The Boundary of a Square Tiling of a Graph coincides with the Poisson Boundary

Agelos Georgakopoulos[*]

Mathematics Institute
University of Warwick
CV4 7AL, UK

January 24, 2014

Abstract

Answering a question of Benjamini & Schramm [8], we show that the Poisson boundary of any planar, uniquely absorbing (e.g. one-ended and transient) graph with bounded degrees can be realised geometrically as a circle, namely as the boundary of a tiling of a cylinder by squares. This implies a conjecture of Northshield [34] of similar flavour. For our proof we introduce a general criterion for identifying the Poisson boundary of a stochastic process that might have further applications.

1 Introduction

1.1 Overview

In this paper we prove the following fact, conjectured by Benjamini & Schramm [8, Question 7.4.]

Theorem 1.1. *If G is a plane, uniquely absorbing[1] graph with bounded degrees, then the boundary of its square tiling is a realisation of its Poisson boundary.*

The main implication of this is that every such graph G can be embedded inside the unit disc $\mathbb{D}$ of the real plane in such a way that random walk on G converges to $\partial\mathbb{D}$ almost surely, its exit distribution coincides with Lebesgue measure on $\partial\mathbb{D}$, and there is a one-to-one correspondence between the bounded harmonic functions on G and $L^\infty(\partial\mathbb{D})$ (the innovation of Theorem 1.1 is the last sentence).

Knowing that a graph G is not uniquely absorbing provides a lot of information about its Poisson boundary; in particular, G is not Liouville. Thus it is not a significant restriction in our context to assume that G is uniquely absorbing.

We use Theorem 1.1 to prove a conjecture of Northshield [34] of similar flavour (Conjecture 8.1), thus obtaining a further geometric realisation of the Poisson boundary of the graphs in question.

In the case where G is hyperbolic, we can say a bit more:

---

[1]See Section 2 for definitions.

Corollary 1.2. *Let G be an infinite, Gromov-hyperbolic, non-amenable, 1-ended, plane graph with bounded degrees and no infinite faces. Then the following five boundaries of G (and the corresponding compactifications of G) are canonically homeomorphic to each other: the hyperbolic boundary, the Martin boundary, the boundary of the square tiling, the Northshield circle, and the boundary $\partial_{\cong}(G)$.*

We present examples showing that all these conditions are necessary to make the statement true and none of them is implied by the others, except that it is not clear whether both the bounded degree and the non-amenability conditions are necessary; see Problems 8.3 and 8.4. The least obvious case is to show that hyperbolicity is not implied by the other properties, contrary to another conjecture of Northshield [34]; we disprove that conjecture by a counterexample in Section 8.2.

The equivalence of the Martin boundary to the hyperbolic boundary in Corollary 1.2 follows from a result of Ancona [1, 2] that can be thought of as a discrete version of the theorem of Anderson & Schoen [3] that for any complete, simply connected Riemannian manifold with bounded and negative sectional curvatures, the Martin boundary with respect to the Laplace-Beltrami operator coincides with the geometric boundary. This motivates

Conjecture 1.1. *Let M be a complete, simply connected Riemannian surface with Gaussian curvatures bounded between two negative constants. Let $f : M \to \mathbb{D}$ be a conformal map from M to the open unit disc in $\mathbb{C}$. Then for every 1-way infinite geodesic γ in M, the image $f(\gamma)$ converges to a point in the boundary $\mathbb{S}^1$ of $\mathbb{D}$, and this convergence determines a homeomorphism from the sphere at infinity of M to $\mathbb{S}^1$. (In particular, images of equivalent geodesics converge to the same point of $\mathbb{S}^1$.)*

For the proof of Theorem 1.1 the following criterion for identifying the Poisson boundary of a general Markov chain is introduced, that might have further applications. A function $h : V \to [0, 1]$ is *sharp*, if its values along the trajectory of the Markov chain converge to 0 or 1 almost surely.

Theorem 1.3. *Let M be an irreducible Markov chain and $\mathcal{N}$ an M-boundary. Then $\mathcal{N}$ is a realisation of the Poisson boundary of G if and only if it is faithful to every sharp harmonic function.*

Loosely speaking, Theorem 1.3 states that in order to check that a candidate space is a realisation of the Poisson boundary, it suffices to consider its behaviour with respect to the sharp harmonic functions rather than all bounded harmonic functions. This fact, and its proof, can be generalised to many continuous stochastic processes. We will have a closer look at Theorem 1.3 and its implications in Section 1.4, after some background information.

1.2 Square tilings —a discrete analogue of conformal uniformization

Motivated by Dehn's problem of dissecting a rectangle into squares of distinct side lengths [13], Brooks et. al. [10] showed how any finite planar graph G can be associated with a tiling of a rectangle by squares in such a way that every edge

e of G corresponds to a square R_e in the tiling and every vertex x corresponds to an interval tangent with all squares corresponding to the edges of x. The construction was made using an electrical current on the graph, and the square R_e is given side length equal to the current going through e, while the position of R_e is determined by the voltages of the vertices (see Section 5 for details). This construction has become a classic, with a lot of applications ranging from recreational mathematics to statistical mechanics; see e.g. [12, 15, 30, 31].

Benjamini and Schramm [8] showed that the construction of [10] can be applied to an infinite planar graph G in the uniquely absorbing case; the electrical current now emanates from a single vertex o and escapes to infinity, and is intimately connected to the behaviour of random walk from o. The square tiling takes place on the cylinder $K = \mathbb{R}/\mathbb{Z} \times [0,1]$, the edges of the graph being mapped to disjoint squares tiling K; see Figure 1 or [8, FIG. 1] for examples. Their motivation was to find discrete analogues of Riemann's mapping theorem; the following quote is from [8]

"The tiling plays the same role as conformal uniformization does for planar domains. In fact, the proof of its existence illustrates parallels with the continuous theory. In a way, it is a discrete analogue of Riemann's mapping theorem."

For more on the relationship between square tilings and Riemann's mapping theorem see [11]. The paper [7] is similar in nature to [8] and provides a further discrete analogue of Riemann's mapping theorem. A well known corollary of both [7, 8] is that every planar transient graph admits non-constant harmonic functions of finite energy. This result, in fact a detailed description of the space of such functions, can now be derived combining Theorem 1.1 with the results of [20].
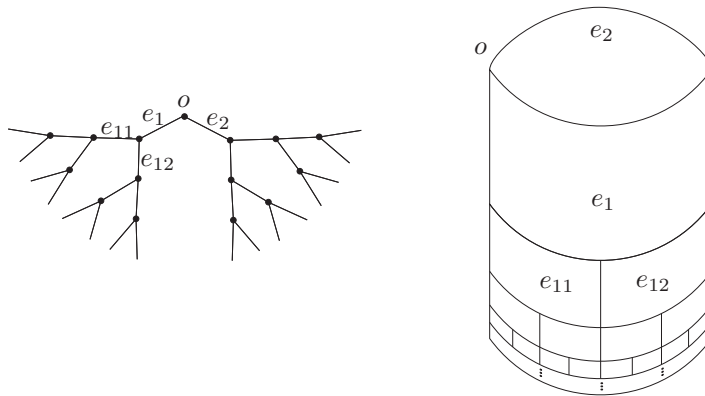

Figure 1: The infinite binary tree and its square tiling.

The quest for discrete occurrences of conformal invariance, either as discretised analogues of the continuous theory, or as original constructions, is imminent in much of the work of Schramm and was followed up by many others, leading to several impressive applications; see [36] for a survey.

1.3 The Poisson boundary of a stochastic process

The Poisson(-Furstenberg) boundary of a transient stochastic process M is a measure space $\mathcal{P}$ associated with M such that every bounded harmonic function on the state space of M can be represented by an integral on $\mathcal{P}$ and, conversely, every function in $L^\infty(\mathcal{P})$ can be integrated into a bounded harmonic function of M. The standard example [18, 24] is when M denotes Brownian motion on the open unit disc $\mathbb{D}$ in the complex plane; then the classical Poisson integral representation formula $h(z) = \int_0^1 \frac{1-|z|^2}{|e^{2\pi i\theta} - z|^2} \hat{h}(\theta) d\theta = \int_0^1 \hat{h}(\theta) d\nu_z(\theta)$ recovers every continuous harmonic function h from its boundary values $\hat{h}$ on the circle $\partial\mathbb{D}$. Here, ν_z is the *harmonic measure* on $\partial\mathbb{D}$, i.e. the exit distribution of Brownian motion started at z, and can be obtained by multiplying the Lebesgue measure by the Poisson kernel $\frac{1-|z|^2}{|e^{2\pi i\theta} - z|^2}$. Thus, we can identify the Poisson boundary of Brownian motion M on $\mathbb{D}$ with $\partial\mathbb{D}$, endowed with the family of measures $\nu_z, z \in \mathbb{D}$.

In this example, the boundary was geometric and obvious. If M is now an arbitrary transient Markov chain, then there is an abstract construction of a measurable space $\mathcal{P}$, called the Poisson boundary of M, endowed with a family of measures ν_z indexed by the state space V of M, such that the formula $h(z) = \int_{\mathcal{P}} \hat{h} d\nu_z$ provides an isometry between the Banach space $H^\infty(M)$ of bounded harmonic functions of M (endowed with the supremum norm) and the space $L^\infty(\mathcal{P})$.

Triggered by the work of Furstenberg [18, 19], a lot of research has concentrated on identifying the Poisson boundary of various Markov chains, most prominently locally compact groups endowed with some measure; see [1, 25, 26, 28] just to mention some examples, and [16] for a survey including many references and some impressive applications. However, I would like to stress that this paper is not about groups (although some of its techniques might be applicable to them), but rather about extending the study of the Poisson boundary to embrace general graphs, a currently emerging quest [4, 20].

In general, one would like to identify the Poisson boundary of a given Markov chain with a geometric object as we did for $\mathbb{D}$, which could for example be a compactification of a Cayley graph on which our group acts. This task can however be very hard. Some general criteria have been developed for the case of groups, mostly by Kaimanovich [24, 26], that have helped to identify 'geometric' Poisson boundaries for certain classes of groups e.g. hyperbolic groups, but others, like the lamplighter group over Z^3 [17], defy such identification despite extensive efforts. A well-known open problem is whether the Liouville property, which can be expressed as triviality of the Poisson boundary, for simple random walk on a Cayley graph $G = Cay(\Gamma, S)$ is a group invariant, i.e. independent of the choice of the generating set S of a given group Γ.

The aforementioned general criteria for geometric identification of the Poisson boundary only apply to groups. Theorem 1.3, and its corollary Corollary 1.4 below, may be the first general criterion for arbitrary Markov chains, while still being applicable in the case of groups.

1.4 More on our results

There are various definitions of the Poisson boundary for general Markov chains in the literature. In this paper, rather than choosing one of them, or introducing our own, we will follow a more flexible approach, accepting any measure space that fulfils the properties expected from "the" Poisson boundary as "a" Poisson boundary. More precisely, given a transient Markov process M, we let an M-*boundary*, or a G-*boundary* if M happens to be random walk on a graph G, be any measurable space $\mathcal{N}$ endowed with a family of measures $(\nu_z), z \in V$ (recall our $\mathbb{D}$ example) and a measurable, measure preserving, shift-invariant function f from the space of random walk trajectories $\mathcal{W}$ to $\mathcal{N}$; see Section 3 for details. This definition is rather standard [24]. We say that an M-boundary $\mathcal{N}$ is a *realisation* of the Poisson boundary of M, if every bounded harmonic function h on the state space V of M can be obtained by integration of a bounded function $\hat{h} \in L^\infty(\mathcal{N})$, where this $\hat{h}$ is unique up to modification on a null-set, and conversely, for every $\hat{h} \in L^\infty(\mathcal{N})$ the function $f : V \to \mathbb{R}$ defined by $z \mapsto \int_{\mathcal{N}} \hat{h}(\eta) d\nu_z(\eta)$ is bounded and harmonic.

Recall the definition of a sharp function from Section 1.1. In Section 3 we introduce 'intersection' and 'union' operations between pairs or families of sharp harmonic functions using probabilistic intuition, and show that the family $\mathcal{S}$ of sharp harmonic functions of a Markov chain is closed under these operations. Thus $\mathcal{S}$ carries a σ-algebra structure, except that there is no ground set. Theorem 1.3 can be interpreted in the following way: any M-boundary $\mathcal{N}$ that can be used as the ground set for this σ-algebra structure, with its measures agreeing with corresponding probabilities defined with respect to sharp harmonic functions, is a realisation of the Poisson boundary. A bit more precisely, if $\mathcal{N}$ is *faithful* to $\mathcal{S}$, that is, if for every measurable subset X of $\mathcal{N}$ there is $s \in \mathcal{S}$ such that almost surely random walk ends up in X if and only if the values of s along its trajectory converge to 1, then $\mathcal{N}$ is a realisation of the Poisson boundary.

A direct implication of the comparison of the σ-algebra structure of $\mathcal{S}$ to that of $\mathcal{N}$ is that any two realisations of the Poisson boundary of M are cryptomorphic, see Corollary 3.5 (this fact will not be surprising to experts).

Given a Markov chain it is often easy to guess a realisation of its Poisson boundary, most often in the form of some topological space naturally associated to the chain, e.g. the boundary of a compactification, but it is much harder to prove that the guess is correct. Our case, the boundary of a square tiling, is such an example. Other examples include the end-compactification of a tree, and the hyperbolic boundary of a hyperbolic group [24]. The following tool, abstracted from the proof of Theorem 1.1 via Theorem 1.3, may be helpful in further such cases. A *topological G-boundary* of a graph G is a topological space $(\mathcal{N}, \mathcal{O})$ endowed with a 'projection' $\tau : V \to \mathcal{O}$ so that $\tau(Z^n)$ converges to a point in $\mathcal{N}$ for almost every random walk trajectory Z^n, and there is a Borel-measurable function $\tau^* : \mathcal{W} \to \mathcal{N}$ mapping almost every $(Z^n) \in \mathcal{W}$ to $\lim_n \tau(Z^n)$. Note that defining ν_z by $O \in \mathcal{O} \mapsto \mu_z(\tau^{*-1}(O))$ turns $\mathcal{N}$ into a G-boundary for $f = \tau^*$. We say that $(\mathcal{N}, \tau)$ is *layered*, if there is a sequence $(G_n)_{n \in \mathbb{N}}$ of finite subgraphs of G with $\bigcup G_n = G$ the boundaries $(B_n)_{n \in \mathbb{N}}$ of which satisfy $\mu_z^n(b) = \nu_z \circ \tau(b)$ for every $b \in B_n$, where μ_z^n denotes the exit distribution of G_n for random walk from z.

