Ergod. Th. & Dynam. Sys. (2015), **35**, 359–373 doi:10.1017/etds.2013.56
© *Cambridge University Press, 2013*

Unimodular random trees

ITAI BENJAMINI†, RUSSELL LYONS‡ and ODED SCHRAMM§

† *Mathematics Department, The Weizmann Institute of Science, Rehovot 76100, Israel*
(e-mail: Itai.Benjamini@weizmann.ac.il)
‡ *Department of Mathematics, 831 E 3rd St, Indiana University, Bloomington, IN 47405-7106, USA*
(e-mail: rdlyons@indiana.edu)

(*Received* 16 *July* 2012 *and accepted in revised form* 20 *May* 2013)

Abstract. We consider unimodular random rooted trees (URTs) and invariant forests in Cayley graphs. We show that URTs of bounded degree are the same as the law of the component of the root in an invariant percolation on a regular tree. We use this to give a new proof that URTs are sofic, a result of Elek. We show that ends of invariant forests in the hyperbolic plane converge to ideal boundary points. We also note that uniform integrability of the degree distribution of a family of finite graphs implies tightness of that family for local convergence, also known as random weak convergence.

1. *Introduction*

The theory of unimodular random rooted networks (URNs) is an outgrowth mainly of two lines of investigation: one is concerned with asymptotics of finite networks, while the other involves group-invariant stochastic processes on infinite Cayley graphs, especially percolation. An important motivation also arises from the class of sofic groups. Parallels with the theory of limits of dense graphs now spur further investigations (see [**20**] for this).

We give full definitions in §2, but here we recount intuitively some of the above motivations. One way to look at a large finite network (which is a labeled graph) is to look at a large neighborhood around a random uniformly chosen vertex. Often such neighborhood statistics capture quantities of interest and their asymptotics. Thus, one is led to take limits of such statistics and thereby define a probability measure on infinite rooted graphs, where the neighborhood of the root has the statistics that arise as the limit statistics of the finite networks. Such a limit of a sequence of finite networks is called the random weak limit, the local (weak) limit, the distributional limit, or the Benjamini–Schramm limit of the sequence. All such limit measures have a property known as unimodularity; it is not known whether all unimodular measures are limits of finite networks. Those that are such limits are called sofic.

§ Deceased.

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

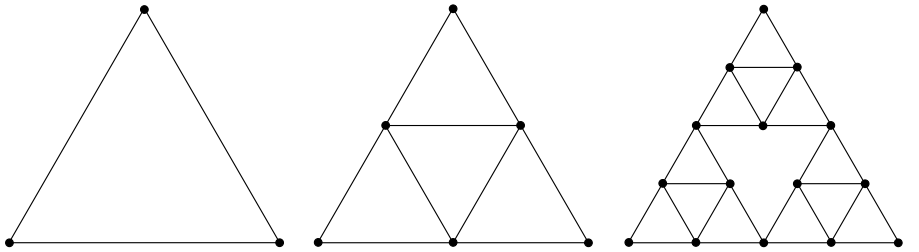


FIGURE 1. First three graphs whose random weak limit is the infinite Sierpiński gasket.

Intuitively, a probability measure on rooted networks is unimodular if and only if its root is chosen 'uniformly' from among all its vertices. This, of course, only makes sense for finite graphs. It is formalized for networks on infinite graphs by requiring a sort of conservation property known as the mass-transport principle.

Unimodularity is an extremely powerful property, especially for studying percolation on infinite graphs. In the present context, for example, the component of the identity in a group-invariant percolation on a Cayley graph has a unimodular law as a random graph rooted at the identity.

Consider the following example of a random weak limit of finite graphs: let $\mathbb{T}_3$ be a 3-regular tree. Let G_n be the ball of radius n in $\mathbb{T}_3$ about any point. Since most points in G_n are near the leaves of G_n, the random weak limit of $\langle G_n;\ n \geq 1\rangle$ is not $\mathbb{T}_3$ but the following probability measure, μ. Let T_n be disjoint binary trees of depth n for $n \geq 0$. Modify T_n by adding a new vertex x_n adjacent to the root of T_n. Also consider an isolated vertex x_{-1}. Now add an edge between x_n and x_{n+1} for each $n \geq -1$. The resulting tree, T, was called the *canopy tree* by Aizenman and Warzel [**1**]. The graph T rooted at x_n is denoted by (T, x_n). We now define μ by letting $\mu(T, x_n) := 2^{-n-2}$ $(n \geq -1)$. Thus, μ is supported on a single tree, which is a proper subtree of $\mathbb{T}_3$. In fact, as we show, μ can be obtained as the component of the root in an automorphism-invariant percolation on $\mathbb{T}_3$. Indeed, one of our main theorems is that every URN that is supported by trees of bounded degree can be obtained as the component of the root in an invariant percolation on a regular tree.

An interesting contrast is provided by other URNs. For example, consider the infinite discrete Sierpiński gaskets characterized by Teplyaev [**24**] (see Lemma 2.3 there). These are obtained as the random weak limit of the graphs that are the natural boundaries of the nth-stage construction of the usual Sierpiński gasket as $n \to \infty$ (see Figure 1). In this case, the limit measure μ has an uncountable support, although all the graphs in the support are still subgraphs of the triangular lattice in the plane. Yet, in this case, there is no invariant percolation on the triangular lattice such that the component of the root has law μ, since the subgraphs in the support of μ have density 0 and, by their topology, only one component in any percolation on the triangular lattice can be a Sierpiński gasket (except for degenerate ones that altogether have μ-measure 0).

We call a URN that is supported by trees a URT. We shall use our result that every URT of bounded degree can be obtained as the law of the component of the root in an invariant

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

labeled percolation on a regular tree to give a new proof that URTs are sofic. This was first shown (in a special case) by Elek [**14**], which answered Bollobás and Riordan [**11**, Question 3.3]. It was then extended from graphs to networks by Elek and Lippner [**15**]. Although not needed for any of these results, we give a sufficient condition for a collection of finite graphs to have a convergent subsequence, namely, that their degree distributions be uniformly integrable.

