$ if and only if $\sigma(g_1, g_2, \dots) \in A$. We will not prove the following claim.

Theorem 2.1. *The shift-invariant and tail sigma-algebras are equal, \mathbb{P} -mod 0. I.e., $L^\infty(G^{\mathbb{N}}, \mathcal{T}, \mathbb{P}) = L^\infty(G^{\mathbb{N}}, \mathcal{I}, \mathbb{P})$.*

Clearly, the shift-invariant sigma-algebra is G -invariant; if $A \in \Sigma$ is shift-invariant then gA is also shift-invariant. Another important property of the shift-invariant sigma-algebra is the following.

Proposition 2.2. *For all $g, k \in G$, shift-invariant S and $n \in \mathbb{N}$ such that $g_*\mathbb{P}[Z_n = k] > 0$ it holds that $g_*\mathbb{P}[S|Z_n = k] = k_*\mathbb{P}[S]$.*

That is, the probability of a shift-invariant event S given that the random walk was at k at time n is the same as the probability of this event for a random walk starting from k . We will not prove this claim.

3. THE TAILS OF RANDOM WALKS ON ABELIAN GROUPS: STATEMENT

Recall that we wanted to prove that $\mathbb{P}[P] = 0$. Clearly, by the symmetry of the symmetric random walk, it is impossible that $\mathbb{P}[P] = 1$. Hence it will suffice to show that $\mathbb{P}[P] \in \{0, 1\}$. We prove the following theorem.

Theorem 3.1. *Let G be an abelian group. Then for every tail event $T \in \mathcal{T}$ it holds that $\mathbb{P}[T] \in \{0, 1\}$.*

That is, all tail events are trivial.

4. HARMONIC FUNCTIONS AND THE FURSTENBERG TRANSFORM

To prove Theorem 3.1 we will need to study *harmonic functions*. We say that $f: G \rightarrow \mathbb{R}$ is μ -harmonic (or just harmonic) if for all $g \in G$

$$f(g) = \sum_{k \in G} f(gk)\mu(k).$$

That is, the average of the values of f around g is equal to the value of f at g , and averages are taken using μ .

We first show how a tail event can be used to define a harmonic function. Let T be a tail event. Define $f: G \rightarrow \mathbb{R}$ by

$$f(g) = g_*\mathbb{P}[T].$$

That is, $f(g)$ is the probability that a random walk starting from g will be in T .

Claim 4.1. *f is harmonic.*

Proof. We condition on X_1 , the first step of the random walk:

$$\begin{aligned} f(g) &= g_*\mathbb{P}[T] \\ &= \sum_{k \in G} g_*\mathbb{P}[T|X_1 = k]\mu(k) \\ &= \sum_{k \in G} g_*\mathbb{P}[T|Z_1 = gk]\mu(k) \end{aligned}$$

by Theorem 2.1 we can assume w.l.o.g. that T is shift-invariant. It therefore follows from Proposition 2.2 that

$$\begin{aligned} f(g) &= \sum_{k \in G} (gk)_* \mathbb{P}[T] \mu(k) \\ &= \sum_{k \in G} f(gk) \mu(k). \end{aligned}$$

□

Note that these functions are bounded in $[0, 1]$, by definition. More generally, given a bounded real random variable Y that is tail measurable, the function

$$f(g) = \mathbb{E} [g^{-1}Y] = \int_{G^{\mathbb{N}}} Y(gg_1, gg_2, \dots) d\mathbb{P}(g_1, g_2, \dots).$$

is bounded and harmonic, by the same argument.

We have thus mapped each bounded tail random variable to a bounded harmonic function. We will now see how to go in the opposite direction.

Given a harmonic function $f: G \rightarrow \mathbb{R}$, define the random variables $\{M_n\}_{n \in \mathbb{N}}$ by

$$M_n = f(Z_n).$$

It is easy to verify that M_n is a martingale with respect to the filtration $\{\mathcal{F}_n\}_{n \in \mathbb{N}}$ where \mathcal{F}_n is the sigma-algebra of sets measurable in the coordinates g_1, \dots, g_n :

$$\begin{aligned} \mathbb{E} [M_{n+1} | Z_1, \dots, Z_n] &= \mathbb{E} [f(Z_{n+1}) | Z_n] \\ &= \sum_{k \in G} \mathbb{E} [f(Z_{n+1}) | Z_{n+1} = Z_n k] \mu(k) \\ &= \sum_{k \in G} f(Z_n k) \mu(k) \\ &= f(Z_n) \\ &= M_n. \end{aligned}$$

Here the first equality follows from the Markov property of the random walk, the second from conditioning on X_{n+1} , the third from the definition of conditional expectation and the fourth from the harmonicity of f .