Corollary 1.4. *Let G be a transient graph and let $\mathcal{N}$ be a layered topological*

G-boundary with projection τ. If for every sharp harmonic function s we have $\lim_{m,n} \nu(\tau(F_m) \triangle \tau(F_n)) = 0$, where $F_i := \{b \in B_i \mid s(b) > 1/2\}$, then $\mathcal{N}$ endowed with the measures (ν_z) as above is a realisation of the Poisson boundary of random walk on G.

The following observation, proved in Section 6.1, is one of the main tools in the proof of Theorem 1.1 and might be of independent interest. Here, a graph is considered as a metric space where every edge is a copy of the real unit interval, and so each square of the tiling is foliated into horizontal intervals, one for each inner point of the corresponding edge.

Observation 1.5. *Let G be a plane, uniquely absorbing graph and consider its square tiling T of the cylinder K. For any circle $L \subset K$ parallel to the base of K, let B be the set of points of G the images of which lie in L. Then the widths $w(T(b))$, $b \in B$ of these images coincide with the exit probabilities of standard brownian motion on G started at the reference vertex o and killed at B.*

Our proof of Theorem 1.1 applies to a larger class of graphs where the degrees are not necessarily bounded, see Corollary 7.8. Such graphs have attracted a lot of interest lately, see [23] and references therein.

We prove Theorem 1.3 in Sections 3 and 4 and Theorem 1.1 in Section 7, after constructing the square tiling in Section 5 and proving some general properties, including interesting probabilistic interpretations of its geography like Observation 1.5, in Section 6. We deduce Corollary 1.4 in Section 7.3.

2 Preliminaries

Let $G = (V, E)$ be a graph fixed throughout this section, where $V = V(G)$ is its set of vertices and $E = E(G)$ its set of edges. A *walk* on G is a (possibly finite) sequence $(v_n)_{n \in \mathbb{N}}$ of elements of V such that v_i is always connected to v_{i+1} by an edge. More generally, we define a walk on the state space V of a Markov chain in a similar manner, where we might or might not demand that the transition probabilities $p_{v_i \to v_{i+1}}$ be positive.

A *plane graph* is a graph G endowed with a fixed embedding in the plane $\mathbb{R}^2$; more formally, G is a plane graph if $V(G) \subset \mathbb{R}^2$ and each edge $e \in E(G)$ is an arc between its two vertices that does not meet any other vertices or edges. A graph is *planar* if it admits an embedding in $\mathbb{R}^2$. Note that a given planar graph can be isomorphic (in the graph-therotic sense) to various plane graphs that cannot necessarily be mapped onto each other via a homeomorphism of $\mathbb{R}^2$.

A plane graph $G \subset \mathbb{R}^2$ is *uniquely absorbing*, if for every finite subgraph G_0 there is exactly one connected component D of $\mathbb{R}^2 \backslash G_0$ that is *absorbing*, that is, random walk on G visits $G \backslash D$ only finitely many times with positive probability (in particular, G is transient). Uniquely absorbing graphs are precisely those admitting a square tiling. Example classes include all transient 1-ended planar graphs, which includes the bounded-degree 1-ended planar Gromov-hyperbolic graphs, all transient trees and, more generally, all transient graphs that can be embedded in $\mathbb{R}^2$ without accumulation points of vertices.

We assume that G is endowed with an assignment $c : E \to \mathbb{R}_+$ of *conductances* to its edges, which are used to determine the behaviour of random walk on G as follows.

6

A *random walk* on G begins at some vertex and when at vertex x, traverses one of the edges $\overrightarrow{xy}$ incident to x according to the probability distribution

$$p_{x \to y} := \frac{c(xy)}{\pi_x}, \tag{1}$$

where $\pi_x := \sum_{y \in N(x)} c(xy)$ and $N(x)$ denotes the set of *neighbours* of x, that is, the vertices connected to x by an edge. When $c = 1$, which is the case most often considered, π_x coincides with the *degree* of x, and we have *simple random walk*, i.e. y is chosen according to the uniform distribution on $N(x)$.

Formally, there are two standard ways of rigorously formalising random walk as a probability space: the first is as a Markov chain in the obvious way. The second, and the one that we will adhere to in this paper, is by considering random walk on G as a measurable space $(\mathcal{W}, \Pi)$, endowed with a family of measures $(\mu_z)_{z \in V}$ indexed by the vertices of G, where $\mathcal{W}$ is the set of 1-way infinite walks on G, called *path space*, Π is the σ-algebra on $\mathcal{W}$ generated by the *cylinder sets*, i.e. subsets of $\mathcal{W}$ comprising all walks having a common finite initial subwalk, and μ_z is the probability measure on $(\mathcal{W}, \Pi)$ corresponding to fixing z as the starting vertex. Note that once z is fixed, (1) uniquely determines μ_z; see [39] for details.

For convenience, we will also assume that every Markov chain M in this paper is formally given in the above form, that is, a choice of a (possibly random) starting point o and a family of measures $(\mu_z)_{z \in V}$ on path space $(\mathcal{W}, \Pi)$, where V denotes the state space of M and μ_z is the law of M conditioning on the starting point being z. In other words, we formalize M as a random walk on V; in contrast to the random walk on a graph defined above, such random walk need not be reversible, see [32].

It will be convenient in some cases to think of our graph G as a metric space constructed as follows. Start with the discrete set V, and for every edge $xy \in E$ join x to y by an isometric copy of a real interval of length $1/c(e)$ (the *resistance* of e). Then, one can consider a brownian motion on this space (as defined e.g. in [5, 33]) that behaves locally like standard brownian motion on $\mathbb{R}$, and it turns out that the sequence of distinct vertices visited by this brownian motion has the same distribution as random walk governed by (1).

A function $h : V \to \mathbb{R}$ on the vertex-set of a graph G, or more generally on the state space V of a Markov chain, is *harmonic at* $x \in V$ if it satisfies

$$h(x) = \sum_{y \in V} p_{x \to y} h(y), \tag{2}$$

where again $p_{x \to y}$ denotes the transition probability, and it is called *harmonic* if it is harmonic at every x.

A fundamental property of harmonic functions is that their values inside a set are determined by the values at the boundary of that set; to make this more precise, let B be a subset of V, and x a vertex such that random walk from x visits B almost surely; for example, B could be the boundary of a ball of G containing x. Then, letting b be the first vertex of B visited by random walk from x, we have

$$h(x) = \mathbb{E}\left[h(b)\right]. \tag{3}$$

In other words, the boundary values of a harmonic function uniquely determine the function.

We say that random walk *hits* B at $b \in B$ if the first element of B it visit is b.

The following fact, which is a special case of Levy's zero-one law [27], will come in handy in many occasions

> For every μ-measurable event $A \subseteq \mathcal{W}$, random walk Z^n satisfies $\lim_n \mu_{Z^n}(A) = \Vdash_A$ almost surely. In particular, $\mu_{Z^n}(A)$ converges to 0 or 1. $\qquad(4)$

Here, $\Vdash_A$ is as usual the characteristic function (from $\mathcal{W}$ to $0,1$) of A. In other words, if we observe at each step i of our random walk the probability p_i that A will occur given the current position (the past does not matter because of the Markov property), then almost surely p_i will converge to 1 and A will occur or p_i will converge to 0 and A will not occur.

For a walk $W \in \mathcal{W}$ define the *shift* $t(W)$ to be the walk obtained from W by deleting the first step. An event $A \subseteq \mathcal{W}$ is called if *tail event* if $t(A) = A$ (we could be less strict here and write $\mu(t(A) \triangle A) = 0$ instead, where $\triangle$ denotes symmetric difference).

3   Sharp functions and the Poisson boundary

Let $M = Z^1, Z^2, \ldots$ be an irreducible Markov chain with state space V fixed throughout this section. Call a harmonic function $h : V \to \mathbb{R}^+$ *sharp*, if the range of h is $[0, 1]$ and $\lim_n h(Z^n)$ equals 0 or 1 almost surely (the limit exists almost surely by the bounded martingale convergence theorem). Note that whether h is sharp or not does not depend on the (possibly random) starting point o, for if r is any other element of V then, by irreducibility, the probability to visit r from o is positive. Let $\mathcal{S} = \mathcal{S}(V)$ denote the set of sharp harmonic functions on V.

Lemma 3.1. *If $h(z) : V \to [0, 1]$ equals the probability that random walk from z will satisfy a tail event A for every z, then h is a sharp harmonic function.*

Proof. The fact that h is harmonic follows immediately from the fact that A is a tail event and the Markov property. Sharpness follows from (4). $\qquad\square$

Let $s \in \mathcal{S}$ and fix $z \in V$. Define the real valued random variable X_n to be $s(Z^n)$ where Z^n denotes random walk from z. Define the random variable X by letting $X = \lim s(Z^n)$ if this limit exists (which it does almost surely by the bounded martingale convergence theorem), and $X = 0$ otherwise. Since almost sure convergence implies weak convergence [29], we immediately obtain

> The sequence X_n converges weakly to X. $\qquad(5)$

Alternatively, let $(B_n)_{n \in \mathbb{N}}$ be a sequence of subsets of V such that our Markov chain visits every B_n almost surely, and let $X_n = s(Z^{t_n})$, where t_n is the first time t such that $Z^t \in B_n$ (if no such t exists, which happens with probability 0, we can let $X_n = 0$ to make sure X_n is always defined). Then $\lim s(X_n)$ still exists almost surely by the choice of $(B_n)_{n \in \mathbb{N}}$, and so (5) also holds in this case.

As s is harmonic, we have

$$s(z) = \mathbb{E}[X_n]. \tag{6}$$

Given $z \in V$, let 1^s be the event, in the path space $\mathcal{W}_z$, that for random walk Z^n from z we have $\lim s(Z^n) = 1$. Note that this event is μ_z-measurable. Define the event 0^s similarly.

Corollary 3.2. *If s is a sharp harmonic function, then $s(z) = \mu_z(1^s)$ for every vertex z.*

Proof. Since the X_n are uniformly bounded, their weak convergence (5) implies convergence in L^1, that is, $\lim_n \mathbb{E}[X_n] = \mathbb{E}[X]$, where X is defined as in (5). Since s is sharp, $\mathbb{E}[X]$ equals the probability to have $\lim_n s(Z^n) = 1$ by the definition of X. Combined with (6) and the definition of 'sharp' this completes our proof. $\square$

Corollary 3.3. *If s is a sharp harmonic function that is not identically 0, then for every $\epsilon > 0$ there is $z \in V$ with $s(z) > 1 - \epsilon$.*

Proof. Since $s \neq 0$, there is o with $s(o) > 0$. By Corollary 3.2 the probability to have $\lim_n s(Z^n) = 1$ for random walk Z^n from o is positive, in particular there are vertices z with $s(z)$ arbitrarily close to 1. $\square$

Given a sequence $(s_n)_{n \in \mathbb{N}}$ of sharp harmonic functions, we define their *union* $\bigcup s_i$ by $z \mapsto \mu_z(\bigcup 1^{s_i})$ and their *intersection* $\bigcap s_i$ by $z \mapsto \mu_z(\bigcap 1^{s_i})$. The *complement* s_1^c of s_1 is the function $1 - s_1$; note that, by Corollary 3.2, $s_1^c(x) = \mu_x(0^{s_1})$.

Lemma 3.4. *Let (s_i) be a sequence of sharp harmonic functions. Then the functions*

(i) $\bigcup s_i$,

(ii) s_1^c,

(iii) $\bigcap s_i$,

are also harmonic and sharp.

Proof. Since all these functions are probabilities of tail events of random walk, they are harmonic and sharp by Lemma 3.1. $\square$

This means that the family of sharp harmonic functions satisfies the axioms of a σ-algebra, except that it is formally not a family of subsets of a given set. One way to intuitively interpret Theorem 1.3 is to say that if a measure space $\mathcal{N}$ has the 'same' σ-algebra structure as the family of sharp harmonic functions of G, and some obvious requirements are fulfilled, then $\mathcal{N}$ can be identified with the Poisson boundary of G. Let us make this idea more precise.

A *M-boundary*, or a *G-boundary* if M happens to be random walk on a graph G, is a measurable space $(\mathcal{N}, \Sigma)$ endowed with a family of probability measures $\{\nu_z, z \in V\}$ and a measurable, measure preserving, shift-invariant function $f : \mathcal{W} \to \mathcal{N}$, where $\mathcal{W} := \bigcup_{z \in V} \mathcal{W}_z$ is the set of 1-way infinite walks in G. Here, we say that f is *measure preserving* if $\nu_z(X) = \mu_z(f^{-1}(X))$ for every $z \in V$ and $X \in \Sigma$. The term *shift-invariant* means that if the walk W'

is obtained from W by skipping the first step, then $f(W) = f(W')$; in other words, f can be thought of as a function from the set of equivalence classes with respect to the shift (called 'ergodic components' in [24]) to $\mathcal{N}$.

The fact that f is shift-invariant and measure preserving implies

$$\nu_z = \sum_{y \sim z} p_{zy} \nu_y. \tag{7}$$

Indeed, since f is shift-invariant, random walk from z 'finishes' in $X \in \Sigma$ if and only if its subwalk after the first step finishes in X, and so we have $\mu_z(f^{-1}(X)) = \sum_{y \sim z} p_{zy} \mu_y(f^{-1}(X))$. Since we are demanding that f is measure preserving, our assertion follows.