We remark that our theorem showing that every URT of bounded degree can be obtained via invariant percolation on a regular tree has a counterpart in the other direction: that is, rather than put a URT into a regular tree, one can put a regular tree (or forest) on a URT. More precisely, Hjorth [**18**] proved that every treeable (probability-measure-preserving) equivalence relation of cost at least 2 can be generated by a free action of a free group $\mathbb{F}_2$ on 2 generators. If, instead, the cost is assumed only to be larger than 1, then there is a subrelation that is generated by a free action of $\mathbb{F}_2$: see [**16**, Proposition 14]. In the remaining case where the treeable equivalence relation has cost 1, the equivalence relation is amenable, whence a theorem of Connes *et al* [**13**] shows that it is generated by a free action of $\mathbb{Z}$. A URT is essentially the same as a treeable equivalence relation. We do not use any of these notions here, however, so we leave these terms undefined.

Finally, we turn from results about general URTs to somewhat specific URTs. Consider a discrete forest in the hyperbolic plane. A simple infinite path in the forest is called a ray. Clearly, a ray can have fairly arbitrary limiting behavior; in particular, though it must tend to the ideal boundary because the forest is discrete, the ray need not converge to any ideal boundary point. However, we show that with the condition solely that the forest is random with a law that is invariant under hyperbolic isometries, almost surely (a.s.) all its rays converge to ideal boundary points. We do not know whether this holds in higher dimensions or, more generally, in word-hyperbolic groups. We do know that rays do not necessarily converge with positive speed. Note that if we fixed a point in the hyperbolic plane and took the nearest vertex of the forest to that point as the root of its component, then we would obtain a URT.

2. *Definitions*

We review a few definitions from the theory of unimodular random rooted networks; for more details, see [**2**]. A *network* is a (multi-)graph $G = (\mathsf{V}, \mathsf{E})$ together with a complete separable metric space Ξ, called the *mark space*, and maps from V and E to Ξ. Images in Ξ are called *marks*. Each edge is given two marks, one associated to ('at') each of its endpoints. The only assumption on degrees is that they are finite. We omit the mark maps from our notation for networks.

A *rooted network* (G, o) is a network G with a distinguished vertex o of G, called the *root*. A *rooted isomorphism* of rooted networks is an isomorphism of the underlying networks that takes the root of one to the root of the other. We do not distinguish between a rooted network and its isomorphism class. Let $\mathscr{G}_*$ denote the set of rooted isomorphism classes of rooted *connected* locally finite networks. Define a separable complete metric on $\mathscr{G}_*$ by letting the distance between (G_1, o_1) and (G_2, o_2) be $1/(1+\alpha)$, where α is the supremum of those $r > 0$ such that there is some rooted isomorphism of the balls of (graph-distance) radius $\lfloor r \rfloor$ around the roots of G_i, such that each pair of corresponding

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

marks has distance less than $1/r$. For probability measures μ, μ_n on $\mathscr{G}_*$, we write $\mu_n \Rightarrow \mu$ when μ_n converges weakly with respect to this metric.

For a (possibly disconnected) network G and a vertex $x \in \mathsf{V}(G)$, write G_x for the connected component of x in G. If G is finite, then write U_G for a uniform random vertex of G and $U(G)$ for the corresponding distribution of (G_{U_G}, U_G) on $\mathscr{G}_*$. Suppose that G_n are finite networks and that μ is a probability measure on $\mathscr{G}_*$. We say that the *random weak limit* of G_n is μ if $U(G_n) \Rightarrow \mu$.

A probability measure that is a random weak limit of finite networks is called *sofic*. In particular, a finitely generated group is called sofic when its Cayley diagram is sofic. It is easy to check that this property does not depend on the generating set chosen. All sofic measures are unimodular, which we now define. Similarly to the space $\mathscr{G}_*$, we define the space $\mathscr{G}_{**}$ of isomorphism classes of locally finite connected networks with an ordered pair of distinguished vertices and the natural topology thereon: the distance between (G_1, o_1, o_1') and (G_2, o_2, o_2') is $1/(1+\alpha)$, where α is the supremum of those $r > 0$ such that there is some isomorphism of the balls of radius $\lfloor r \rfloor$ around o_i that takes o_1 to o_2 and o_1' to o_2', such that each pair of corresponding marks has distance less than $1/r$. We shall write a function f on $\mathscr{G}_{**}$ as $f(G, x, y)$. We refer to $f(G, x, y)$ as the *mass* sent from x to y in G.

Definition 2.1. Let μ be a probability measure on $\mathscr{G}_*$. We call μ *unimodular* if it obeys the *mass-transport principle*: for all Borel $f : \mathscr{G}_{**} \to [0, \infty]$, we have

$$\int \sum_{x \in \mathsf{V}(G)} f(G, o, x)\, d\mu(G, o) = \int \sum_{x \in \mathsf{V}(G)} f(G, x, o)\, d\mu(G, o). \tag{2.1}$$

It is easy to see that every sofic measure is unimodular, as observed by Benjamini and Schramm [**10**], who introduced this general form of the mass-transport principle under the name 'intrinsic mass-transport principle'. The converse is open and was posed as a question by Aldous and Lyons [**2**].

A special form of the mass-transport principle was considered, in different language, by Aldous and Steele [**3**]. Namely, they defined μ to be *involution invariant* if (2.1) holds for those f supported on (G, x, y) with $x \sim y$. In fact, the mass-transport principle holds for general f if it holds for these special f, as shown by Aldous and Lyons [**2**].

PROPOSITION 2.2. *A measure is involution invariant if and only if it is unimodular.*

If G is μ-a.s. regular, then involution invariance of μ is equivalent to the following: if o' is a uniform random neighbor of the root, then the law of (G, o, o') is the same as the law of (G, o', o) when (G, o) has the law μ.

See also [**6**] for a discussion of unimodularity.

We call a measure a *URT* if it is a unimodular probability measure on rooted networks whose underlying graphs are trees. We call a probability measure a *labeled percolation* on a graph G if it is carried by the set of networks on G whose marks are pairs, with the second coordinate, called *color*, of a mark being 0 or 1. Edges colored 0 or 1 are called *closed* and *open*, respectively.

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

3. *Tightness and degree*

One of our theorems is that URTs are sofic. Although, for this purpose, we shall not need results that imply random weak convergence of a subsequence of finite graphs, such results have not been stated in the literature before except in the easy case of bounded degree and the harder case of exponential tails of the degree distribution [**4**]. On the other hand, it does not suffice that the mean degrees be bounded. For example, consider the complete bipartite graphs $K_{1,n}$ (stars): no subsequence converges. Yet, also, it is not necessary that the mean degrees be bounded. In fact, the mean degree can be infinite for extremal sofic unimodular random rooted graphs, even trees. Here, *extremal* means that the probability measure is not a convex combination of other unimodular probability measures on rooted graphs. We impose that condition since it is trivial to take a mixture of finite-mean-degree URTs to get a URT of infinite mean degree.