If f is also bounded then M_n is a bounded martingale, and therefore almost surely converges. Hence we can define

$$M = \lim_n M_n = \lim_n f(Z_n),$$

and M is clearly \mathcal{T} -measurable. Hence we have shown how every bounded harmonic function gives rise to a bounded tail random variable, and vice versa. Indeed, let

$$\begin{aligned} \Phi &: H^\infty(G, \mu) &\longrightarrow & L^\infty(G^\mathbb{N}, \mathcal{T}, \mathbb{P}) \\ &f &\longmapsto & \lim_n f(Z_n) \end{aligned} .$$

be a map from the bounded μ -harmonic function to the bounded tail random variables. This transform is known as the *Furstenberg transform*, and the next claim states that its inverse is one-to-one. We leave to the reader to show that the inverse is onto.

Theorem 4.2. *Let Y be a bounded tail random variable. Let $f(g) = \mathbb{E}[g^{-1}Y]$ and let $M = \lim_n f(Z_n)$. Then M and Y are \mathbb{P} -a.s. equal.*

Proof. We prove for the case that Y is the indicator of a tail event T . The general proof is identical, but requires more cumbersome notation.

$$\begin{aligned} M &= \lim_n f(Z_n) \\ &= \lim_n \mathbb{P}[Z_n^{-1}T] \\ &= \lim_n Z_{n*} \mathbb{P}[T] \\ &= \lim_n \mathbb{P}[T|Z_n]. \end{aligned}$$

Here the last equality uses the shift-invariance of T and Proposition 2.2. By the Markov property of the random walk and the fact that T is a tail event,

$$= \lim_n \mathbb{P}[T|Z_1, \dots, Z_n].$$

But T is a function of (Z_1, Z_2, \dots) , and so this converges to the indicator of T . \square

In light of this transform, we can hope to study the tail sigma-algebra by studying the bounded harmonic functions on G . In the case of abelian groups this turns out to be a simple endeavor.

Theorem 4.3. *Let G be abelian and let μ be any probability measure on G . Then every bounded μ -harmonic function is constant.*

Proof. Let $H = H^{[0,1]}(G, \mu)$ be the set of μ -harmonic functions with range in $[0, 1]$. We note that harmonicity is invariant to multiplication by a constant and addition, and so if we show that every $h \in H$ is constant then we have proved our theorem.

We state three properties of H that are easy to verify.

- (1) H is invariant to the g action; for any $f \in H$ and $g \in G$, the function $f^g: G \rightarrow \mathbb{R}$ given by $[f^g](k) = f(g^{-1}k)$ is also in H .

- (2) H is compact in the topology of pointwise convergence.
- (3) H is convex.

As a convex compact space, H is the closed convex hull of its extremal points; this is the Krein-Milman Theorem. Let $f \in H$ be an extremal point. Then, since f is harmonic,

$$f(g) = \sum_{k \in G} f(gk)\mu(k).$$

Since G is abelian then

$$= \sum_{k \in G} f(kg)\mu(k) = \sum_{k \in G} f^{k^{-1}}(g)\mu(k).$$

By the first property of H each $f^{k^{-1}}$ is also in H , and thus we have written f as a convex combination of functions in H . But f is extremal, and so $f = f^{k^{-1}}$ for all k in the support of μ . But μ is generating, and so f is invariant to the G -action, and hence constant. \square

5. THE TAILS OF RANDOM WALKS ON ABELIAN GROUPS: PROOF

We have thus shown that every bounded harmonic function on G is trivial. It will now quickly follow that every tail event is trivial.

Proof of 3.1. Let T be a tail event and let $f = \Phi^{-1}(\mathbb{1}_{\{T\}})$. Since f is constant (Theorem 4.3) then $M = \lim_n f(Z_n)$ is constant. But M is the indicator of T (Theorem 4.2) and so either $\mathbb{P}[T] = 0$ or $\mathbb{P}[T] = 1$. \square

6. A NON-TRIVIAL EXAMPLE

Let F_2 be the free group generated by $S = \{a, b, a^{-1}, b^{-1}\}$, and let μ be the uniform distribution over S . This random walk is the simple random walk on the Cayley graph of F_2 , which is the four regular tree. We will state without proof that this random walk is transient.