This means in particular that the measures ν_z are pairwise equivalent when G is connected: if $\nu_z(X) > 0$ for some set $X \in \Sigma$, then $\nu_y(X) > 0$ for any neighbour y of z.

Let s be a sharp harmonic function. We say that $\mathcal{N}$ is *faithful* to s, if there is $X \in \Sigma$ such that $\mu_z(1^s \triangle f^{-1}(X)) = 0$ for every $z \in V$, where $\triangle$ denotes symmetric difference. Note that such a set X is unique up to modification by a null-set of ν.

For every measurable subset X of $\mathcal{N}$, the function $s = s_X$ defined by $z \mapsto \nu_z(X)$ is harmonic and sharp by Lemma 3.1. We claim that

$$\mu(1^s \triangle f^{-1}(X)) = 0 \tag{8}$$

To see this, set $\Phi := f^{-1}(X) \triangle 1^s$, and recall that both $f^{-1}(X)$ and 1^s are μ-measurable events in path space, hence so is Φ. Now note that $\mu_z(f^{-1}(X) \triangle 1^s) = \mu_z(1^s \cap f^{-1}(X^c)) + \mu_z(0^s \cap f^{-1}(X))$,
where we used the fact that $\mu_z(1^s \cup 0^s) = 1$ as s is sharp. It is easy to see that both these summands equal 0 using the definition of s, the Markov property of random walk, and the fact that f is measure preserving.

Note that if $\mathcal{N}$ is faithful to every sharp harmonic function, then combined with (8) this means that, up to perturbations by null-sets, the correspondence between measurable subsets of $\mathcal{N}$ and sharp harmonic functions is one-to-one. Combined with Theorem 1.3, this observation yields

Corollary 3.5. *If $(\mathcal{N}, \Sigma)$ is a realisation of the Poisson boundary of a graph G, then there is a bijection σ from the set of equivalence classes of Σ (where two elements are equivalent if they differ by a null-set) to the set of sharp harmonic functions of G such that $\nu_z(X) = \mu_z(1_z^{\sigma([X])})$ for every $X \in \Sigma$ and $z \in V(G)$. Thus any two realisations of the Poisson boundary of G are cryptomorphic.*

4 Proof of the Poisson boundary criterion

We start by proving the easier direction of Theorem 1.3, namely that if $\mathcal{N}$ is a realisation of the Poisson boundary then it is faithful to every sharp harmonic function s. For this, given s let $\hat{s} \in L^\infty(\mathcal{N})$ be such that $s(z) = \int_{\mathcal{N}} \hat{s}(\eta) d\nu_z(\eta)$ for every z.

We claim that $\hat{s}$ equals 0 or 1 almost everywhere on $\mathcal{N}$. Indeed, since f is measure preserving, it suffices to prove that $\mu(f^{-1}(X) \cup f^{-1}(Y)) = 1$, where $X := \{\eta \mid \hat{s}(\eta) = 1\}$ and $Y := \{\eta \mid \hat{s}(\eta) = 0\}$. For $\epsilon \in (0, 1)$, let $X_\epsilon := \{\eta \mid \hat{s}(\eta) \in [\epsilon, 1 - \epsilon]\}$, and define the function $s_\epsilon : V \to [0, 1]$ by $z \mapsto \mu_z(f^{-1}(X_\epsilon))$.

By Lemma 3.1 s_ϵ is harmonic and sharp. Now note that if $s_\epsilon(W^n)$ converges to 1 for some walk W^n, then $s(W^n)$ does not converge to 0 or 1. But as s is sharp, this occurs with probability 0 for our random walk. Since s_ϵ is sharp, this implies that $s_\epsilon(W^n)$ converges to 0 for almost every random walk W^n. But now Levy's zero-one law (4) implies that $\mu_z(f^{-1}(X_\epsilon)) = 0$ for every z. Since this holds for every ϵ, our claim follows.

As $\hat{s}$ equals 0 or 1 almost everywhere on $\mathcal{N}$, we have $s(z) = \nu_z(X)$ by the choice of $\hat{s}$ and X. Thus (8) yields $\mu_z(1^s \triangle f^{-1}(X)) = 0$, which means that $\mathcal{N}$ is faithful to s as desired.

4.1 Splitting $\mathcal{N}$ according to the values of h

In this section collect some lemmas that will be useful in the proof of the other direction of Theorem 1.3. The reader may choose to skip to Section 4.2 at this point and come back later.

We denote the set of bounded harmonic functions of G by $BH(G)$.

Lemma 4.1. *Let G be a transient network. For every $h \in BH(G)$ and every $a < b \in \mathbb{R}$, the function $h_{[a,b)}$ defined by $z \mapsto \mu_z(\lim_n h(Z^n) \in [a,b))$ is harmonic and sharp.*

Proof. To begin with, it is easy to check that the event $\{\lim_n h(Z^n) \in [a,b)\}$ is measurable. The assertion now follows from Lemma 3.1. □

Definition 4.2. *Let $(\mathcal{N}, (\nu_z)_{z \in V})$ be an M-boundary that is faithful to every sharp harmonic function, and let $h \in BH(G)$. Recall that for every measurable $X \subseteq \mathcal{N}$, there is a sharp function s_X such that $f^{-1}(X) = 1^{s_X} \triangle \Phi$ where Φ is a null-set in path space (8). Given a bounded interval $[a,b) \subset \mathbb{R}$, we define the sharp harmonic function $y = y_{[a,b)} := s_X \cap h_{[a,b)}$ where $h_{[a,b)}$ is as in Lemma 4.1 and intersection as in Lemma 3.4. Since $\mathcal{N}$ is faithful to y, there is a set $Y \subseteq \mathcal{N}$ such that $\mu(1^y \triangle f^{-1}(Y)) = 0$, and we let $X \restriction_{[a,b)}$ denote such a set Y. By Corollary 3.2 we have*

$$\nu_z(Y) = \mu_z(f^{-1}(Y)) = \mu_z(1^y) = y(z) \tag{9}$$

Lemma 4.3. *Let $(\mathcal{N}, (\nu_z)_{z \in V})$ be an M-boundary that is faithful to every sharp harmonic function, and let $h \in BH(G)$. Let $a_0 < a_1 < \ldots a_k < a_{k+1} \in \mathbb{R}$ be points such that the range of h is contained in (a_0, a_{k+1}), then $\nu_z(X) = \sum_{0 \le i \le k} \nu_z(X \restriction_{[a_i, a_{i+1})})$ (hence X equals $\bigcup X \restriction_{[a_i, a_{i+1})}$ up to a null-set).*

Proof. By (9) we have $\sum \nu_z(X \restriction_{[a_i, a_{i+1})}) = \sum_i y_{[a_i, a_{i+1})}(z)$. By our definitions, we have

$$y_{[a_i, a_{i+1})}(z) = s(z) \cap h_{[a_i, a_{i+1})} = \mu_z(1^s \cap 1^{h_{[a_i, a_{i+1})}}) =$$
$$\mu_z(1^s \cap \{\lim_n h_{[a_i, a_{i+1})}(W^n) = 1\})$$

where $s(\cdot) := \nu.(X)$. Now recall that $h_{[a_i, a_{i+1})}$ is a probability, namely of the event that random walk W'^m from (the random vertex) W^n will have its h values converge to $[a_i, a_{i+1})$. Now note that the distribution of W'^m is the same as that of the continuation of W_n after the nth step. This means that $h_{[a_i, a_{i+1})}(W^n)$ equals the probability that W^n itself displays this behaviour.

Applying Levy's zero-one law (4) to the latter probability, we can thus deduce that the last expression above equals

$$\mu_z(1^s \cap \{\lim_n h(W^n) \in [a_i, a_{i+1})\}).$$

Plugging this into the above sum, we obtain

$$\sum \nu_z(X \upharpoonright_{[a_i, a_{i+1})}) = \sum_i \mu_z(1^s \cap \{\lim_n h(W^n) \in [a_i, a_{i+1})\}).$$

The latter sum however equals $\mu_z(1^s)$ since, by the bounded martingale convergence theorem, with probability 1 exactly one of the events $\{\lim_n h(W^n) \in [a_i, a_{i+1})\}$ occurs. Finally, recall that $\mu_z(1^s) = \mu_z(f^{-1}(X)) = \nu_z(X)$ by (8) and the fact that f is measure preserving. $\qquad\square$

4.2 Main proof of Theorem 1.3

We proceed with the proof of the backward direction of Theorem 1.3, that if $\mathcal{N}$ is faithful to every sharp harmonic function then it is a realisation of the Poisson boundary.

Let us start with the easier assertion we have to prove, namely that for every $\hat{h} \in L^\infty(\mathcal{N})$ the function $f : V \to \mathbb{R}$ defined by $z \mapsto \int_{\mathcal{N}} \hat{h}(\eta) d\nu_z(\eta)$ is bounded and harmonic.

It is immediate from its definition that the range of f is contained in the range of $\hat{h}$, and so f is bounded. The fact that f is harmonic follows easily from (7).

Next, we prove that for every $h \in BH(G)$ there is $\hat{h} : \mathcal{N} \to [0, 1]$ such that for every $z \in V$ we have $h(z) = \int_{\mathcal{N}} \hat{h}(\eta) d\nu_z(\eta)$, which is the core of Theorem 1.3. Assume without loss of generality that the range of h is the interval $[0, 1]$.

Recall that for every $0 < a < b < 1$ and any measurable $X \subseteq \mathcal{N}$, we can define the measurable set $X \upharpoonright_{[a,b)}$ (Definition 4.2). Using this we can, for every $z \in V$, and any measurable $X \subseteq \mathcal{N}$, induce a measure ν_z^X on $[0, 1]$ by letting $\nu_z^X(I) = \nu_z(X \upharpoonright_I)$ for every subinterval I of $[0, 1]$ and extending to all Borel subsets of $[0, 1]$ using Caratheodory's extension theorem. Now let

$$H_z(X) := \int_{[0,1]} a\nu_z^X(da). \tag{10}$$

It is straightforward to check that H_z is a measure on $\mathcal{N}$. Easily, H_z is uniformly continuous to ν_z. Thus we can let

$$R_z(\eta) = \frac{\partial H_z}{\partial \nu_z}(\eta)$$

be the corresponding Radon-Nikodym derivative. The range of R_z is contained (up to a null-set) in the closure of the range of h. Thus $R_z \in L^\infty(\mathcal{N})$.

Recall that we would like to find $\hat{h} : \mathcal{N} \to [0, 1]$ such that for every $z \in V$ we have $h(z) = \int_{\mathcal{N}} \hat{h}(\eta) d\nu_z(\eta)$. Thus it suffices to prove the following two claims.

Claim 1: For every z, o we have $R_z(\eta) = R_o(\eta)$ for almost every η.

This allows us to define $\hat{h} := R_o$ for a fixed $o \in V$. Recall that $R_o \in L^\infty(\mathcal{N})$.

Claim 2: For every $z \in V$ we have $h(z) = \int_{\mathcal{N}} dH_z = \int_{\mathcal{N}} R_z(\eta) d\nu_z(\eta) = \int_{\mathcal{N}} \hat{h}(\eta) d\nu_z(\eta)$.

To prove Claim 1, suppose to the contrary that there is $X \subseteq \mathcal{N}$ of positive measure such that $R_z(\eta) > R_o(\eta) + \epsilon$ for some $\epsilon > 0$ and every $\eta \in X$. By Lemma 4.3, we can decompose $\mathcal{N}$ into a union $\bigcup_{0 \leq i \leq k} \mathcal{N} \restriction_{[a_i, a_{i+1})}$ of measurable subsets, with $a_0 = 0$ and $a_{k+1} = 1$, where we are free to choose the a_i as we wish. So let us choose them in such a way that $a_{i+1} - a_i < \epsilon$ for every i.

Now note that, by the definition of H_z, we have $\frac{H_z(\mathcal{N} \restriction_{[a_i, a_{i+1})})}{\nu_z(\mathcal{N} \restriction_{[a_i, a_{i+1})})} \in [a_i, a_{i+1})$; indeed, if $I \cap J = \emptyset$ then $\nu((\mathcal{N} \restriction_I) \restriction_J) = 0$, and so $\nu_z^{\mathcal{N} \restriction_{[a_i, a_{i+1})}}$ is supported on $[a_i, a_{i+1})$. Even more, if Y is any measurable subset of $\mathcal{N} \restriction_{[a_i, a_{i+1})}$ of positive measure, we also have $\frac{H_z(Y)}{\nu_z(Y)} \in [a_i, a_{i+1})$. Thus $R_z(\eta) \in [a_i, a_{i+1})$ for almost every $\eta \in \mathcal{N} \restriction_{[a_i, a_{i+1})}$. But as this holds for every z, and $a_{i+1} - a_i < \epsilon$, this means that $|R_z(\eta) - R_o(\eta)| < \epsilon$ for almost every $\eta \in \mathcal{N}$, contradicting the existence of X as defined above. This proves Claim 1.

To prove Claim 2, we have to show that $h(z) = \int_{\mathcal{N}} dH_z = H_z(\mathcal{N}) := \int_{[0,1]} a \nu_z^{\mathcal{N}}(da)$. Since h is harmonic, we have $h(z) = \int_{b \in V} h(b) \mu_z^n(b)$ for every n, where μ_z^n denotes the distribution of the nth step of random walk from z. Easily, the latter sum equals $\int_{[0,1]} a \mu_z^n(da)$ by a double-counting argument, where, with a slight abuse of notation, we treat μ_z^n as a probability measure on $[0,1]$ by making the convention $\mu_z^n(da) := \mu_z^n(\{b \in V \mid h(b) \in da\})$. Comparing the latter integral with the one above, we see that it suffices to find a family $\mathcal{I}$ of intervals of $[0,1]$ that is a basis for its topology and $\lim \mu_z^n(da) = \nu_z^{\mathcal{N}}(da)$ for every interval $da \in \mathcal{I}$.