For example, let $\langle p_n;\ n \geq 1\rangle$ be a probability distribution on $\mathbb{Z}^+$ with infinite mean. For each integer k, join k to $k+1$ by n parallel edges with probability p_n, independently for different k. This is easily seen to be an extremal sofic probability measure. To get an extremal URT with infinite mean degree, take the universal cover rooted at 0 of the resulting multigraph; see [**2**, Example 9.3].

Thus, it may be useful to present the following result on tightness. For simplicity, and with no essential loss of generality, we shall assume that all our graphs have no isolated vertices. The proof we present was suggested by Omer Angel and simplifies our original proof. The proof in Angel and Schramm [**4**] would also work.

THEOREM 3.1. *If A is a family of finite graphs such that the random variables $\{\deg_G U_G;\ G \in A\}$ are uniformly integrable, then $\{U(G);\ G \in A\}$ is tight.*

Proof. Let $f(d) := \sup_{G\in A} \mathbf{E}[\deg_G U_G;\ \deg_G U_G > d]$. By assumption, $\lim_{d\to\infty} f(d) = 0$. Write $m(G) := \mathbf{E}[\deg_G U_G]$. Thus, $1 \leq m(G) \leq f(0) < \infty$. Write $D(G)$ for the degree-biased probability measure on $\{(G, x);\ x \in V(G)\}$, that is,

$$D(G)[(G, x)] = \frac{\deg_G x}{m(G)} \cdot U(G)[(G, x)],$$

and D_G for the corresponding root. Since $U(G) \leq m(G)D(G) \leq f(0)D(G)$, it suffices to show that $\{D(G);\ G \in A\}$ is tight. Note that $\{\deg_G D_G;\ G \in A\}$ is tight.

For $r \in \mathbb{N}$, let $F_r^M(x)$ be the event such that there is some vertex at distance at most r from x whose degree is larger than M. Let X be a uniform random neighbor of D_G. Because $D(G)$ is a stationary measure for a simple random walk, $F_r^M(D_G)$ and $F_r^M(X)$ have the same probability. Also, $\mathbf{P}[F_{r+1}^M(D_G) \mid \deg_G D_G] \leq (\deg_G D_G)\mathbf{P}[F_r^M(X) \mid \deg_G D_G]$. We claim that for all $r \in \mathbb{N}$ and $\epsilon > 0$, there exists $M < \infty$ such that $\mathbf{P}[F_r^M(D_G) < \epsilon]$ for all $G \in A$; this clearly implies that $\{D(G);\ G \in A\}$ is tight. We prove the claim by induction on r.

The statement for $r = 0$ is trivial. Given that the property holds for r, let us now show it for $r+1$. Given $\epsilon > 0$, choose d so large that $\mathbf{P}[\deg_G D_G > d] < \epsilon/2$ for all $G \in A$. Also, choose M so large that $\mathbf{P}[F_r^M(D_G)] < \epsilon/(2d)$ for all $G \in A$. Write F for the event such

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

that $\deg_G D_G > d$. Then, by conditioning on $\deg_G D_G$, we see that

$$\begin{aligned}
\mathbf{P}[F_{r+1}^M(D_G)] &\le \mathbf{P}[F] + \mathbf{E}\big[\mathbf{1}_{F^c}\mathbf{P}[F_{r+1}^M(D_G) \mid \deg_G D_G]\big] \\
&\le \epsilon/2 + \mathbf{E}\big[\mathbf{1}_{F^c}(\deg_G D_G)\mathbf{P}[F_r^M(X) \mid \deg_G D_G]\big] \\
&\le \epsilon/2 + \mathbf{E}\big[\mathbf{1}_{F^c} d\mathbf{P}[F_r^M(X) \mid \deg_G D_G]\big] \\
&\le \epsilon/2 + \mathbf{E}\big[d\mathbf{P}[F_r^M(X) \mid \deg_G D_G]\big] \\
&= \epsilon/2 + d\mathbf{P}\big[F_r^M(X)\big] \\
&= \epsilon/2 + d\mathbf{P}\big[F_r^M(D_G)\big] \\
&< \epsilon/2 + d\epsilon/(2d) = \epsilon
\end{aligned}$$

for all $G \in A$, which proves the claim. □

In this proof, the only way that we used finiteness of the graphs was that the degree-biased uniform distribution on vertices gave a stationary measure for a simple random walk. Thus, the result applies also to any collection of probability measures bounded by a fixed multiple of stationary probability measures on rooted graphs, such as unimodular probability measures on graphs.

4. *Invariant percolation*

We now prove that every URT of bounded degree arises as the open component of the root in an invariant percolation on a regular tree.

We use the following lemma that is straightforward to check from the definitions. The technical definition of adding independent and identically distributed (i.i.d.) marks is explained in Aldous and Lyons [**2**, §6].

LEMMA 4.1. *Suppose that μ is a unimodular probability measure on rooted networks. Let ϕ be a measurable map on rooted networks that takes each network to an element of the mark space. Define Φ to be the map on rooted networks that takes a network (G, o) to another network on the same underlying graph, but replaces the mark at each vertex $x \in G$ by $\phi(G, x)$. Then the pushforward measure $\Phi_*\mu$ is also unimodular. If instead we add a second coordinate to each vertex mark by an i.i.d. mark according to some probability measure on the mark space, then the resulting measure is again unimodular.*

THEOREM 4.2. *Let μ be a probability measure on rooted networks whose underlying graphs are trees of degree at most d. Then μ is unimodular if and only if μ is the law of the open component of the root in a labeled percolation on a d-regular tree whose law is invariant under all automorphisms of the tree.*

Proof. The 'if' part of the assertion is well known and not dependent on the fact that the underlying graph is a tree. See, for example, [**7**] or [**2**, Theorem 3.2].