It follows that eventually the random walk will leave the identity into one of the four branches and never return to the identity. Since the graph is a tree, the random walk will always remain in that same branch. By the symmetry of the random walk it follows that each of the four branches has equal probability of being the branch into which the random walk is absorbed, and so, if we define

$$A = \{(g_1, g_2, \dots) : g_n \text{ starts with } a \text{ for all } n \text{ large enough}\}$$

then $\mathbb{P}[A] = 1/4$. Now A is clearly a tail event, and so we have found a non-trivial tail event. Accordingly,

$$f(g) = g_*\mathbb{P}[A]$$

is non-constant bounded harmonic function.

7. THE FURSTENBERG-POISSON BOUNDARY

Consider again the simple random walk on \mathbb{Z}^2 . Let $F \subset \mathbb{Z}^2$ be a finite connected component of the Cayley graph of \mathbb{Z}^2 that contains the origin, and let ∂F be the elements of F that, in the Cayley graph, have edges to the complement of F . We modify the random walk so that, once it hits ∂F , it stops and stays in the same location in all subsequent time periods. Denote by

$$Z_\infty := \lim_n Z_n \in \partial F$$

the (random) location at which the random walk hits the boundary ∂F and stops. Note that Z_∞ is a tail random variable.

Let $f: F \rightarrow \mathbb{R}$ be harmonic on $F \setminus \partial F$. Then, by an argument almost identical to the one used above,

$$M_n = f(Z_n)$$

is a bounded martingale, with limit

$$M = \lim_n M_n = f(Z_\infty).$$

Now, by the martingale property,

$$M_n = \mathbb{E}[M|Z_n] = \mathbb{E}[f(Z_\infty)|Z_n].$$

Hence, by the definition of M_n ,

$$f(Z_n) = \mathbb{E}[f(Z_\infty)|Z_n]$$

and in particular for any fixed $g \in F$,

$$f(g) = \mathbb{E}[f(Z_\infty)|Z_n = g].$$

In words, $f(g)$ is the expectation of f at Z_∞ , for a random walk starting at g . Since Z_∞ takes values on ∂F , it follows that f is uniquely determined by its values on the boundary ∂F .

The Furstenberg-Poisson boundary serves the role of ∂F , for the case that F is the entire group G and f is a bounded harmonic function on G .

To define the Furstenberg-Poisson boundary we will use the Mackey Point Realization Theorem. We rephrase this theorem as follows.

Theorem 7.1 (Mackey). *Let (Ω, Σ, η) be a standard probability space. Let a locally compact, second countable group G act on Ω by measurable transformations. Let Λ be a sub-sigma-algebra of Σ that is G -invariant, and such that for any $A \in \Lambda$ and $g \in G$, $\eta(A) = 0$ iff $\eta(gA) = 0$. Then there exists a standard probability G -space (Ω', Σ', ν) and a Λ -measurable, G -equivariant map $\pi: \Omega \rightarrow \Omega'$ such that $\pi_*\eta = \nu$ and such that the image of $L^\infty(\Omega', \Sigma', \nu)$ under π^{-1} is $L^\infty(\Omega, \Lambda, \eta)$.*

Applying this theorem to the tail sigma-algebra \mathcal{T} yields the Furstenberg-Poisson boundary.

Theorem 7.2. *Let G be a discrete group and let μ be a probability measure on G . Then there exists a standard probability G space (Π, ν) and a \mathcal{T} -measurable G -map $B: G^{\mathbb{N}} \rightarrow \Pi$ such that $B_*\mathbb{P} = \nu$ and such that*

$$B^{-1}L^\infty(\Pi, \nu) = L^\infty(G^{\mathbb{N}}, \mathcal{T}, \mathbb{P}).$$

The space (Π, ν) is called the *Furstenberg-Poisson boundary* and B is called the *boundary map*.

To prove this theorem we need simply apply Mackey's Theorem; the only complication is that we need to check that the G -action on \mathcal{T} preserves null sets (i.e., that for every $T \in \mathcal{T}$ and $g \in G$ it holds that $\mathbb{P}[T] = 0$ iff $\mathbb{P}[gT] = 0$). We leave this as an exercise to the reader; it is a consequence of Proposition 2.2.

Using B we can define a random variable

$$Z_\infty = B(Z_1, Z_2, \dots).$$

Since B is \mathcal{T} -measurable, Z_∞ is a tail random variable. It takes values in Π , and we can think of it as the point in which the random walk hits the boundary. Its distribution is ν , which we therefore call the hitting (or harmonic) measure. Since B is G -equivariant we get that $B_*g_*\mathbb{P} = g_*\nu$. Hence we can think of $g_*\nu$ as the hitting measure for a random walk that starts at g .