For this, call a number $a \in \mathbb{R}$ h-singular, if $\mu_o(\lim_n h(Z_n) = a) > 0$. Let $\mathcal{I}$ be the family of intervals contained in $[0,1]$ the endpoints of which are not h-singular. Since there are at most countably many h-singular points, $\mathcal{I}$ is clearly a basis of $[0,1]$.

To see that for every $da \in \mathcal{I}$ we have $\lim \mu_z^n(da) = \nu_z^{\mathcal{N}}(da)$, recall that $\nu_z^{\mathcal{N}}(da) = \nu_z(\mathcal{N} \restriction da) = \mu_z(\{\lim f(W^n) \in da\})$ where we used (9), and note that $\mu_z(\{\lim f(W^n) \in da\}) \leq \liminf_n \mu_z^n(da)$ because, as the endpoints of da are not h-singular, subject to $\{\lim f(W^n) \in da\}$ our random walk almost surely visits $da = \{b \in B^n \mid h(b) \in da\}$ for almost every n. We claim that, conversely, $\nu_z^{\mathcal{N}}(da) \geq \limsup_n \mu_z^n(da)$. To see this, note that $\mu_z^n(da) + \mu_z^n(da^c) = 1$, where da^c is the complement $[0,1] \backslash da$ of da, because μ_z^n is a probability measure. By Lemma 4.3 we have $\nu_z^{\mathcal{N}}(da) + \nu_z^{\mathcal{N}}(da^c) = 1$ as well. Thus, applying the above arguments to da^c instead of da, which yields $\nu_z^{\mathcal{N}}(da^c) \leq \liminf_n \mu_z^n(da^c)$, we conclude that $\nu_z^{\mathcal{N}}(da) \geq \limsup_n \mu_z^n(da)$ as claimed. This means that $\lim \mu_z^n(da)$ exists and equals $\nu_z^{\mathcal{N}}(da)$. This proves Claim 2, completing the proof of the existence of the desired function $\hat{h}$.

It remains to check that $\hat{h}$ is unique up to modification on a null-set. If this is not the case, then there is another candidate $\hat{h}'$ such that, for every $z \in V$, we have

$$h(z) = \int_{\mathcal{N}} \hat{h}(\eta) d\nu_z(\eta) = \int_{\mathcal{N}} \hat{h}'(\eta) d\nu_z(\eta).$$

Define the function $\hat{k}(\eta) := \hat{h}(\eta) - \hat{h}'(\eta)$ on $\mathcal{N}$, and note that $k(z) := \int_{\mathcal{N}} \hat{k}(\eta) d\nu_z(\eta) = 0$ for every z. Now if $\hat{h}$ does not coincide with $\hat{h}'$ ν-almost everywhere, there is some $\epsilon > 0$ such that the measurable set $X := \{\eta \in \mathcal{N} \mid \hat{k}(\eta) > \epsilon\}$ is not a null-set.

Let $s_X(z) := \nu_z(X) = \mu_z(f^{-1}(X))$, which is a sharp harmonic function by Lemma 3.1. Thus, by Corollary 3.3, there is $x \in V$ with $s_X(x) = \nu_x(X) > 1 - \epsilon'$ for any $\epsilon' > 0$ we choose. But this means that $k(x) > \epsilon(1 - \epsilon') - \epsilon \inf \hat{k}$ by the choice of X. Choosing ϵ' large enough compared to $\epsilon \inf \hat{k}$, we obtain a contradiction to $k = 0$ that completes the proof.

5 Construction of the tiling

In this section we show how any plane transient graph can be associated with a tiling of the cylinder $K := \mathbb{R}/\mathbb{Z} \times [0,1]$ with squares, or rectangles if the edges of G have various resistances. Our construction follows the lines of [8], but we will be pointing out many properties of this tiling that we will need later. Thus this section could be useful to the reader already acquainted with [8].

5.1 The Random Walk flow

Fix a vertex $o \in V$ and for every vertex $v \in V$ let $h(v)$ be the probability $p_v(o)$ that random walk from v will ever reach o. Thus $h(o) = 1$. We will use $h(v)$ as the 'height' coordinate of v in the construction of the square tiling in the next section.

Recall that the *Green function* $G(x, y)$ is defined as the expected number of visits to y by random walk from x. Let

$$h'(v) := \frac{\pi_o G(o, v)}{G(o, o) \pi_v},$$

where as usual $\pi_x := \sum_{y \in N(x)} c(xy)$.

We claim that

$$h'(v) = h(v). \tag{11}$$

Indeed this is a consequence of the reversibility of our random walk: it is well-known, and not hard to prove (see [32, Exercise 2.1]), that $\pi_v G(v, o) = \pi_o G(o, v)$. Observing that $G(v, o) = p_v(o) G(o, o)$ now immediately yields (11).

It is no loss of generality to assume that the constant $\frac{\pi_o}{G(o,o)}$ appearing in the definition of h' equals 1: multiplying the conductances c by a constant does not affect the behaviour of our random walk, and hence $G(o, o)$, and so we can achieve $\pi_o = G(o, o)$ by multiplying with the appropriate constant. Thus, from now on we can assume that

$$h(v) = \frac{G(o,v)}{\pi_v} = \frac{1}{\pi_v} \mathbb{E}\left[\# \text{ of visits to } v \text{ by random walk from } o\right]. \tag{12}$$

A *directed edge* $\overrightarrow{xy}$ of G is an ordered pair (x, y) of vertices such that $\{x, y\}$ is an edge of G. Define the *random walk flow* to be the function $w(\vec{e})$ on the set of directed edges $\vec{e}$ of G equaling the expected net number of traversals of $\vec{e}$ by random walk from o:

$$w(\overrightarrow{xy}) := G(o, x) p_{x \to y} - G(o, y) p_{y \to x}.$$

Since our random walk is transient, $w(\vec{e})$ is always finite. Note that, by (12) and (1), we have

$$w(\overrightarrow{xy}) = c(xy)(h(x) - h(y)). \tag{13}$$

In electrical network terminology, (13) says that the pair h, w satisfies Ohm's law. Thus w is 'antisymmetric', i.e. $w(\overrightarrow{xy}) = -w(\overrightarrow{yx})$. Moreover, w is a *flow* from o to infinity, by which we mean that it satisfies the following conservation condition, known as Kirchhoff's node law, at every vertex x other than o:

$$w^*(x) := \sum_{y \in N(x)} w(\overrightarrow{xy}) = 0, \tag{14}$$

where $N(x)$ is the set of neighbours of x. This is equivalent to saying that h is harmonic at every vertex except o.

The following fact can be found in [32, Exercise 2.87].

Lemma 5.1. *For almost every random walk (Z^n), we have $\lim_n h(Z^n) = 0$.*

Note that up to now we did not use the planarity of G, so all above statements hold for an arbitrarily transient graph.

5.2 The dual graph G^*

Let $G = (V, E)$ be a planar graph, and fix a proper embedding of G into the plane $\mathbb{R}^2$. Our graph can now be considered as a *plane graph*, in other words, V is now a subset of $\mathbb{R}^2$ and E a set of arcs in $\mathbb{R}^2$ each joining two points in V. It is a standard fact that one can associate with G a further plane graph $G^* = (V^*, E^*)$, called the (geometric) *dual* of G, having the following properties:

(i) Every face of G contains precisely one vertex of G^* and vice versa;

(ii) There is a bijection $e \mapsto e^*$ from E to E^* such that $e \cap G^* = e^* \cap G$ consists of precisely one point, namely a point at which the edges e, e^* meet.

Note that if G^* is a geometric dual of G then G is a geometric dual of G^*, but we will not need this fact. For example, the geometric dual of a hexagonal lattice is a triangular lattice with 6 triangles meeting at every vertex. See [14] for more details on dual graphs.

The orientability of the plane allows us to extend the bijection $e \mapsto e^*$ to a bijection between the directed edges of G and G^* in such a way that if C is a directed cycle of G and $\vec{E}(C)$ its set of directed edges, then $\{\vec{e}^* \mid \vec{e} \in \vec{E}(C)\}$ is a *directed cut*, that is, it coincides with the set of edges from a subset A of V to $V \setminus A$, directed from the endvertex in A to the endvertex in $V \setminus A$ (where A is the set of vertices of G^* contained in one side of C).

5.3 The Tiling

Given a plane, transient, and uniquely absorbing graph G, we now construct a tiling of the cylinder $K := \mathbb{R}/\mathbb{Z} \times [0, 1]$ by associating to each edge of N a rectangle $R_e \subseteq K$, with sides parallel to the boundary of K.

In order to specify R_e we will use four real coordinates, corresponding to the endvertices x, y of e in G and the endvertices x', y' of e^* in the dual G^*: two of these coordinates will be the values $h(x), h(y)$, where h is the function from Section 5.1 defined by means of the random walk flow; these values will be used

as 'height' coordinates. We now specify a 'width' function w on the vertices of G^* to be used for the other coordinates.

Fix an arbitrary vertex ζ of G^* and set $w(\zeta) = 0$. For every other vertex $z \in V(G^*)$, pick a z–ζ path $P_z = z_0 z_1 \ldots z_k$, where $z_0 = z$ and $z_k = \zeta$, and let $w(z) = \sum_{i<k} w(\overrightarrow{z_i z_{i+1}}^*) \mod 1$, where $\overrightarrow{z_i z_{i+1}}^*$ denotes the directed edge $\vec{e}$ of G such that $\vec{e}^* = \overrightarrow{z_i z_{i+1}}$ (recall the remark on orientability at the end of the previous section).

This value $w(z)$ does not depend on the choice of P_z, but only on the endpoint z; this fact is a consequence of a well-known duality between Kirchhoff's node law on a plane network G and Kirchhoff's cycle law on its dual G^*. To be more precise, if $C = \vec{e_1}\vec{e_2}\ldots\vec{e_k}$ is a directed cycle in G^* such that the set of vertices U of G contained in one of the sides of C is finite, then $\sum w(\vec{e_i}) = \sum_{x \in U} w^*(x)$. Using this, and the fact that G is uniquely absorbing, it is not hard to check that the latter sum equals a multiple $k\eta$ of the total flow $\eta := w^*(o)$ out of o, where k is the 'winding number' of C around o; see [8, Lemma 3.2] for a detailed proof. As we are assuming that $\eta = 1$, our definition of $w(z)$ does indeed not depend on the choice of P_z.

Having defined $w : V(G^*) \to \mathbb{R}/\mathbb{Z}$, we can now specify the rectangle R_e corresponding to an edge e as above in our tiling: R_e is one of the two rectangles in K bounded between the horizontal lines $h = h(x)$ and $h = h(y)$ and the perpendicular lines $w(x')$ and $w(y')$. To decide which of the two, orient e from its endvertex of lower h value into the one with higher value, recall that this induces an orientation of e^*, and choose that rectangle in which the w values increase as we move from the initial vertex of e^* to the terminal vertex inside the rectangle.

This completes the definition of the tiling, which from now on we denote by $T = T_{G,o}$. Formally, T can be defined as a function on E mapping each e to a rectangle $R_e \subset K$, but it will be more convenient in the sequel to assume T to be a function on G, seen as a metric space (recall the discussion in Section 2), mapping any interior point p of an edge e to the maximal horizontal interval at height $h(p)$ contained R_e, and any vertex x to the maximal horizontal interval at height $h(p)$ contained in $\bigcup_{x \in e} R_e$.

Note that $T(G)$ does not meet the base $\mathcal{C} := \mathbb{R}/\mathbb{Z} \times \{0\}$ of our cylinder K. It is the main aim of this paper to show that $\mathcal{C}$, to be thought of as the *boundary* of K, is a realisation of the Poisson boundary of G.

Let us point out some further properties of our tiling T. By (13), the aspect ratio of R_e equals the conductance of e:

$$\frac{w(e)}{dh(e)} = c(e). \tag{15}$$

In particular, if c is identically 1 then we obtain a square tiling.

Let $\vec{E}(x)$ be the set of directed edges emanating from vertex x. Let $E^+(x) \subseteq \vec{E}(x)$ be the set of directed edges $\overrightarrow{xy}$ with $h(y) \geq h(x)$ and $E^-(x)$ be the set of directed edges $\overrightarrow{xy}$ with $h(y) < h(x)$. Note that the random walk flow flows into x along the edges in $E^+(x)$ and out of x along the edges in $E^-(x)$.

For a vertex x, define $w(x)$ to be the *width* of its image in T; that is, $w(x) := \sum_{\vec{e} \in E^-(x)} w(\vec{e}) = \sum_{\vec{e} \in E^+(x)} w(\vec{e}) = 1/2 \sum_{\vec{e} \in E(x)} |w(\vec{e})|$ ($w(x)$ should not be confused with $w^*(x)$). By the definitions, we have

$$w(x) = \mathbb{E}\left[\text{net \# of particles arriving to } x \text{ from above}\right]. \tag{16}$$

The energy of a function $u : V \to \mathbb{R}$ is defined by $E(u) := \sum_{xy \in E(G)} (u(x) - u(y))^2$ (or more generally, $\sum_{xy \in E(G)} (u(x) - u(y))^2 / c(xy)$ if the conductances are non-constant). Note that for our height function h, $E(h)$ equals the area of K, which is 1, as the contribution of each edge e to $E(h)$ is the area of R_e by definition. This, combined with the bounded degree condition, implies

Lemma 5.2. *If π is bounded from above then w converges to 0, i.e.*

$$\lim_n w(x_n) = 0 \text{ for every enumeration } (x_n) \text{ of } V. \tag{17}$$

Proof. If there are infinitely many x_i with $w(x_i) \geq \epsilon > 0$, then each of them is incident with an edge e_i with $w(e_i) > \epsilon/D$ where D is the maximal degree. But each such edge contributes more than $(\epsilon/D)^2$ to $E(h)$, contradicting the fact that the latter is finite. $\qquad\square$

We formulated Theorem 1.1 for graphs of bounded degree only, but most of our proof applies to all graphs satisfying the weaker condition (17), the only exception being the part on convergence to the boundary (Section 6.3).