The idea for proving the converse is as follows. First, sample $(T, o) \sim \mu$. Of all the possible ways to embed it in the d-regular tree $\mathbb{T}_d$ such that o maps to the root $\mathbf{o}$ of $\mathbb{T}_d$, choose one uniformly (i.e. choose one arbitrarily and then apply a uniform automorphism of $\mathbb{T}_d$ preserving $\mathbf{o}$). The embedded image of T is marked open. Now, for every edge e in $\mathbb{T}_d$ that is not in the image of T but has one endpoint in T, mark e closed and

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

sample an independent copy $(T', o') \sim \mu$ with o' embedded as the endpoint of e that is not in T. However, this choice of T' has to be biased so that the degree of o' is not d. In fact, we sample instead $(T', o') \sim \mu'$, where μ' is absolutely continuous with respect to μ with a Radon–Nikodým derivative at (T, o) equal to $(d - \deg_T o)/\alpha$, where α is a normalizing constant. Continue in this way to cover all of the vertices of $\mathbb{T}_d$ by weighted independent copies of $(T', o') \sim \mu'$. Of course, all edges in embedded copies of T or T' are marked open, while the rest are marked closed. To prove that this is invariant, we first show involution invariance of the constructed marked tree and then appeal to Aldous and Lyons [**2**, Theorem 3.2] to get that it is an invariant percolation on $\mathbb{T}_d$. Proving involution invariance involves two cases: one case involves crossing a closed edge; that is where the biased measure μ' comes in. The other case involves crossing an open edge; that is where the unimodularity of the original measure μ comes in.

Here are the details. Let p_k be the μ-probability that the root has k children. Let $\alpha := \sum_k p_k(d - k)$ be a normalizing constant. Let μ' be absolutely continuous with respect to μ with a Radon–Nikodým derivative at (T, o) equal to $(d - \deg_T o)/\alpha$. Given two probability measures ν_1 and ν_2 supported by networks on rooted trees, where the root has degree at most $d - 1$ and all other vertices have degree at most d, write $Q(\nu_1, \nu_2)$ for the probability measure supported by networks on the rooted d-ary tree constructed as follows, similar to a Galton–Watson branching process: choose $(T', o) \sim \nu_1$, whose edges are colored open. To each vertex $x \neq o$ of T', adjoin $d - \deg_{T'} x$ edges colored closed whose other endpoint is the root of an independent sample from ν_2, while to the root o of T', adjoin $d - 1 - \deg_{T'} o$ edges colored closed whose other endpoint is the root of an independent sample from ν_2. Call the result the network (T, o). Then $Q(\nu_1, \nu_2)$ is the law of (T, o). Write ν for the measure on rooted networks defined by the equation $\nu = Q(\mu', \nu)$.

Let ρ be the measure constructed as follows: choose $(T', o) \sim \mu$, whose edges are colored open. To each vertex x of T', adjoin $d - \deg_{T'} x$ edges colored closed whose other endpoint is the root of an independent sample from ν. The result is a network whose underlying graph is $\mathbb{T}_d$. We claim that this measure ρ is unimodular, which we show by proving that ρ is involution invariant. This suffices to prove the theorem by appeal to Aldous and Lyons [**2**, Theorem 3.2] (in which the averaging over automorphisms is taken).

To prove this claim, it will be convenient to use the following technical modification of ρ to deal with counting issues: given $(T, o) \sim \rho$, assign independently and uniformly marks to the closed edges in each direction so that each vertex is surrounded by outgoing closed edges marked $1, \ldots, k$ when it is incident to k closed edges. Call ρ' the resulting measure on networks. It clearly suffices to prove that ρ' is involution invariant.

For k such that $p_k > 0$, let μ_k be the measure constructed as follows: choose $(T', o) \sim \mu$, whose edges are colored open, conditioned on $\deg_{T'} o = k$. To each vertex $x \neq o$ of T', adjoin $d - \deg_{T'} x$ edges colored closed whose other endpoint is the root of an independent sample from ν. Let $\mathcal{N}_i$ denote the class of networks supported on a rooted tree with all vertices except the root having degree d and the root having degree i. Consider $i, i' \in [1, d - 1]$ and Borel sets $A \subseteq \mathcal{N}_i$, $A' \subseteq \mathcal{N}_{i'}$, and $B_1, \ldots, B_{d-i-1}, B'_1, \ldots, B'_{d-i'-1} \subseteq \mathcal{N}_{d-1}$. Now, let $(T, o) \sim \rho'$ and let o' be a uniform neighbor of o. Then, the chance that we see (a) i open edges at o, (b) the edge (o, o') is closed with (c) mark j in the

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

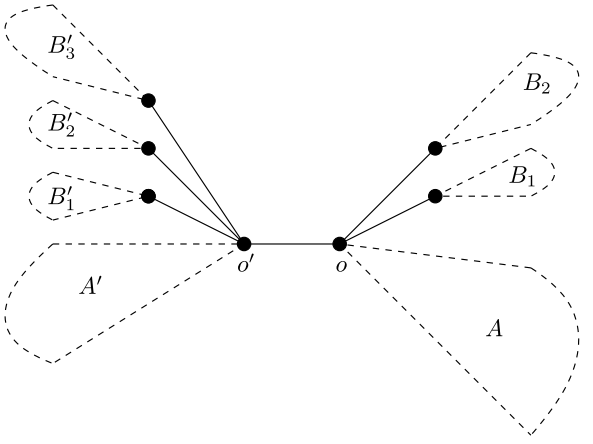


FIGURE 2. All edges drawn with solid lines are closed. The root o is incident to three closed edges, while o' is incident to four closed edges.

direction (o, o') and (d) mark j' in direction (o', o), see (e) i' open edges at o', and see the event where (f) the open edges at o are part of a network in A, (g) the open edges at o' are part of a network in A', while (h) the other endpoints of the closed edges at o belong to networks in $(B_1, \ldots, B_{d-i-1})$ in increasing order of their marks from o, and (i) similarly the other endpoints of the closed edges at o' belong to networks in $(B_1', \ldots, B_{d-i'-1}')$ in increasing order of their marks from o' (see Figure 2) equals

$$\underbrace{p_i}_{\text{(a)}} \cdot \underbrace{\mu_i(A)}_{\text{(f)}} \cdot \underbrace{\prod_{m=1}^{d-i-1} \nu(B_m)}_{\text{(h)}} \cdot \underbrace{\frac{d-i}{d}}_{\text{(b)}} \cdot \underbrace{\frac{1}{(d-i)}}_{\text{(c)}} \cdot \underbrace{\frac{1}{(d-i')}}_{\text{(d)}} \cdot \underbrace{\prod_{r=1}^{d-i'-1} \nu(B_r')}_{\text{(i)}}$$
$$\cdot \underbrace{p_{i'} \frac{d-i'}{\alpha}}_{\text{(e)}} \cdot \underbrace{\mu_{i'}(A')}_{\text{(g)}}.$$

This is invariant under the involution exchanging o and o'.