Given $f \in H^\infty(G, \mu)$, recall that $\Phi(f) = \lim_n f(Z_n)$ is in $L^\infty(G^{\mathbb{N}}, \mathcal{T}, \mathbb{P})$. Since we can identify this space with $L^\infty(B, \nu)$ using B , we can define the Furstenberg transform

$$\Psi : \begin{array}{ccc} H^\infty(G, \mu) & \longrightarrow & L^\infty(B, \nu) \\ f & \longmapsto & B \circ \Phi(f) \end{array} .$$

The inverse transform $\Psi^{-1} \in H^\infty(G, \mu)$ is given by

$$[\Psi^{-1}(Y)](g) = g_*\nu(Y),$$

for $Y \in L^\infty(B, \nu)$. Note that ν is G -quasi-invariant, since the G -action on \mathcal{T} preserves \mathbb{P} -null sets. Note also that if ν is G -invariant then every bounded harmonic function on G is constant. In this case we know that the tail sigma-algebra is trivial, and so the sigma-algebra of B has to be trivial. Hence B has to be (mod 0) a map to a single point, making the boundary trivial. We state this formally:

Proposition 7.3. *If ν is G -invariant then the Furstenberg-Poisson boundary is trivial.*

Given that ν is not G -invariant, it might be interesting to measure how non-invariant it is. To this end, recall that the *Kullback-Leibler divergence* or *relative entropy* between two measures β and γ on a measurable space (Ω, Σ) , where β is absolutely continuous with respect to γ , is

$$D_{KL}(\beta||\gamma) = \int_{\Omega} -\log \frac{d\gamma}{d\beta}(\omega) d\beta(\omega).$$

The relative entropy is non-negative, and is zero only if the two measures are equal.

To measure how non-invariant ν is, we will calculate how much it is deformed by G *on average*, where this average will be taken using μ :

$$\begin{aligned} h_{\mu}(B, \nu) &= \sum_{g \in G} D_{KL}(g_*\nu||\nu)\mu(g) \\ &= \sum_{g \in G} \int_B -\log \frac{d\nu}{dg_*\nu}(b) dg_*\nu(b) \mu(g). \end{aligned}$$

This quantity is called the *Furstenberg entropy* of (B, ν) , and, as we will see below, plays an important role in the study of the Furstenberg-Poisson boundary.

8. LAMPLIGHTERS

Fix $d \geq 1$, and let the set of “lamp configurations” L_d be the direct sum $\bigoplus_{z \in \mathbb{Z}^d} \mathbb{Z}/2\mathbb{Z}$. This is the set of finite subsets of \mathbb{Z}^d , equipped with the operation of symmetric difference.

\mathbb{Z}^d acts on L_d in the obvious way, and we can define the *lamplighter group* $\Lambda_d = L_d \rtimes \mathbb{Z}^d$. L_d is a normal subgroup of Λ_d , and we denote by $\text{pr}_2: \Lambda_d \rightarrow \mathbb{Z}^d$ the homomorphism which has L_d as its kernel. Let $\text{pr}_1: \Lambda_d \rightarrow L_d$ be the projection on the first coordinate.

We think of elements of L_d as functions $\ell: \mathbb{Z}^d \rightarrow \mathbb{Z}/2\mathbb{Z}$ with finite support, and let s be the function that is zero everywhere but at the origin. Then it is easy to check that s , together with the standard generating set of \mathbb{Z}^d , is a generating set for Λ_d . We denote this generating set by S_d .

Let μ_d be the uniform distribution on S_d , and consider the μ_d random walk on Λ_d . We can project this walk using pr_2 to a simple (lazy) random walk on \mathbb{Z}^d . This random walk will be recurrent for $d = 1, 2$ and transient for $d \geq 3$.

Let the set of random walk paths A_0 be those paths in which the lamp at the origin is eventually on:

$$A_0 = \{(\ell_n, z_n)_{n \in \mathbb{N}} : \ell_n(0) = 1 \text{ for all } n \text{ large enough}\}.$$

This set is clearly a tail event. When the projected random walk is recurrent then it has probability 0, since the lamp at the origin will change states infinitely often with probability one. However, when the random walk is transient, then the random walk will only visit the origin a finite number of times, and with non-trivial probability the lamp will be left on during the last visit.

We have thus proved the following claim.