6 Probabilistic interpretations of the geography of the tiling cylinder

Our next observation is that if we modify our graph by subdividing some edge e into two edges the total resistance of which equals the resistance $1/c(e)$ of e, then the tiling remains practically unchanged. Note that in our metric space model of G introduced in Section 2, such an operation is tantamount to declaring some interior point x of the arc e to be a vertex while leaving the metric unchanged.

Lemma 6.1. *Let G be a transient plane graph and x be an interior point of an edge e of G, and let G' be the plane graph obtained from G by declaring x to be a vertex. Then $T_o(G')$ can be obtained from T_G by cutting the rectangle R_e into two rectangles along the horizontal line $h(x)$.*

Proof. Observe that the functions h and w of Section 5.1 remain unchanged at every old vertex or edge. This can be checked directly, or by using the fact that our random walk can be obtained as the sequence of distinct vertices visited by brownian motion on G, defining $h(x)$ to be the probability that brownian motion from any point x of N will ever reach o, noticing that this definition of h coincides with that of Section 5.1 on V, and observing that brownian motion cannot tell the difference between G and G'.

The assertion now is an immediate consequence of the construction of the tilings. $\qquad\square$

Another way of stating Lemma 6.1 is that if we consider our tiling T to be a function on G as described above, then declaring interior edge points to be vertices does not change T.

17

6.1 Parallel circles

We are now going to consider a sequence $(G_n)_{n \in \mathbb{N}}$ of subgraphs of G with nice properties using our tiling T of the cylinder K: let $(l_n)_{n \in \mathbb{N}}$, $0 < l_n < 1$, be a sequence of numbers monotonely converging to 0, define the *parallel circle* L_n to be the set of points of K at height l_n, and let G_n be the subgraph mapped by T to the strip between the top of the cylinder K and L_n; in other words, G_n is the subgraph spanned by the vertices x with $h(x) \geq l_n$. Similarly, for $m < n \leq \infty$ define the *strip* G_n^m to be the subgraph of G (and G_n) spanned by the vertices x with $l_m \geq h(x) \geq l_n$, where we set $l_\infty = 0$. Note that $V(G_n) \cup V(G_\infty^n) = V(G)$.

We may, and will, assume that each L_n meets G at vertices only (thus $G_n \cup G_\infty^n = G$ and $G_n \cap G_\infty^n = L_n$); for if e is an edge dissected by L_n, then we can, by Lemma 6.1, put a dummy vertex at the point x in the interior of e with $h(x) = l_n$. Let B_n be the preimage of L_n, i.e. the set of (possibly dummy) vertices b with $h(b) = l_n$. We claim that

random walk from o visits every B_n almost surely. $\qquad$ (18)

Indeed, by Lemma 5.1 the heights of the vertices visited by random walk almost surely converge to 0; in other words, random walk converges to the base of K. Thus it almost surely visits every parallel circle L_n. Note that (18) is trivially true if each G_n is a finite graph, which is the case for a large class of planar graphs G but not always.

The advantage of our assumption that each L_n meets G at vertices only can be seen in the following lemma. Denote by $\mu_o^n(b)$ the exit distribution (harmonic measure) of random walk from o killed at B_n.

Lemma 6.2. *For every n and every $b \in B_n$, $w(b) = \mu_o^n(b)$.*

This important observation does not presuppose G to be planar; in the planar case we can visualise w as width with respect to T, but w is well-defined, and Lemma 6.2 true, for any transient graph. Note that Lemma 6.2 is a reformulation of Observation 1.5 when considering the brownian motion mentioned in Section 2; the advantage of that formulation is that it is not affected by the dummy vertices we introduced above (recall the discussion in the proof of Lemma 6.1).

We will also show that $w(b), b \in B_n$ also coincides with the distribution of the last vertex of B_n visited by random walk from o (and so the distribution of the first vertex of B_n visited coincides with that of the last).

Before we prove all this, we need to introduce a further tool.

If G is any finite (or recurrent) planar graph, then it is possible to construct a tiling as we did in the infinite case, except that we now have to stop the random walk at some point because if we do not, then $h(v)$ (the probability to reach o from v) will be 1 for every v, yielding a trivial tiling. One possibility is to stop our random walk upon its first visit to a fixed vertex o'; this is in fact what Brooks et. al. did when they first introduced square tilings [8, 10].

We will follow a slightly different approach: we will construct a tiling T_{G_n} corresponding to G_n by starting our random walk at o and stopping it upon its first visit to the set B_n (recall that such a visit exists almost surely (18)). For this, we repeat the construction of Section 5 except that we now define $h_n(v)$ to be the probability for random walk from v to reach o before B_n. Note that

now we can reformulate (12) as

$h_n(v) = \frac{1}{\pi_v} \mathbb{E} [\# \text{ of visits to } v \text{ by random walk from } o \text{ stopped at } B_n]$.

The rest of the construction of the tiling remains the same. Note that we now have $h_n(v) = 0$ for every $v \in B_n$; in fact this tiling T_{G_n} can be obtained from our tiling T_G of G by linearly stretching the part of T_G bounded between the lines $h = l_0$ and $h = l_n$ to bring the latter line to height zero:

Lemma 6.3. T_{G_n} can be obtained from T_G by the following transformation: if T_G maps a point p of G_n to the interval $I_p \times \{h(p)\}$, then T_{G_n} maps p to the interval $I_p \times \{\frac{h(p)-l_n}{1-l_n}\}$.

Proof. Let $f(v) := \frac{h(v)-l_n}{1-l_n}$. Note that the formula $h_n(v) = f(v)$ is trivially correct for $v \in B_n \cup \{o\}$, and that both h and h_n are harmonic on the rest of G_n. Moreover, f is harmonic since h is, because the former is a linear transformation of the latter. Recall that random walk from o visits B_n almost surely (18). Thus, since the values of a harmonic function are uniquely determined by its boundary values (recall (3) and the remark after it), h_n must coincide with f everywhere.

Recall that w is uniquely determined by h (13). Thus the width of each vertex is the same in the two tilings, and as we are using the same embedding in both cases, the position of the interval corresponding to any vertex is also the same. $\qquad\square$

Lemma 6.3 easily implies Lemma 6.2: recall that $w(b)$ equals the net number of particles ariving to v from above by (16). But in the killed random walk we used in the construction of T_{G_n}, this number coincides with the exit probability $\mu_o^n(b)$. Since by Lemma 6.3 $w(b)$ is the same in T_{G_n} and T_G, Lemma 6.2 follows.

For $m \in \mathbb{N}$ and a walk W from o, we define an m-*subwalk* of W to be a maximal subwalk of W that has no interior vertex in B_m. Note that W is the concatenation of all its m-subwalks. The first of them is called the *initial m-subwalk*, and if there is a last one it is called the *final m-subwalk*; all others are *interior m-subwalks*, and start and end in B_m but do not visit B_m in between.

Recall that $w(\vec{e})$ was defined as the expected net number of traversals of $\vec{e}$ by random walk from o. Let $w^m(\vec{e})$ be the contribution to w by the initial and final m-subwalks of random walk from o, i.e. we let W be a random walk from o and let $w^m(\vec{e})$ be the expected net number of times that the initial m-subwalk of W goes through $\vec{e}$ plus the expected net number of times that the final m-subwalk of W goes through $\vec{e}$. We define $w_n^m(\vec{e})$ similarly for $n > m$, except that W is stopped upon its first visit to B_n.

Lemma 6.4. $w(\vec{e}) = w^m(\vec{e}) = w_n^m(\vec{e})$ for every $\vec{e} \in \vec{E}(G)$.

Proof. By the definitions, an equivalent statement is that for every directed edge $\vec{e}$, the interior m-subwalks of random walk from o traverse e in each direction equally often in expectation, where random walk is stopped upon its first visit to B_n when considering the second equation.

To prove that these traversals indeed cancel out, let Q be a walk that starts and ends in B_m and does not visit B_m in between, in other words, a candidate for an interior m-subwalk of our random walk; for the proof of $w = W_n^m$ we also demand Q to avoid B_n here. Let $a, z \in B_m$ be the endvertices of Q. Let p_{az} denote the probability that random walk from a will coincide with Q up to its first re-visit to B_m, and conversely, let p_{za} denote the probability that random

walk from z will coincide with the inverse Q^- of Q up to its first re-visit to B_m. It is well-known, and straightforward to prove using (1), that $\pi(a)p_{az} = \pi(z)p_{za}$.

Now recall that $h(a) = G(o,a)/\pi(a)$ by (12), and similarly $h(z) = G(o,z)/\pi(z)$, where $G(o,a)$ denotes the expected number of visits to a. Since $a, z \in B_m$, we have $h(a) = h(z)$. Now note that the expected number of times that an m-subwalk of our random walk coincides with Q equals $G(o,a)p_{az}$ by linearity of expectation and the Markov property. Similarly, the expected number of times that an m-subwalk of our random walk coincides with Q^- is $G(o,z)p_{za}$. Putting all this together we have

$$G(o,a)p_{az} = h(a)\pi(a)p_{az} = h(a)\pi(z)p_{za} = h(z)\pi(z)p_{za} = G(o,z)p_{za}.$$

Thus, the contribution of the pair of walks Q, Q^- to $w(\vec{e})$ is zero, for even if they contain the edge $\vec{e}$, their contributions cancel out in expectation. Since all walks that are candidates for an interior m-subwalk of our random walk can be organised in such pairs, their overall contribution to $w(\vec{e})$ is zero as claimed. $\square$

This implies the following assertion, which complements Lemma 6.2.

Corollary 6.5. *For every m and every $b \in B_m$, $w(b)$ coincides with the probability distribution of the last visit of random walk (from o) to B_m. This remains true if the random walk is stopped upon its first visit to B_n for $n > m$.*

(The reader might be upset at this point for our use of the letter m here and the letter n in Lemma 6.2, but will be grateful for this when reading Section 7.)

Proof. The difference between the distributions of the first and last visits to B_m is determined by the behaviour of the interior m-subwalks of random walk from o. But by Lemma 6.4 the influences of these m-subwalks cancel out, and so Lemma 6.2 implies our claim. For the second sentence we use Lemma 6.3. $\square$

6.2 Meridians

Having studied the nice probabilistic behaviour of parallel circles, we now turn our attention to *meridians*, i.e. lines in K with a constant width coordinate. We will prove an assertion which is, in a sense, dual to Lemma 6.4: for every meridian M, the expected number of particles crossing M from left to right equals the expected number of particles crossing M from right to left.

Some care needs to be taken before we can make such an assertion, since if our random walk is currently at a vertex x the span $T(x)$ of which (recall that $T(x)$ is a horizontal interval in K) is dissected by M, then we cannot say whether the particle is on the right or left of M. To amend this, we assume that at each step n of our random walk, an additional random experiment is made to choose a random point P^n in the span $T(Z^n)$ of the current position Z^n uniformly among all points in $T(Z^n)$, and all these experiments are independent from each other and $Z^i, i < n$, and we think of P^n as the position of our random walk at step n. With this assumption in mind, we can now state the following

Lemma 6.6. *For every vertex x and every meridian M meeting $T(x)$, the expected number of particles crossing M from left to right at x equals the expected number of particles crossing M from right to left at x.*

Proof. To begin with, suppose for simplicity that M does not dissect any square $R_e \subset K$ associated to an edge e incident with x. Thus locally the position of M is like in the left of Figure 2. By the construction of T, this implies that for the two vertices r, s of G^* lying in the faces of G separating the edges mapped to the left of M from those mapped to its right we have $w(r) = w(s) = w(M)$, where $w(M)$ denotes the common width coordinate of the points in M (Figure 2 right).
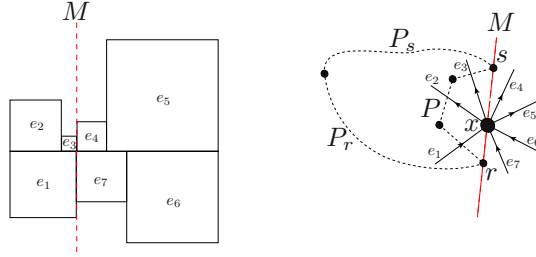

Figure 2: The local situation at a vertex x whose span $T(x)$ is dissected by a meridian M in the tiling and the graph.

Now let P be the r–s path in G^* comprising the set of edges E^r incident with x on the 'left' of M. Recall that the values $w(r), w(s)$ where specified by choosing arbitrary paths P_r, P_s from r, s respectively to the reference vertex ζ of G^*, and adding the $w(\vec{e})$ values along such a path. Now note that we can obtain a candidate for P_r by prefixing P_s by P. But since $w(r) = w(s)$, this implies that $\sum_{\vec{e} \in P} w(\vec{e}) = 0$. The definition of the random walk flow w now implies that the net flow into x along edges in E^r is zero. Likewise, the net flow into x along the 'right' edges E^l is zero. Our claim that the expected number of particles crossing M from left to right at x equals the expected number of particles crossing M from right to left now follows from Kirchhoff's node law (14) and the Sand-bucket Lemma below (brown sand corresponds to particles coming to x from the left; shovel them to B if they leave x from the right).

It is straightforward to adapt this argument to the general case where M does dissect some square R_e; we leave the details to the reader. $\square$

Lemma 6.7 (Sand-bucket Lemma)**.** *Alice has two (not necessarily equally full) buckets of sand A and B, where A contains only brown sand and B only white sand. If she puts one shovelful from A to B, mixes arbitrarily, then puts back one shovelful from B to A, then the amount (measured by volume) of white sand in A equal the amount of brown sand in B.*

6.3 Convergence of random walk to the boundary

In this section we provide a proof of the almost sure convergence of the image of the trajectory or random walk under T to a point in the boundary $\mathcal{C}$ of K which is maybe simpler than the proof of [8]. This proof is the only occasion in this paper where the bounded degree condition cannot be replaced by the weaker (17). Let $D := \max_{x \in V} d(x)$ be the highest degree in G.