The other case to prove is when the edge (o, o') is open. Consider the following measure, σ. Begin with a sample $(T, o) \sim \mu$. Assign a second coordinate $(\tau_1(x), \ldots, \tau_{d-1}(x))$ to the vertex mark at each vertex x given by i.i.d. samples $\tau_i(x) \sim \nu$. This new network is unimodular by Lemma 4.1. Now, replace the second coordinate of the vertex mark at each vertex x by $(\tau_1(x), \ldots, \tau_{d-\deg_T(x)}(x))$. This new network is again unimodular, by Lemma 4.1. We denote by σ its law. Note that we can obtain ρ from σ by replacing the second coordinate of the vertex mark at each x by a tree network rooted at x, where we adjoin $d - \deg_T(x)$ closed edges to x, at the other end of which we adjoin the trees $\tau_i(x)$.

What remains to prove is that involution invariance holds for ρ' across open edges. It suffices to do the same for ρ, but this is clearly the same as unimodularity of σ. □

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

5. *Soficity*

We now use the preceding theorem to prove that URTs are sofic. This result is the same as Elek and Lippner's [**15**, Theorem 4], but in different language. See [**2**, Example 9.9] for a comparison of the different languages†.

THEOREM 5.1. *Every URT is sofic.*

Proof. It follows from Bowen [**12**, Theorem 3.4] that every invariant network on $\mathbb{T}_d$ is sofic; the result is stated there for even d only, but the proof works for all d. Thus, given a URT μ, if the degrees are bounded by d, let ρ be an invariant labeled percolation on $\mathbb{T}_d$ such that the open component of the root has law μ. Let $\langle G_n;\ n \geq 0\rangle$ be finite networks whose random weak limit is ρ. Here, we may assume that the edges of G_n are each colored closed or open. Let G'_n be the result of deleting every closed edge from G_n. Then, clearly, $\langle G'_n\rangle$ has random weak limit μ. Finally, if the degrees are not bounded μ-a.s., then for each d, let μ_d be the law of the component of o when we delete every edge of T incident to some vertex of degree larger than d, where $(T, o) \sim \mu$. Then, μ_d is unimodular and, by what we just proved, sofic. Since the sofic measures form a weakly closed set and μ_d tend weakly to μ, we deduce that μ is sofic as well. □

As noted by Elek and Lippner [**15**], this implies that every treeable group is sofic. Here, a group Γ is *treeable* if there is a probability measure on trees with vertex set Γ that is invariant under the natural action of Γ; such a probability measure is called a *treeing* of Γ. Briefly, the idea is to use a treeing μ of Γ, a generating set S for Γ, and a sofic approximation $\langle G_n\rangle$ of μ to construct a sofic approximation of the Cayley diagram of Γ with respect to S by putting edges labeled $s \in S$ between points x, y of G_n, such that a path from x to y has length at most R_n and has labels that multiply to s, where $R_n \to \infty$ at an appropriately slow rate.

6. *Rays*

Random weak limits of finite trees have mean degree at most 2, are supported by trees with at most two ends, and hence are recurrent for simple random walks; see [**2**, Proposition 6.3]. In the case of URTs with finite mean degree larger than 2, the speed of the simple random walk is positive: see [**2**, Theorem 4.9]. The case of infinite mean degree is open. However, it is interesting in all cases to see whether the rays themselves, rather than the simple random walk, have positive speed when embedded in a larger graph. What we mean by this is the following.

We say that a sequence $\langle x_n;\ n \geq 0\rangle$ in a metric space has *positive (liminf) speed* if there is some constant $c > 0$ such that the distance between x_n and x_0 is at least cn for all $n \geq 1$. A simple infinite path in a tree is called a *ray*. An *end* of a tree is an equivalence class of rays, where two rays are equivalent when they have finite symmetric difference. Of course, any statement about limits of rays applies equally to all rays belonging to the same end and is therefore a statement about limits of ends.

† Actually, the notion of 'sofic' in Elek and Lippner [**15**] applies only to certain URNs, namely, those where the rooted network is a measurable function of the label of the root. Thus, the result we prove is superficially more general.

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

We are interested in the rays in forests that arise either as invariant percolation on a Cayley graph or as random graphs discretely embedded in hyperbolic space $\mathbb{H}^d$ with an isometry-invariant law. When do all the rays have positive speed in the metric of the Cayley graph or in the hyperbolic metric? It does not suffice that the Cayley graph be non-amenable: for example, consider the usual Cayley graph of the group $\mathbb{Z} * \mathbb{Z}^2$. Use the random forest that arises from an independent copy of the uniform spanning tree [**22**] in every copy of $\mathbb{Z}^2$. Then, a.s. each such tree contains only one end and no ray has positive speed. (In fact, the nth vertex in a ray is roughly at distance $n^{4/5}$ from its starting point; see [**5**].) What if we restrict ourselves to word-hyperbolic groups? As we shall see, the answer is still no. Thus, we focus on the following weaker property for hyperbolic groups.

Question 6.1. Does every ray in an invariant forest in a word-hyperbolic group converge a.s. to an ideal boundary point?

See the survey [**19**] for information on the boundary of a word-hyperbolic group. We know the answer only in $\mathbb{H}^d$ for $d = 2$.

THEOREM 6.2. *Let G be a one-ended graph embedded in $\mathbb{H}^2$ such that a group of isometries of $\mathbb{H}^2$ acts quasi-transitively on G. Given an automorphism-invariant forest in G, a.s. every ray in the forest converges to an ideal boundary point. Furthermore, the set of limits of the rays is a.s. the entire ideal boundary.*

Here, to say that G is one-ended means that the complement of each finite set in G has only one infinite component in G. Note that the ideal boundary points of G are the same as those of $\mathbb{H}^2$. We call the set of limit points of the convergent rays in a tree or forest the *limit set* of that tree or forest.

One can prove a similar statement for forests in $\mathbb{H}^2$ whose law is invariant under isometries of $\mathbb{H}^2$, without assuming an underlying graph, G. On the other hand, one could also let the underlying graph G be random with isometry-invariant law; no quasi-transitivity of G is then needed, nor need G have only one end. The lengths of edges should be bounded and the vertices should be separated by a minimum distance in $\mathbb{H}^2$.

We are grateful to Omer Angel for some simplifications to our proof.

Proof. Without loss of generality, we may assume that the forest $\mathfrak{F}$ is spanning and contains only infinite trees. Indeed, we may first delete all finite trees and then independently add an edge at random from each vertex not in the forest to a vertex that is closer to the forest.