Claim 8.1. *The Furstenberg-Poisson boundary of (Λ_d, μ_d) is non-trivial when $d \geq 3$.*

When the random walk is recurrent then the state of every lamp eventually stabilizes. That is, if we embed $\bigoplus_{z \in \mathbb{Z}^d} \mathbb{Z}/2\mathbb{Z}$ into $\Pi_d = \prod_{z \in \mathbb{Z}^d} \mathbb{Z}/2\mathbb{Z}$ (equipped with the profinite topology), then for almost every random walk path $(\ell_n, z_n)_{n \in \mathbb{N}}$ we have that the limit $\lim_n \ell_n$ exists. Let $\varphi: \Lambda_d^{\mathbb{N}} \rightarrow \Pi_d$ be the map that assigns to each random walk path $(\ell_n, z_n)_{n \in \mathbb{N}}$ the limit $\lim_n \ell_n$.

Hence a natural candidate for the Poisson boundary is the space of eventual lamp Π_d , equipped with the measure $\nu_d = \varphi_* \mathbb{P}$.

Theorem 8.2 (Lyons and Peres, 2014). *For $d \geq 3$, the Poisson boundary of (Λ_d, μ_d) can be identified with (Π_d, ν_d) .*

This was earlier shown by Erschler, for $d \geq 5$.

While these tail events are trivial for the cases $d = 1, 2$, it is not immediately obvious that there do not exist other, non-trivial tail events. However, it turns out that this is indeed the case.

Theorem 8.3 (Kaimanovich and Vershik). *For $d = 1, 2$, the Poisson boundary of (Λ_d, μ_d) is trivial.*

Proof. View L_d as a subgroup of Λ_d ; that is, we identify $(f, 0) \in \Lambda_d$ with $f \in L_d$. Since the projected random walk is recurrent, the μ_d random walk visits L_d infinitely often. Let N_1, N_2, \dots be the times of these visits, and let

$$Y_k = \text{pr}_1(Z_{N_k})$$

be the element of L_d that is visited at the k^{th} return to L_d . It is easy to see that the distribution of $Y_n^{-1}Y_{n+1}$ is independent of n , and so (Y_1, Y_2, \dots) is a random walk on L_d . Denote by $\bar{\mu}_d$ the distribution of $Y_n^{-1}Y_{n+1}$.

Fix $f \in H^\infty(\Lambda_d, \mu_d)$, and let \bar{f} be its restriction to L_d . Note that $\bar{f} \in H^\infty(L_d, \bar{\mu}_d)$, since $\bar{f}(Y_k) = f(Z_{N_k})$, N_k is a stopping time and so $\bar{f}(Y_k)$ is a martingale.

Now, since L_d is abelian, it follows that \bar{f} is constant, say C . Hence

$$f(g) = \mathbb{E} \left[\lim_n f(Z_n) \middle| Z_k = g \right] = \mathbb{E} \left[\lim_k \bar{f}(Y_{N_k}) \middle| Z_k = g \right] = C$$

and so f is also constant. Thus $H^\infty(\Lambda_d, \mu_d)$ is trivial. \square

In the proof above we have (almost) proved the following general claim.

Theorem 8.4. *Let H be a recurrent subgroup of the μ random walk on G . Then there exists a probability measure $\bar{\mu}$ on H such that $H^\infty(G, \mu) \cong H^\infty(H, \bar{\mu})$.*

9. FURSTENBERG ENTROPY AND THE KAIMANOVICH-VERSHIK THEOREM

We saw before that the boundary of abelian groups is always trivial. We also saw that the boundary of the simple random walk on F_2 is not trivial. It is natural to ask for which pairs (G, μ) is the boundary trivial? This turns out to be a deep and important question that is not yet completely resolved. The most important tool at our disposal is Furstenberg entropy and the Kaimanovich-Vershik Theorem.

We recall some basic information theoretical notions. Let X and Y be countably supported random variables, denote $p_x = \mathbb{P}[X = x]$, $p_y = \mathbb{P}[Y = y]$, and $p_{x|y} = \mathbb{P}[X = x|Y = y]$. The *Shannon entropy* of X is given by

$$H(X) = \sum_x -\log(p_x) \cdot p_x.$$

An important property of entropy that we will use is the following: if φ is a function of the support of X , then

$$H(\varphi(X)) \leq H(X)$$

with equality if and only if φ is one-to-one.