To begin with, recall that the image of our random walk converges to the boundary $\mathcal{C}$ of K by Lemma 5.1 (we do not need the bounded degree condition for this). So it just remains to show that the width coordinates converge too.

For this, let X be an interval of $\mathcal{C}$, and let M_L, M_R be the two meridians corresponding to the endpoints of X. Let L be the set of vertices x such that M_L intersects $T(x)$, and define R similarly. Note that unless M_L, M_R happen to meet some vertex at an endpoint of its span —a possibility we can exclude since there are countably many such meridians— L, R are the vertex sets of two 1-way infinite paths of G.

We claim that the expected number of times that our random walk alternates between L and R is finite. This implies that almost surely the number of such alternations is finite, and as this holds for every X, convergence to the boundary follows.

It is not hard to see that the expected number of times that random walk from o goes from L to R equals $\sum_{x \in L} G(o, x) \sum_{y \in R} p_{xy}$, where as usual $G(o, x)$ denotes the expected number of visits to x, and p_{xy} denotes the probability that random walk from x exits $L \cup R$ at a given vertex $y \in R$. Thus, we can write the above claim as follows:

$$\sum_{x \in L, y \in L} G(o, x) p_{xy} < \infty. \tag{19}$$

We are going to prove this exploiting the relations between random walks and electrical networks.

Recall that $G(o, x) = h(x)\pi_x$ (12), and so the above sum equals $\sum_{x \in L, y \in L} h(x)\pi_x p_{xy}$. The quantity $\pi_x p_{xy}$ was shown in [21] to equal the 'effective conductance' C_{xy} between x and y when the network is finite. Effective conductance is closely related to energy (recall the physical formula $E = I^2 R = I^2/C$), and we are going to exploit this fact using an argument similar to the proof of Lemma 5.2. We cannot directly apply the results of [21] as the graph is infinite in our case, however, there is an easy way around this: we can consider an increasing sequence of finite subgraphs $H_1 \subseteq H_2 \subseteq \ldots \subseteq G$ such that $\bigcup H_n = G$, apply the results of [21] to H_n and take a limit.

To make this more precise, let p_{xy}^n be the probability that random walk on H_n from a vertex $x \in L \cap H_n$ exits $(L \cup R) \cap H_n$ at $y \in R \cap H_n$, and let $C_{xy}^n := \pi_x p_{xy}^n$. It is a well known property of electrical networks, called Reyleigh's monotonicity law, that C_{xy}^n is monotone increasing with n [32]. Thus $\lim p_{xy}^n = \lim C_{xy}^n/\pi_x$ exists. It is now not hard to prove that $p_{xy} \le \lim p_{xy}^n$ using coupling and elementary probabilistic arguments. Thus, to prove (19) it suffices to find a uniform upper bound for $\sum_n := \sum_{x \in L, y \in L} C_{xy}^n$ (where we used the fact that h is bounded). It is shown in [21] that the above sum $\sum_n$ equals the energy E_n of the harmonic function (or the electric current) on H_n with boundary conditions 1 at $L \cap H_n$ and 0 at $R \cap H_n$. Thus, all we need to do is to prove that the E_n are bounded.

To achieve this, we will invoke the following well known fact: on a finite network, the harmonic function with given boundary conditions minimises energy among all function satisfying the boundary conditions. This means that it suffices in our case to find arbitrary functions v^n with boundary conditions 1 at $L \cap H_n$ and 0 at $R \cap H_n$ and uniformly bounded energies $E(v^n)$. We will do so by defining a single function $v : V \to \mathbb{R}$ with finite energy $E(v)$ equaling 1 on L and 0 on R, and letting v^n be its restriction to H_n. Since $E(v^n) \le E(v)$, the $E(v^n)$ will indeed be bounded.

To define this v, let $v(x) = 1$ if $x \in L$, let $v(x) = o$ if $x \in R$, and let $v(x)$ be any value in $T(x)$ otherwise, where we think of $T(x)$ as an interval of $\mathbb{R}/Z$

by projecting it to the base of K. To check that $E(v) < \infty$, we will compare $E(v)$ with the energy $E(h)$ of the height function h we used in the construction of the tiling. Recall that $E(h)$ equals the area of K, which is finite.

Now consider an edge $e = wz$, set $a := |v(w) - v(z)|$ and recall that the contribution of e to $E(v)$ is $(v(w) - v(z))^2 = a^2$. Note that as e joins w to z, the construction of the tiling implies that $T(w), T(z)$ meet when projected vertically. Combined with our choice of v, this implies that at least one of $T(w), T(z)$ has length at least $a/2$; assume without loss of generality this is $T(z)$. Now since $d(z)$ is bounded by D, there is at least one edge $f = f(e)$ incident with z such that the corresponding square R_f in the tiling has side length at least $a/2D$. This means that the area of R_f is at least $a^2/(2D)^2$, which is a constant times the contribution a^2 of e to $E(v)$. Now note that each edge f of G can be considered as $f(e)$ for a bounded number of edges e, namely those sharing an endvertex with f. This easily implies that $E(v)$ is bounded by a constant times $E(h)$, and so $E(v)$ is finite since $E(h)$ is.

This proves our claim that the expected number of times that our random walk alternates between L and R is finite. Applying this to a countable family of pairs L_i, R_i such that the corresponding intervals generate the topology of $\mathcal{C}$, and combining this with the almost sure convergence of h along a random walk trajectory, immediately yields

Theorem 6.8. *For random walk Z^n on G, the intervals $T(Z^n)$ almost surely converge to a point in $\mathcal{C}$.*

This allows us to view $\mathcal{C}$ as a G-boundary as defined in Section 3: define $f : \mathcal{W} \to \mathcal{C}$ as follows. For a walk W on G such that $T(W_n)$ converges to a point $p \in \mathcal{C}$, we let $f(W) = p$; otherwise we let $f(W)$ be a fixed arbitrary point of $\mathcal{C}$ (such walks form a null-set, so it does not really matter). It is easy to check that f is measurable and shift-invariant. Naturally, we define the measures ν_z on $\mathcal{C}$ by $\nu_z(X) := \mu_z(f^{-1}(X))$, making sure that f is measure preserving.

Combining Theorem 6.8 with Lemma 6.2, it is an easy exercise to deduce the following fact, already proved in [8], that will be useful later

Corollary 6.9. *ν_o is equal to Lebesgue measure on $\mathcal{C}$.*

7 Faithfulness of the boundary to the sharp harmonic functions

In this section we prove that the boundary $\mathcal{C}$ of the tiling constructed above is faithful to every sharp harmonic function, which we plug into Theorem 1.3 to complete the proof that $\mathcal{C}$ is a realisation of the Poisson boundary of our planar graph G (Theorem 1.1).

Let s be such a function, fixed throughout this section.

7.1 Some basic random walk lemmas

We first collect some basic general lemmas on random walks that will be useful later. The results of this section hold for general Markov chains, but we will only formulate them for random walk.

Let $\mathcal{A}$ be a tail event of our random walk, i.e. an event not depending on the first n steps for every n. (The only kind of event we will later consider is the event 1^s that $s(Z_n)$ converges to 1, where s is our fixed sharp harmonic function.)

For $r \in (0, 1/2]$, let

$$A_r := \{v \in V \mid \mathbb{P}[\mathcal{A}] > 1 - r\} \text{ and}$$
$$Z_r := \{v \in V \mid \mathbb{P}[\mathcal{A}] < r\}.$$

Note that $A_r \cap Z_r = \emptyset$ for every such r.

By Corollary 3.2, if we let $\mathcal{A} := \mathcal{X}$ then we have $A_r = \{v \in V \mid s(v) > 1 - r\}$ and $Z_r = \{v \in V \mid s(v) < r\}$.

Lemma 7.1. *For every $\epsilon, \delta \in (0, 1/2]$, and every $v \in A_\epsilon$, we have $\mathbb{P}_v[\text{visit } V \backslash A_\delta] < \epsilon/\delta$. Similarly, for every $v \in Z_\epsilon$, we have $\mathbb{P}_v[\text{visit } V \backslash Z_\delta] < \epsilon/\delta$.*

Proof. Start a random walk (Z_n) at v, and consider a stopping time τ at the first visit to $V \backslash A_\delta$. If τ is finite, let $z = Z_\tau$. Since $z \notin A_\delta$, we have $\mathbb{P}[\mathcal{X}_z^c] \geq \delta$ by the definition of A_δ. Thus, subject to visiting $V \backslash A_\delta$, the event $\mathcal{A}$ fails with probability at least δ since it is a tail event. But $\mathcal{A}$ fails with probability less than ϵ because $v \in A_\epsilon$, and so $\mathbb{P}_v[\text{visit } V \backslash A_\delta] < \epsilon/\delta$ as claimed.

The second assertion follows by the same arguments applied to the complement of $\mathcal{A}$. $\square$

Corollary 7.2. *If random walk from $v \in A_\epsilon$ (respectively, $v \in Z_\epsilon$) visits a set $W \subset V$ with probability at least κ, then there is a v–W path all vertices of which lie in $A_{\epsilon/\kappa}$ (resp. $Z_{\epsilon/\kappa}$).*

Proof. Apply Lemma 7.1 with $\delta = \epsilon/\kappa$. $\square$

Lemma 7.3. *For every $r \in (0, 1/2)$, the probability that random walk alternates k times between A_r and Z_r is less than $(2r)^k$.*

Proof. Applying Lemma 7.1 for $\epsilon = r$ and $\delta = 1/2$, we deduce that for every $v \in A_r$, $\mathbb{P}_v[\text{visit } V \backslash A_{1/2}] < 2r$. Since $r < 1/2$ we have $Z_r \subseteq V \backslash A_{1/2}$, and so $\mathbb{P}_v[\text{visit } Z_r] < 2r$. Similarly, if $v \in Z_r$, then $\mathbb{P}_v[\text{visit } A_r] < 2r$.

Using this, and the Markov property of random walk, we can prove by induction that the probability for random walk from any vertex to alternate k times between A_r and Z_r is less than $(2r)^k$ as claimed. $\square$

Lemma 7.4. *Let $x, y \in B_m$ and $\epsilon, \delta \leq 1/2$. There is no x–y path in $A_\delta \cap G - G_m$ separating a vertex $z \in B_m \cap Z_\epsilon$ from C.*

Proof. If there is such a path, then random walk from z in $G - G_m$ would have to visit it with probability 1. But by the Markov property of Z_n, this would imply that $z \in A_\delta$, contradicting $z \in Z_\epsilon$ since $A_\delta \cap Z_\epsilon = \emptyset$. $\square$

7.2 The main lemma

Let s be a sharp harmonic function of G. For every $n \in N$, we let

$$F_n := \{b \in B_n \mid s(b) > 1/2\}$$
$$F_n' := \{b \in B_n \mid s(b) < 1/2\}.$$

The following lemma is the main ingredient of the proof that $\mathcal{C}$ is a realisation of the Poisson boundary. Let $\overline{F_n}$ denote the projection of $T(F_n)$ to the base $\mathcal{C}$ of K.

Lemma 7.5. $\lim_{m,n} w(\overline{F_n} \triangle \overline{F_m}) = 0$.

Proof. Suppose to the contrary there is $\lambda > 0$ such that for arbitrarily large m there is an arbitrarily large $n > m$ with $w(\overline{F_n} \triangle \overline{F_m}) > 2\lambda$.

Since s is sharp, the weak convergence of (5) —and the remark after it— easily implies that there are sequences $(\epsilon_n)_{n \in \mathbb{N}}, (\zeta_n)_{n \in \mathbb{N}}$ such that $\lim \epsilon_i = 0$, $\lim \zeta_i = 0$ and, defining X_i by

$$X_i := \{b \in B_i \mid s(b) > 1 - \epsilon_i\}, \tag{20}$$

we have $w(F_i \triangle X_i) < \zeta_i$. Note that $X_i = F_i \cap A_{\epsilon_i}$ by Corollary 3.2, where A_r is defined as in the previous section for $\mathcal{A} := 1^s$. We may assume that both $(\epsilon_n)_{n \in \mathbb{N}}, (\zeta_n)_{n \in \mathbb{N}}$ are decreasing. Fix such sequences, and define Y_i similarly to X_i by $Y_i := \{b \in B_i \mid s(b) < \epsilon_i\}$. Assume that $\epsilon_1 < 1/2$, to ensure that X_i, Y_i are always disjoint and that $X_i \subseteq F_i$ and $Y_i \subseteq F_i'$. We may also assume that $w(F_i' \triangle Y_i) < \zeta_i$.

Let $\chi := \lim \mu_o^i(F_i) = \lim \mu_o^i(X_i)$, which exists by (5) and (3). Note that $1 - \chi = \lim \mu_o^i(X_i^c) = \lim \mu_o^i(Y_i)$. In other words, setting $W_i := B_i \backslash (X_i \cup Y_i)$, we have

$$\mu_o^i(W_i) \text{ converges to } 0. \tag{21}$$

Given λ as above, we can choose $m \in \mathbb{N}$ large enough that ϵ_m, ζ_m are very small compared to λ. We will later be able to specify how much smaller we want them to be. In order to avoid tedious delta-epsilontics, we introduce the notation $\Psi \simeq \Omega$, where Ψ, Ω are functions of m, or quantities bounded by functions of m, to express the assertion that $|\Psi - \Omega|$ is bounded from above by a function of m that converges to 0 as m goes to infinity. Similarly, the notation $\Psi \gtrsim \Omega$ means that $\Omega - \Psi$ is bounded above by such a function if it is positive. We will use this notation as a means of simplification, the point being that when we introduce an error that can be bounded by some function of ϵ_m, ζ_m, we do not have to write down that function explicitly to know that the error is small, if we know that we can make the error as small as we wish by choosing m accordingly large.