Let $G^\dagger$ be the planar dual graph of G; its edges are in bijective correspondence with those of G in such a way that each edge e of G crosses only its corresponding edge $e^\dagger$. Let $\mathfrak{F}^\times$ be the dual spanning forest in $G^\dagger$ defined by $e^\dagger \in \mathfrak{F}^\times$ if and only if $e \notin \mathfrak{F}$. Now, G and $G^\dagger$ are unimodular since the isometry group of $\mathbb{H}^2$ is unimodular, whence so are the co-compact subgroups of isometries that fix G and $G^\dagger$.

Note that the process of adding edges in the first paragraph a.s. does not increase the set of ends of any tree: if it did, we could transport mass 1 from every vertex not originally in a tree to the vertex x in the tree T where it joins T. This would give x infinite mass whenever there is a new ray beginning at x that uses no edge from T. By the mass-transport principle, the probability of this event is 0.

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

Let $\mathfrak{F}_3$ be the set of trees in $\mathfrak{F}$ that have at least three ends. If $\mathbf{P}[\mathfrak{F}_3 \neq \varnothing] > 0$, then by conditioning on $\mathfrak{F}_3 \neq \varnothing$, we may assume (temporarily) that this probability is 1. Now a simple random walk on the forest $\mathfrak{F}_3$ a.s. has positive speed in the metric of the forest (as we noted above or, e.g., by Lyons with Peres [**21**, Theorem 16.4]), whence it also has positive speed in the graph metric of G by Benjamini *et al* [**8**, Lemma 4.6], and hence also in the hyperbolic metric. Therefore, as in Benjamini and Schramm [**9**, proof of Theorem 4.1], a simple random walk on the forest $\mathfrak{F}_3$ a.s. converges to an ideal boundary point. It follows that, a.s. for every tree $T \in \mathfrak{F}_3$, we have that μ_T-almost every ray converges to an ideal point, where μ_T is a harmonic measure on the boundary of T.

Let A be the limit set of $\mathfrak{F}$. We have just shown that $A \neq \varnothing$ a.s. It follows that A is dense a.s. Indeed, for every $\epsilon > 0$, choose B_ϵ to be a non-dense subset of the ideal boundary for which $\mathbf{P}[A \cap B_\epsilon \neq \varnothing] > 1 - \epsilon$. Now, let B be any non-empty open subset of the ideal boundary. There is an isometry of $\mathbb{H}^2$ that induces an automorphism of G and carries B_ϵ into B. Since our probability measure is invariant under automorphisms, it follows that $\mathbf{P}[A \cap B \neq \varnothing] > 1 - \epsilon$. Since this holds for every $\epsilon > 0$, we deduce that $\mathbf{P}[A \cap B \neq \varnothing] = 1$. Since the ideal boundary is separable, the claim follows.

The density of A now implies that all rays converge a.s., whence A is the entire ideal boundary. To see this, consider any two points $\xi \neq \eta \in A$. Let P, Q be rays that converge to ξ, η, respectively. Let S be a path between the initial vertices of P and Q. No ray can cross $P \cup Q \cup S$ infinitely many times, whence its limit set must be contained in one of the two closed arcs determined by ξ and η. Our freedom in choosing ξ and η from the density of A now gives the result. The same argument shows that the limit set of $\mathfrak{F}^\times$ is the entire ideal boundary a.s.

If the number of trees in $\mathfrak{F}$ with one end is at least three, then either $\mathfrak{F}$ or $\mathfrak{F}^\times$ must contain a tree with at least three ends. Hence, we may again conclude that the limit set of $\mathfrak{F}$ is the entire ideal boundary a.s. On the other hand, the probability is 0 that the number of trees with one end is positive and finite, since if the probability is positive, then, as before, we may assume that the entire forest consists of such trees a.s. In fact, by choosing just one of the trees at random, we may assume that there is only one tree with one end. This contradicts [**7**, Theorem 5.3] since G is non-amenable.

The only remaining case, then, is that all trees in $\mathfrak{F}$ and $\mathfrak{F}^\times$ have two ends. We claim that this has probability 0, for, when they do, we can order the trees as the integers in the following sense. Each tree T in $\mathfrak{F}$ separates the plane into two pieces since it has two ends. The dual of the edge boundary of T lies in $\mathfrak{F}^\times$ and has two connected components, each one being part of a tree in $\mathfrak{F}^\times$. The same applies to each of those trees in turn, which means that on each side of those trees, besides T, there is another tree in $\mathfrak{F}$ that includes the dual of part of its edge boundary. Those two trees in $\mathfrak{F}$ are the ones next to T in the integer ordering of all the trees in $\mathfrak{F}$. This allows us to define an invariant percolation with all clusters finite yet with arbitrarily high marginal, contradicting non-amenability by Benjamini *et al* [**7**, Theorem 2.12]. To see this, call the unique bi-infinite simple path in a tree with two ends the *trunk* of that tree. If a vertex x of the trunk is deleted and y is in a finite component of what is left of the tree (or $y = x$), then call x the *trunk attachment* of y. Now, given $\epsilon > 0$, delete each tree of $\mathfrak{F}$ with probability ϵ independently and in each tree that is left, delete each vertex on the trunk with probability ϵ independently, and delete all vertices not on a

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

trunk whose trunk attachment was deleted. Thus, each vertex is deleted with probability $\epsilon + (1-\epsilon)\epsilon$. It remains to show that the graph induced by the remaining vertices has no infinite clusters a.s. Number the trees by $\mathbb{Z}$ as indicated above, where we choose arbitrarily which tree is numbered 0 and in which direction the integers increase. Suppose that trees numbered m and $m+n+1$ are deleted, while the n trees numbered i are not deleted for $m < i < m+n+1$. Consider a vertex x_1 on the trunk of tree number $m+1$. Then, there is at least one vertex x_2 on the trunk of tree number $m+2$ such that, for some y_1 whose trunk attachment is x_1, and some y_2 whose trunk attachment is x_2, there is an edge of G between y_1 and y_2. Likewise, we may choose x_3 on the trunk of the tree number $m+3$ such that some y_3 whose trunk attachment is x_3 is adjacent to a vertex whose trunk attachment is x_2, etc. If all vertices $x_1, \ldots, x_n$ are deleted, then there is a path in G of deleted vertices stretching from tree number m to tree number $m+n+1$. The probability that $x_1, \ldots, x_n$ are all deleted is ϵ^n (recall that trees number $m+1, \ldots, m+n$ are not deleted). We may choose infinitely many such sequences $\langle x_1, \ldots, x_n \rangle$ that are pairwise disjoint, so that the corresponding events when these sequences are deleted are independent, each having the same probability ϵ^n. Therefore, infinitely many such events occur a.s., and, when they do, they separate the remaining vertices between trees m and $m+n+1$ into finite components. Since this happens between each consecutive pair of deleted trees, all components are finite a.s. This completes the proof of the theorem. □

We can give some additional information about the limit points of the rays in those trees with at least three ends.