The *conditional entropy* of X conditioned on Y is

$$H(X|Y) = \sum_y \left(\sum_x -\log(p_{x|y}) \cdot p_{x|y} \right) p_y.$$

The mutual information between X and Y is

$$I(X; Y) = H(X) - H(X|Y).$$

It is easy to see that $I(X; Y) = I(Y; X)$ and that $I(X; Y) \leq H(X)$, with equality iff $X = Y$. Intuitively, $I(X; Y)$ is a measure of how much information Y contains regarding X and vice versa.

An important property of mutual information is related to Markov chains. Let (Y_1, Y_2, \dots) be a Markov chain. Then for any $n \in \mathbb{N}$

$$I(Y_1; Y_n, Y_{n+1}, \dots) = I(Y_1; Y_n).$$

Intuitively, the information on Y_1 contained in (Y_n, Y_{n+1}, \dots) is the same as that contained in just Y_n .

An equivalent definition of mutual information is the following. Let X and Y have respective distributions β and γ , and let β_y be the distribution of X conditioned on $Y = y$. Then it is easy to verify that

$$I(X; Y) = \int D_{KL}(\beta_y || \beta) d\gamma(y).$$

The advantage of this form is that it is also well defined when X is not countably supported, and has all the properties we observed before. In particular, it is straightforward to check that when Y is countably supported but X is not then still $I(X; Y) \leq H(Y)$.

Recall that $Z_\infty = B(Z_1, Z_2, \dots)$ is the ‘‘hitting point’’ of the random walk on the Poisson boundary. A natural question is the following: how much information does Z_∞ contain on Z_1 , the first step of the random walk? In particular if this is non-zero then Z_∞ will be non-trivial.

Recall that $g_*\nu$ is the distribution of Z_∞ for a random walk starting at g . By reasoning similar to one used above, it is also the distribution of Z_∞ conditioned on $Z_1 = g$. Hence, by the relative entropy characterization of mutual information,

$$\begin{aligned} I(Z_\infty; Z_1) &= \sum_{g \in G} D_{KL}(g_*\nu || \nu) \mathbb{P}[Z_1 = g] \\ &= \sum_{g \in G} \int_B -\log \frac{d\nu}{dg_*\nu}(b) dg_*\nu(b) \mu(g) \\ &= h_\mu(B, \nu), \end{aligned}$$

the *Furstenberg entropy* of the boundary (B, ν) . Since relative entropy is zero only for two equal measures, and since the support of μ generates G , we have that $h_\mu(B, \nu) = 0$ iff ν is G -invariant. Hence $h_\mu(B, \nu) = 0$ iff the boundary is trivial.

Now, by a continuity argument,

$$I(Z_1; Z_\infty) = \lim_n I(Z_1; Z_n, Z_{n+1}, Z_{n+2}, \dots).$$

Since $\{Z_n\}$ is a Markov process then

$$I(Z_1; Z_\infty) = \lim_n I(Z_1; Z_n).$$

We have thus reduced this calculation to calculating the mutual information of two countably supported random variables. Now,

$$I(Z_1; Z_n) = H(Z_n) - H(Z_n|Z_1) = H(Z_n) - \sum_{g \in G} H(gX_2X_3 \cdots X_n)\mu(g).$$

Note that

$$H(gX_2 \cdots X_n) = H(X_2 \cdots X_n) = H(X_1 \cdots X_{n-1}) = H(Z_{n-1}),$$

and therefore we have shown that

$$I(Z_1; Z_n) = \lim_m H(Z_n) - H(Z_{n-1}).$$

Now, $H(Z_n)$ is a subadditive sequence:

$$\begin{aligned} H(Z_{n+m}) &= H(X_1 \cdots X_n \cdot X_{n+1} \cdots X_m) \\ &\leq H(X_1 \cdots X_n) + H(X_{n+1} \cdots X_m) \\ &= H(Z_n) + H(Z_m). \end{aligned}$$

Hence, assuming these quantities are finite,

$$\lim_m H(Z_n) - H(Z_{n-1}) = \lim_n \frac{1}{n} H(Z_n).$$

We denote $h(\mu) = \lim_n \frac{1}{n} H(Z_n)$ and call it the *random walk entropy* or *Avez entropy*. We have thus proved the following theorem.

Theorem 9.1 (Kaimanovich and Vershik). *If*

$$H(\mu) = \sum_g -\log(\mu(g)) \cdot \mu(g)$$

is finite then

$$h_\mu(B, \nu) = h(\mu).$$

In particular the Furstenberg-Poisson boundary is trivial iff the random walk entropy vanishes.