Note that $\epsilon_m, \zeta_m \simeq 0$ by their definition. This implies, for example, that $\epsilon_m / \lambda \simeq 0$ since λ is fixed, an observation we will use later. As another example, we have $\mu_o^m(X_m) \simeq \chi$ by the definition of χ. Similarly, we have $\mu_o^m(W_m) \simeq 0$ by (21), and $\mu_o^m(Y_m) \simeq 1 - \chi$. Even more, we have

$$\begin{aligned} &\sup_{n \geq m} \mu_o^n(X_n) \simeq \chi, \\ &\sup_{n \geq m} \mu_o^n(W_n) \simeq 0, \text{ and} \\ &\sup_{n \geq m} \mu_o^n(Y_n) \simeq 1 - \chi. \end{aligned} \tag{22}$$

By Lemma 6.2 we have $\mu_o^m(X_m) = \sum_{x \in X_m} w(x)$ and $\mu_o^n(X_n) = \sum_{x \in X_n} w(x)$. It thus follows from our assumption $w(\overline{X_n} \triangle \overline{X_m}) \simeq w(\overline{F_n} \triangle \overline{F_m}) > 2\lambda$ and (22) that the proportion of the $m \to n$ projection of $T(X_m)$ which is not in $T(X_n)$ is $\gtrsim \lambda$; to make this more precise, we denote by $x \downarrow_n$, where x is an element or a subset of X_m, the set of vertices $y \in B_n$ such that $T(y)$ and $T(x)$ meet when

25

vertically projected to $\mathcal{C}$. Then we have $w(X_m) - \sum_{y \in X_m \downarrow_n} w(y) \gtrsim \lambda$, where we used the fact that vertex widths converge to 0 as we move down the tiling (17).

By (21) we have $\mu_o^n(W_n) \simeq 0$, and so $\mu_o^n(X_m \downarrow_n) \simeq \mu_o^n(X_m \downarrow_n \cap X_n) + \mu_o^n(X_m \downarrow_n \cap Y_n)$. The above fact now implies that a $\gtrsim \lambda$ proportion of the $m \to n$ projection of $T(X_m)$ is in $T(Y_n)$, i.e.

$$\mu_o^n(X_m \downarrow_n \cap Y_n) \gtrsim \lambda. \tag{23}$$

We will now use Lemma 7.1 to show that for most vertices $b \in B_m$ there is a b–B_n path in G_n^m along which the values of s are close to $s(b)$. To make this more precise, define A_r, Z_r as in the previous section, putting $\mathcal{A} = 1^s$. Then Lemma 7.1 and the Markov property implies that, for any $\delta < 1/2$,

$$\mu_o\{\text{visit } A_{\epsilon_m} \text{ and then } V \backslash A_\delta\} \leq \mu_o\{\text{visit } A_{\epsilon_m}\} \frac{\epsilon_m}{\delta} \leq \frac{\epsilon_m}{\delta}. \tag{24}$$

Now let $\delta = \lambda$, and call a X_m–B_n path in $G_n^m \cap A_\delta$ an A-barrier, and a Y_m–B_n path in $G_n^m \cap Z_\delta$ a Z-barrier. Recall that $X_m = F_m \cap A_{\epsilon_m}$ and $Y_m = F_m' \cap Z_{\epsilon_m}$. Call a vertex *good* if it is in X_m and is incident with an A-barrier or it is in Y_m and incident with a Z-barrier; let $W_m' \supseteq W_m$ be the set of all vertices of B_m that are not good. We claim that

$$\sum_{x \in W_m'} w(x) \simeq 0. \tag{25}$$

Indeed, recall that, by Lemma 6.5, w coincides with the distribution of the last visit of our random walk to B_m. Now subject to this last visit being in $W_m' \cap X_m$, random walk will visit $V \backslash A_\delta$ with probability 1 because random walk eventually reaches B_n but there is no $(W_m' \cap X_m)$–B_n path in $G_n^m \cap A_\delta$, because such a path would be an A-barrier, from which we conclude $\mu_o\{\text{visit } W_m' \cap X_m \text{ and then } V \backslash A_\delta\} \geq \sum_{x \in W_m' \cap X_m} w(x)$. On the other hand, we have $\mu_o\{\text{visit } X_m \text{ and then } V \backslash A_\delta\} \leq \mu_o\{\text{visit } A_{\epsilon_m} \text{ and then } V \backslash A_\delta\}$, but the latter probability is bounded from above by ϵ_m/δ by (24). Putting the two inequalities together we obtain $\sum_{x \in W_m' \cap X_m} w(x) \leq \epsilon_m/\delta \simeq 0$. Similarly, we can prove that $\sum_{x \in W_m' \cap Y_m} w(x) \simeq 0$. As $W_m' = (W_m' \cap X_m) \cup (W_m' \cap Y_m) \cup W_m$, our claim follows.

We are now going to partition B_m into classes, called *clusters*, in such a way that every cluster is incident with a barrier, and vertices in the same cluster have similar s-values except for a subset of small width. For this, define first a *pre-cluster* to be either a maximal subset of $X_m \backslash W_m'$ spanning an interval of L_m in the tiling, or a maximal subset of $Y_m \backslash W_m'$ spanning an interval of B_m.

Note that every pre-cluster of the first type is, by definition, incident with an A-barrier, and every pre-cluster of the second type is incident with a Z-barrier.

Next, we associate each vertex in W_m' (these are precisely the vertices not belonging to any pre-cluster) to the first pre-cluster to their 'left' along L_m, and define a *quasi-cluster* to be the union of a pre-cluster with the set of vertices associated with it. Note that each quasi-cluster is incident with at least one A-barrier or Z-barrier, but not with both. This allows us to classify quasi-clusters into two types, called type A and type Z accordingly. Finally, we define a *cluster* to be the union of a maximal set of quasi-clusters that spans an interval of L_m and is not incident with both an A-barrier and a Z-barrier.

26

Note that clusters can be still classified into types as above, as each cluster is incident with either an A-barrier or a Z-barrier. Let Ξ_m denote the set of clusters of type A and Υ_m the set of clusters of type Z. By construction, no two clusters of the same type are adjacent as intervals of L_m.

A cluster of type A may in principle have more of its width in Y_m than X_m, but such clusters are rare when measured by w. For a typical cluster, its type agrees with the kind of most of its width:

$$w(\bigcup \Xi_m \backslash X_m) + w(\bigcup \Upsilon_m \backslash Y_m) \simeq 0, \tag{26}$$

because $w(W'_m) \simeq 0$ and $(\bigcup \Xi_m \backslash X_m) \cup (\bigcup \Upsilon_m \backslash Y_m) \subseteq W'_m$.

Combined with (23), this implies that the proportion of the $m \to n$ projection of the clusters of type A that is in Y_n is at least $\gtrsim \lambda$, i.e.

$$\sum_{x \in \bigcup \Xi_m} \mu_o^n(x \downarrow_n \cap Y_n) \gtrsim \lambda. \tag{27}$$

Let

$$\Xi'_m := \{c \in \Xi_m \mid \frac{w(c \downarrow_n \cap Y_n)}{w(c)} \geq \lambda/2\},$$

and note that Ξ'_m consists of those clusters of type A that have a relatively large contribution to (27). We claim that Ξ'_m is a relatively large set, namely, $w(\bigcup \Xi'_m) \gtrsim \frac{\lambda(1-\chi/2)}{1-\lambda/2}$. For this, let $P = w(\bigcup \Xi_m \downarrow_n \cap Y_n)$, and recall that $P \gtrsim \lambda$ by (27). Note that $w(\bigcup \Xi'_m) + \lambda/2(\chi - w(\Xi'_m)) \gtrsim P$ by the definition of Ξ'_m. Since $P \gtrsim \lambda$, an easy calculation proves our claim $w(\Xi'_m) \gtrsim \frac{\lambda(1-\chi/2)}{1-\lambda/2}$. Since $\chi \leq 1$, this implies that

$$w(\Xi'_m) \gtrsim \lambda/2. \tag{28}$$

Since the proportion of the $m \to n$ projection of the clusters of type A that is in Y_n is at least $\gtrsim \lambda$ (27), the results of Section 6.1 (Lemmas 6.5 and 6.2) imply that with positive probability our random walk has its last visit to B_m in some cluster c of type A and then diverts to hit B_n outside $c \downarrow_n$. Let us make this more precise. We say that random walk from o *goes astray*, if its final m-subwalk (this terminology was introduced in the proof of Lemma 6.3) starts at a cluster c and does not hit B_n in $c \downarrow_n$. We claim that

Conditioning random walk from o to have its last visit to B_m in Ξ'_m, it goes astray with probability $\gtrsim \lambda/2$. $\tag{29}$

To prove this, recall that the distribution of the last vertex on B_m visited (which coincides with the first vertex of the final m-subwalk) coincides with w by Lemma 6.5, and that the hitting distribution to B_n coincides with w (Lemma 6.2). Thus the definition of Ξ'_m implies that if our claim (29) is wrong, then the probability that the final m-subwalk starts in Ξ'_m and exits in Y_m is not $\simeq 0$. This probability however is bounded from above by $\mu_o\{\text{visit } X_m \text{ and then } Y_m\}$, which by (24) is at most $\frac{\epsilon_m}{1/2} \simeq 0$, which leads to a contradiction. This proves (29).

The fact that random walk goes astray with positive probability combined with Lemma 7.1 implies the existence of certain L_m–L_n paths with vertices in A_δ that our random walk can follow, which can be proved by arguments similar

to those that we used above to prove the existence of barriers. To make this more precise, let Ξ''_m be the set of those clusters c in Ξ'_m such that there is a path P^c_{ℓ} in the strip G^m_n with vertices in A_δ from c to the set $B_n \backslash c \downarrow_n$ of vertices not lying below c in the tiling (Figure 3). We claim that most of Ξ'_m is in Ξ''_m:
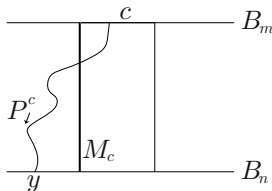

Figure 3: The path P^c_{ℓ}.

$$w(\bigcup \Xi'_m \backslash \bigcup \Xi''_m) \simeq 0. \tag{30}$$

To see this, note that subject to going astray having started the final m-subwalk in $\bigcup \Xi'_m \backslash \bigcup \Xi''_m$, our random walk visits $V \backslash A_\delta$ with probability 1, and so this happens with probability $\simeq 0$ by (24).

Note that every such path P^c_{ℓ} is an A-barrier by definition. But what is special about it, is that it must cross one of the meridians of c, in fact the part of the meridian contained in the strip G^m_n, which part we denote by M_c. A more careful examination of the proof of the existence of P^c_{ℓ} shows that we can strengthen (29) to demand that random walk goes astray crossing M_c.

Since $w(\Xi''_m) \gtrsim \lambda/2$ by (28) and (30), we have just proved that the final m-subwalk of random walk from o crosses an M_c as above with probability $\gtrsim \lambda^2/4$.

These crossings account for a flow of particles from one side of these meridians to the other. By Lemma 6.6, this flow has to be compensated for by other particles, crossing the meridian in the other direction. Such particles have to cross not only M_c, but also P^c_{ℓ} by a topological argument; in particular, they have to visit A_δ. To see that such particles do indeed have to cross P^c_{ℓ}, suppose for a moment that Figure 3 is not showing part of a cylinder, but part of a rectangular strip, i.e. we have cut the cylinder along a perpendicular line not visible in the figure. Let P be a path that starts in B_m, finishes in B_n, and crosses M_c from left to right. For our purposes, it is enough to consider such a P whose width coordinates grow monotonely, for if a particle crosses M_c in one direction and then back in the other direction then it has no compensating contribution to the aforementioned flow. Thus P starts on B_m on the left of M_c, and finishes on B_n on the right of M_c. But then P^c_{ℓ} separates the two endpoints of P, and so P has to cross P^c_{ℓ}.

Now since we are on a cylinder rather than a strip, the words 'left,right' make less sense. However, for this argument it makes little difference; for example, we can consider a covering of the cylinder K by an infinite strip, lift P, M_c and P^c_{ℓ} to the covering, and then apply the above argument.

Recall that only the initial and final m-subwalk can have a non-zero contribution to the expected number of traversals of an edge by Lemma 6.4, and so only the final m-subwalk can have a compensating contribution for the above flow since M_c lies below L_m. We can bound the probability with which the

particle starts its final m-subwalk in Y_m and displays this behaviour, i.e. crosses an M_c and hence visits A_δ, applying (24) with the roles of A and Z reversed: this probability is at most $\epsilon_m/\delta \simeq 0$.

Note however, that although a particle visiting Y_m crosses an M_c with low probability, it might be that when it does so it crosses the M_c of several clusters c, and so its expected contribution in compensating the above flow becomes considerable. But we can use Lemma 7.3 to bring this possibility under control as follows. Again, only the final m-subwalk is relevant. Now suppose that the final m-subwalk of random walk from o crosses the M_c of several $c \in \Xi''_m$, say k of them. Then, as no two clusters of type A can be adjacent as intervals of L_m, and as every cluster has a barrier of its own type, such a subwalk must, by a geometric argument, visit at least k A-barriers and $k-1$ Z-barriers alternatingly. Thus, such a subwalk alternates between A_δ and Z_δ k times. The probability of this event though can be bounded using Lemma 7.3: using an argument similar to the one we used to derive (24) from Lemma 7.1, we see that the expected number of alternating visits between A_δ and Z_δ is less than 2. Combining this fact with the above bound for the probability of even visiting a single A-barrier from Y_m, and using the Markov property of random walk, we deduce that the expected number of A-barriers crossed by final m-subwalks starting in Y_m, which equals the compensating contribution of Y_m to the above flow, is $\simeq 0$.

Now consider the case where the particle starts it final m-subwalk at $x \in X_m$. If x happens to be in Ξ''_m, then it cannot cross the A-barrier of its own cluster c in a compensating manner, i.e. coming into the vertical strip under c from outside, unless it travels all the way around the cylinder. In any case, in order for such a subwalk to have a compensating contribution, it has to cross the A-barrier of some other cluster $z \in \Xi''_m$. But then it would have to visit the Z-barriers of all the cluster of type Z lying between c and z, and there is at least one such cluster by construction. Using the same arguments as in the previous case, we deduce that the expected compensating contribution of X_m to the above flow is also less than $\simeq 0$.