THEOREM 6.3. *Let G be a one-ended graph embedded in $\mathbb{H}^2$ such that a group of isometries of $\mathbb{H}^2$ acts quasi-transitively on G. Let $\mathfrak{F}$ be an invariant forest in G such that each tree has at least three ends. A.s., for each tree in $\mathfrak{F}$, the map from ends to limit points never maps more than two ends to the same limit. In addition, given that there is more than one tree in $\mathfrak{F}$, a.s. for each tree, the limit set is a perfect nowhere-dense set.*

Proof. Consider a point ξ of the ideal boundary that is a limit of at least two rays of some tree T of $\mathfrak{F}$. Among all the bi-infinite paths in T both of whose ends converge to ξ, there is only one, call it P, such that all others lie inside the closed curve defined by $P \cup \{\xi\}$. Indeed, if not, it would follow that ξ is the only limit point of the ends of $\mathfrak{F}$. Since there is no invariant probability measure on the ideal boundary, the event that such a ξ exists has probability 0. Now each end with a limit ξ contains a unique ray that starts at a point in P. Let each vertex in such a ray send mass 1 to its starting point in P. Then some points in P get infinite mass, so, by the mass-transport principle, this has probability 0. This establishes that the map from ends to limit points a.s. never maps more than two ends to the same limit. In particular, the limit set is a.s. infinite.

Now suppose that $\mathfrak{F}$ has more than one tree a.s. Let T be one of them. If its limit set is not nowhere dense, then it contains a maximal proper arc. Let I be one such arc. Among all the bi-infinite paths in T both of whose ends converge to points in I, there is only one, call it P, such that all others lie on the same side of P as I. If not, I would be the entire limit set of $\mathfrak{F}$, and by choosing an endpoint of I at random, we would again obtain an invariant probability measure on the ideal boundary, an impossibility. Now, each end with

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

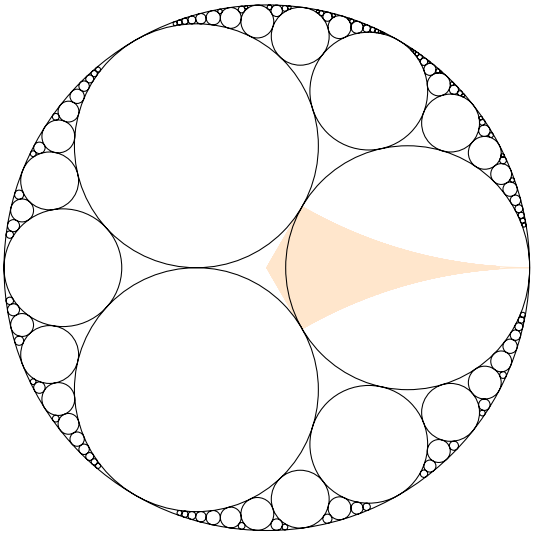


FIGURE 3. Ford horocycle tiling with a fundamental domain.

a limit in I contains a unique ray that starts at a point in P. Let each vertex in such a ray send mass 1 to its starting point in P. Then, some points in P get infinite mass, so, by the mass-transport principle, this has probability 0. This establishes that the limit set is a.s. nowhere dense.

If the limit set is not perfect, then there is an isolated limit point, ξ, and a vertex x of T such that three rays from x that are disjoint other than at x have distinct limit points, one being ξ. In fact, for each isolated limit point ξ, there is a unique such x that is 'closest' to ξ, in that the ray from x to ξ contains no other vertex with these properties, but then we can transport to x mass 1 from each vertex on the ray from x to ξ and so, by the mass-transport principle, this has probability 0. □

In order to give examples where forests have rays that do not have positive speed, it will be convenient to work first in the context given after the statement of Theorem 6.2, that is, we consider first forests in $\mathbb{H}^d$ whose law is invariant under isometries of $\mathbb{H}^d$. The lengths of edges are bounded and the vertices are separated by a minimum distance in $\mathbb{H}^d$.

Note that one cannot exhibit a trivial example of a forest with zero-speed rays by subdividing edges in a random forest in such a way that the number of subdivisions has infinite mean, for then there is no way to re-embed the forest in an invariant fashion. Our examples were discovered in conversation with David Fisher. Take a random invariant collection of disjoint horoballs. For example, in $\mathbb{H}^2$, one can apply a random isometry to the Ford circles in the upper half-plane model; this is possible since the stabilizer of the set of Ford circles has co-finite measure in the full isometry group. A fundamental domain (up to rotation in the tangent bundle) is shown in Figure 3 for the Poincaré disc model of $\mathbb{H}^2$. The fundamental domain has area $\pi/3$, since it is composed of two congruent geodesic triangles, each of which has angles of measure $\pi/3$, $\pi/2$, and 0. For general $d \geq 2$, Garland and Raghunathan [**17**] show that for every discrete subgroup Γ of isometries of $\mathbb{H}^d$ whose

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

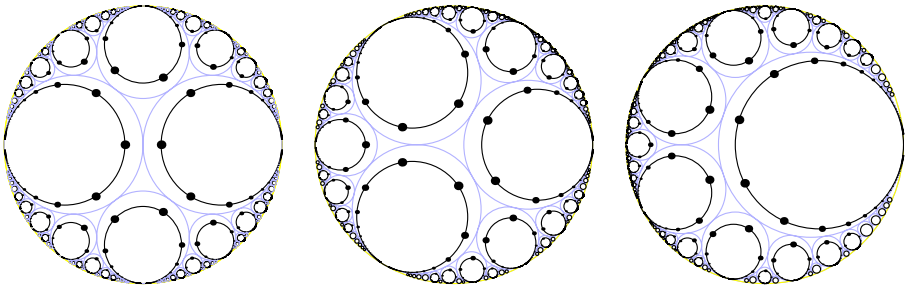


FIGURE 4. The forest is in black. The Ford horocycles are in light blue. The left and middle figures show examples from the corners of the fundamental domain, while the right figure shows a more typical example.

quotient $\mathbb{H}^d/\Gamma$ is non-compact with finite volume, there exists a Γ-invariant collection of disjoint horoballs in $\mathbb{H}^d$. For a proof that such subgroups Γ exist, see [**23**, Ch. 14]. By choice of Γ, we then obtain a probability measure on collections of disjoint horoballs that is invariant under all isometries of $\mathbb{H}^d$. A horosphere is geodesically flat, so just inside each horoball and at a fixed distance δ from the horosphere, we put a copy of $\mathbb{Z}^{d-1}$ lying on another horosphere. Note that we are free to choose any distance δ we want. Again, this is done in an invariant way. Now take any random invariant forest i.i.d. in each copy of $\mathbb{Z}^{d-1}$. None of the rays have positive speed, though all converge. See Figure 4 for three examples in $\mathbb{H}^2$. (Such examples can be constructed similarly in complex and quaternionic hyperbolic spaces, as well as the octonionic plane. Although the horospheres are not then flat, one can take them to be δ-separated for any $\delta > 0$, and one may use an embedded Cayley graph of the stabilizer of a horosphere, a group that is finitely generated, in which one can take, say, a minimal spanning forest independently in each horosphere.)

Now, we transfer such examples to the setting of a graph G that is quasi-isometric to $\mathbb{H}^d$ and is randomly embedded in $\mathbb{H}^d$ in an isometry-invariant fashion. For each vertex x of the forest, let $\phi(x)$ be a vertex in G that is nearest to x; a.s., there is only one choice for $\phi(x)$. If x and y are neighbors in the forest, then join $\phi(x)$ to $\phi(y)$ by a shortest path in G; when there is more than one choice, choose at random uniformly and independently. Choose the distance δ so that ϕ is injective and these shortest paths are disjoint when they come from distinct edges. In this fashion, ϕ induces an embedding of the forest to a subgraph in G. The choice of δ ensures that this subgraph is a forest as well. It is the desired example.

Acknowledgements. We are grateful to Omer Angel for simplifications to some of our proofs and to David Fisher for permission to include the example of zero-speed rays in hyperbolic space. The research was partially supported by Microsoft Research and NSF grant DMS-1007244.

REFERENCES

[1] M. Aizenman and S. Warzel. The canopy graph and level statistics for random operators on trees. *Math. Phys. Anal. Geom.* **9** (2006), 291–333.

[2] D. J. Aldous and R. Lyons. Processes on unimodular random networks. *Electron. J. Probab.* **12**(54) (2007), 1454–1508 (electronic).

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press

[3] D. J. Aldous and J. M. Steele. The objective method: probabilistic combinatorial optimization and local weak convergence. *Probability on Discrete Structures (Encyclopaedia Mathematical Sciences, 110)*. Ed. H. Kesten. Springer, Berlin, 2004, pp. 1–72.
[4] O. Angel and O. Schramm. Uniform infinite planar triangulations. *Comm. Math. Phys.* **241** (2003), 191–213.
[5] M. T. Barlow and R. Masson. Exponential tail bounds for loop-erased random walk in two dimensions. *Ann. Probab.* **38** (2010), 2379–2417.
[6] I. Benjamini and N. Curien. Ergodic theory on stationary random graphs. *Electron. J. Probab.* **17** (2012), Article 93, 20 pp. (electronic).
[7] I. Benjamini, R. Lyons, Y. Peres and O. Schramm. Group-invariant percolation on graphs. *Geom. Funct. Anal.* **9** (1999), 29–66.
[8] I. Benjamini, R. Lyons and O. Schramm. Percolation perturbations in potential theory and random walks. *Random Walks and Discrete Potential Theory (Symposia Mathematica, XXXIX)*. Eds. M. Picardello and W. Woess. Cambridge University Press, Cambridge, 1999, pp. 56–84. Papers from the workshop held in Cortona, 1997.
[9] I. Benjamini and O. Schramm. Percolation in the hyperbolic plane. *J. Amer. Math. Soc.* **14** (2001), 487–507.
[10] I. Benjamini and O. Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.* **6**(23) (2001), 13 pp. (electronic).
[11] B. Bollobás and O. Riordan. Sparse graphs: metrics and random models. *Random Structures Algorithms* **39** (2011), 1–38.
[12] L. Bowen. Periodicity and circle packings of the hyperbolic plane. *Geom. Dedicata* **102** (2003), 213–236.
[13] A. Connes, J. Feldman and B. Weiss. An amenable equivalence relation is generated by a single transformation. *Ergod. Th. & Dynam. Sys.* **1** (1981), 431–450 (1982).
[14] G. Elek. On the limit of large girth graph sequences. *Combinatorica* **30** (2010), 553–563.
[15] G. Elek and G. Lippner. Sofic equivalence relations. *J. Funct. Anal.* **258** (2010), 1692–1708.
[16] D. Gaboriau and R. Lyons. A measurable-group-theoretic solution to von Neumann's problem. *Invent. Math.* **177** (2009), 533–540.
[17] H. Garland and M. S. Raghunathan. Fundamental domains for lattices in (R-)rank 1 semisimple Lie groups. *Ann. of Math.* (2) **92** (1970), 279–326.
[18] G. Hjorth. A lemma for cost attained. *Ann. Pure Appl. Logic* **143** (2006), 87–102.
[19] I. Kapovich and N. Benakli. Boundaries of hyperbolic groups. *Combinatorial and Geometric Group Theory (Contemporary Mathematics, 296)*. Eds. S. Cleary, R. Gilman, A. G. Myasnikov and V. Shpilrain. American Mathematical Society, Providence, RI, 2002, pp. 39–93. Papers from the AMS Special Sessions on Combinatorial Group Theory and on Computational Group Theory held in New York, November 4–5, 2000 and in Hoboken, NJ, April 28–29, 2001.
[20] L. Lovász. *Large Networks and Graph Limits*. American Mathematical Society, Providence, RI, 2012.
[21] R. Lyons and Y. Peres. *Probability on Trees and Networks*. Cambridge University Press, 2013, in preparation. Current version available at http://mypage.iu.edu/~rdlyons/.
[22] R. Pemantle. Choosing a spanning tree for the integer lattice uniformly. *Ann. Probab.* **19** (1991), 1559–1574.
[23] M. S. Raghunathan. *Discrete Subgroups of Lie Groups (Ergebnisse der Mathematik und ihrer Grenzgebiete, 68)*. Springer, New York, 1972.
[24] A. Teplyaev. Spectral analysis on infinite Sierpiński gaskets. *J. Funct. Anal.* **159** (1998), 537–567.

https://doi.org/10.1017/etds.2013.56 Published online by Cambridge University Press