Finally, random walk starts its final m-subwalk in W_m with probability $\simeq 0$, and again its compensating contribution is $\simeq 0$ too by similar arguments.

This proves that our alleged flow crossing the M_c cannot be compensated for in accordance with Lemma 6.6, yielding a contradiction that completes our proof. □

Note that we can replace F_m, F_n with X_m, X_n, defined as above (20), in the assertion of Lemma 7.5 because $w(F_i \triangle X_i) < \zeta_i$ and $\lim \zeta_i = 0$, which means that

$$\lim w(F_i \triangle X_i) = 0. \tag{31}$$

An (m,n)-*impurity* is a vertex in F_n that has an m-predecessor in F'_m, or a vertex in F'_n that has an m-predecessor in F_m, where an m-predecessor of y is a vertex $x \in B_m$ such that $y \in x \downarrow_n$. Let $M^m_n \subseteq B_n$ be the set of m, n-impurities. Using this terminology, Lemma 7.5 can be reformulated as follows

$$\lim_{m,n} w(M^m_n) = 0. \tag{32}$$

29

7.3 Faithfulness

Using (32) we can now show that $\mathcal{C}$ is faithful to $\mathcal{S}$.

Lemma 7.6. *For every sharp harmonic function s there is a measurable subset X of $\mathcal{C}$ such that $\mu_o(1_o^s \triangle f^{-1}(X)) = 0$.*

Proof. Let $\epsilon_i = 2^{-i}$ and apply Lemma 7.5, or (32), to obtain a sequence $(n_{\epsilon_i})_{i \in \mathbb{N}}$ such that for every $m, n \geq n_{\epsilon_i}$ we have $w(M_n^m) < 2^{-i}$. Since our choice of the levels l_n was arbitrary, we may assume without loss of generality that $n_{\epsilon_i} = i$, and we will make this assumption to avoid subscripts in our notation. Our choice of (ϵ_i) has the effect that $\sum_m := \sum_{n > m} w(M_n^m) < \infty$ for every m, and $\lim_m \sum_m = 0$. In other words, we have

$$\lim_m \mu_o(\{\text{exit some } n > m \text{ at an } (n-1, n)\text{-impurity}\}) = 0. \tag{33}$$

Recall the definition of F_i from the previous section, and let $\overline{F_i}$ denote the projection of $T(F_i)$ to the base $\mathcal{C}$ of K. Let X be the set of points eventually in $\overline{F_i}$, i.e.

$$X := \{x \in \mathcal{C} \mid \text{there is } k \text{ such that } x \in \overline{F_i} \text{ for every } i > k\},$$

and define Y similarly replacing F_i by F_i'.

Note that $X = \liminf_i \overline{F_i} := \bigcup_j \bigcap_{i \geq j} \overline{F_i}$. Thus X is a borel subset of C, hence ν-measurable. We will prove that X has the desired property $\mu_o(1_o^s \triangle f^{-1}(X)) = 0$.

Similarly to (33) it follows from the summability of (ϵ_i), and the definition of M_n^m, that $\lim_m \sum_{n > m} w(\overline{F_n} \triangle \overline{F_{n-1}}) = 0$. This implies

$$\lim_m w(\overline{F_m} \triangle X) = 0 \text{ and } \lim_m w(\overline{F_m'} \triangle Y) = 0 \tag{34}$$

because any point $p \in \overline{F_m} \triangle X$ lies in $\overline{F_n} \triangle \overline{F_{n-1}}$ for some $n > m$ by the definition of X.

Next, we claim that there is a sequence $(n_i)_{i \in \mathbb{N}}$ of levels such that

$$\sum_i \mu_o(\{\text{visit } F_{n_i} \text{ and then visit } F_{n_j}' \text{ for some } j > i\}) < \infty. \tag{35}$$

For this, recall that visiting F_{n_i} is almost the same event as visiting X_{n_i} by (31) (recall the definition (20) of X_n); similarly, visiting F_{n_j}' is almost the same event as visiting Y_{n_j}, and it is unlikely to visit any Y_{n_j} after having visited X_{n_i} by Lemma 7.1. We can thus use the above technique of choosing a sequence $(n_i)_{i \in \mathbb{N}}$ corresponding to the sequence ϵ_i where the ϵ_i are used as upper bounds of probabilities of the unlikely events we want to avoid, to achieve (35).

Again, we may assume without loss of generality that $n_i = i$ as we did above.

Using (33), (35) and the Borel-Cantelli lemma, we deduce that our random walk almost surely exits finitely many B_n at an $(n-1, n)$-impurity and switches between the F_i and the F_i' finitely often. This means that almost surely there is some m such that its trajectory converges to a point in $\overline{F_m}$ if it has hit B_m in F_m and to a point in $\overline{F_m'}$ if it has hit B_m in F_m', having exited each $B_n, n > m$ at F_n in the former case and at F_n' in the latter. Note that conditioning on the first case we have almost sure occurrence of the event 1^s, while conditioning on the latter we have 0^s almost surely.

Moreover, the probability with which the limit of the random walk trajectory lies in $\overline{F_m}$ but not in X (and similarly with F_m' and Y) is at most $w(\overline{F_m} \triangle X)$ because w coincides with the exit measure by Lemma 6.2. Since, clearly, we can choose m arbitrarily large, this probability can be made arbitrarily small by (34). Putting these remarks together, we obtain $\mu_o(1_o^s \triangle f^{-1}(X)) = 0$ as desired. $\qquad\square$

The following is essentially a reformulation of Corollary 1.4.

Corollary 7.7. *Let G be a locally finite transient graph and let μ_o be the distribution of random walk from $o \in V$. Let $(\mathcal{N}, \mathcal{O})$ be a topological space endowed with Borel measures $(\nu_z)_{z \in V}$ and a 'projection' $\tau : V \to \mathcal{O}$ so that the following are satisfied*

(i) there is a Borel-measurable function $\tau^* : \mathcal{W} \to \mathcal{N}$ mapping almost every $(Z^n) \in \mathcal{W}$ to $\lim_n \tau(Z^n)$ —in particualr, for random walk Z^n on G, $\tau(Z^n)$ converges in $\mathcal{N}$ μ-almost surely— and $\mu_z(\{\tau^*((Z^n)) \in O\}) = \nu_z(O)$.

(ii) there is a sequence $(G_n)_{n \in \mathbb{N}}$ of subgraphs of G with $\bigcup G_n = G$ such that random walk visits all boundaries B_n of G_n almost surely (this is automatically satisfied when G_n is finite), and $\mu_o^n(b) = \nu_o \circ \tau(b)$ for every $b \in B_n$, where μ_o^n denotes the exit distribution of G_n for random walk from o.

(iii) for every sharp harmonic function s and every z, we have $\lim_{m,n} \nu_z(\tau(F_m) \triangle \tau(F_n)) = 0$, where $F_i := \{b \in B_n \mid s(b) > 1/2\}$.

Then $\mathcal{N}$ is a realisation of the Poisson boundary of G.

Proof. Note that $\mathcal{N}$ endowed with the measures $(\nu_z)_{z \in V}$ satisfies the axioms for being a G-boundary. The proof of Lemma 7.6 applies almost verbatim, replacing $\mathcal{C}$ with $\mathcal{N}$ and w with ν, the role of τ being played by the 'horizontal' position in the tiling. The set M_n^m of (m, n)-impurities can be defined in general as $\tau(F_m) \triangle \tau(F_n)$. This yields that $\mathcal{N}$ is faithful to every sharp harmonic function, and so by Theorem 1.3 it is a realisation of the Poisson boundary. $\qquad\square$

In condition (iii) we can replace F_m, F_n by sets X_m, X_n defined as above (20) because of (31), and it will probably be helpful to do so in applications.

Recall that the only cases in which we used the bounded degree assumption on G were in proving that the widths $w(x)$ converge to zero (17) and that the image of random walk almost surely converges to $\mathcal{C}$ (Theorem 6.8). Thus we can strengthen Theorem 1.1 as follows

Corollary 7.8. *If G is a plane, uniquely absorbing graph such that $\limsup_n \max_{b \in B_n} \mu_o^n(b) = 0$, and the image of random walk in the square tiling of G almost surely converges to the boundary $\mathcal{C}$, then $\mathcal{C}$ is a realisation of the Poisson boundary of G.*

Here, B_n is the boundary of the nth member of a sequence G_n as in our proof, and we applied Lemma 6.2 to rewrite (17) in terms of the exit distributions μ_o^n.

8 The Northshield circle

Northshield [34] endowed every plane graph G satisfying properties (i)–(iii) below with compactifications $|G|_\sim$ and $|G|_\cong$, and essentially conjectured that the

boundary $\partial_\sim(G)$ of the former, which he showed to be homeomorphic to a circle, is a realisation of the Poisson boundary, while the boundary $\partial_\cong(G)$ of the latter is homeomorphic to the Martin boundary of G. We will confirm the former conjecture by showing that it is canonically homeomorphic to the boundary of the square tiling of G.

Every (plane) graph G considered by Northshield has the following properties

(i) G is *accumulation-free*, i.e. no point of $\mathbb{R}^2$ is an accumulation point of vertices of G^2;

(ii) G is *non-amenable*, i.e. $\inf_K \frac{|\partial K|}{|K|} > 0$, where K ranges over all finite vertex-sets of G, and its boundary ∂K is the set of vertices outside K sending an edge to K;

(iii) G has bounded degrees.

We now repeat the definition of $|G|_\sim$ and $|G|_\cong$ from [34]. The corresponding boundaries are defined as equivalence classes of geodesics in G with a common starting vertex o. Let $\Gamma = \Gamma_o$ denote the set of these geodesics. For any two elements $R, L \in \Gamma$ with finite intersection, $\{R \cup L\}$ divides the plane, and hence G, into two unbounded open domains by (i), and we will denote by (R, L) and (L, R) the intersection of each of these domains with G, where we use a fixed orientation of the plane to decide which is which. Note that (R, L) and (L, R) can be empty; this occurs exactly when $R \cup L$ bounds an infinite face of G. We use $[R, L]$ to denote the subgraph of G spanned by $R \cup (R, L) \cup L$.

Given $R, L \in \Gamma$, we write $R \sim L$ whenever (R, L) or (L, R) is a recurrent graph. We write $R \cong L$ whenever there exists a *transient* 1-way infinite path Q such that both $R \cap Q$ and $L \cap Q$ contain infinitely many vertices; here, we call Q *transient*, if random walk on G almost surely visits Q finitely often. We let $\partial_\sim(G)$ denote the set of equivalence classes of $\sim$ and let $\partial_\cong(G)$ denote the set of equivalence classes of $\cong$. Then we let $|G|_\sim = G \cup \partial_\sim(G)$ and $|G|_\cong = G \cup \partial_\cong(G)$. It remains to specify the topologies of these spaces, and this is done in a natural way using the embedding of G in the plane as follows.

Let $\simeq$ denote any of the relations $\sim$ or $\cong$. Given an equivalence class $[Q] \in \Gamma/\simeq$, we will, with a slight abuse of notation, write $[Q] \in (R, L)$ to denote that either (a) $|R \cap L| = \infty$, or (b) $R, L \notin [Q]$ and for every $W \in [Q]$, a cofinal subgeodesic of W is contained in (R, L). Using this notation, we specify a base for the topology of $|G|_\simeq$ by declaring all sets of the form $(R, L) \backslash S$ to be open, where $R, L \in \Gamma$ and $S \subset G$ is finite.

By the following theorem, $\partial_\sim(G)$ is homeomorphic to $\S^1$; we will call it the *Northshield circle* of G.

Conjecture 8.1 ([34, CONJECTURE 2.6]). *Let G be an accumulation-free plane, non-amenable graph with bounded vertex degrees. Then every bounded harmonic function is the integral of a bounded measurable function with respect to harmonic measure on the Northshield circle.*

Theorem 8.1 ([34, THEOREMS 2.2, 2.3 and 2.7]). *Let G be a accumulation-free plane non-amenable graph with bounded vertex degrees. Then*

(i) $\partial_\sim(G)$ is homeomorphic to a circle;

---

[2]Northshield requires instead that every cycle in G surrounds only finitely many vertices; that this is equivalent to (i) as proved in [38, Lemma 7.1].

(ii) random walk on G converges to a point in $\partial_\sim(G)$ almost surely, and

(iii) for every geodesic R and $v \in V(G)$, we have $\mu_v(\{[R]\}) = 0$.

Lemma 8.2 ([34, LEMMA 2.1])**.** *Let G be as in Theorem 8.1 and $R \not\sim L$ be geodesics in Γ. Then (R, L) contains a geodesic Q with $Q \not\sim R, L$.*

Lemma 8.3 ([34, LEMMA 4.2])**.** *Let G be as in Theorem 8.1 and $R, L \in \Gamma$. If $R \cong L$ then $R \sim L$.*

8.1 Proof of Northshield's conjecture

In this section we prove the following result, which provides a positive answer to Conjecture 8.1.

Theorem 8.4. *Let G be an accumulation-free plane, non-amenable graph with bounded vertex degrees. Then the Northshield circle $\partial_\sim(G)$ is canonically homeomorphic to the boundary $\mathcal{C}$ of the square tiling of G. Thus the former is a realisation of the Poisson boundary.*

Proof. To begin with, note that if G is accumulation-free then it is uniquely absorbing by the definitions. Moreover, it is transient by Theorem 8.1, and so the square tiling and its boundary are well-defined.

We will define a homeomorphism $h : \partial_\sim G \to \mathcal{C}$ directly, by showing that the image of any geodesic R of G in the tiling converges to a point of $\mathcal{C}$ depending only on the limit of R in $\partial_\sim G$ and not on R itself.

Suppose, to the contrary, that for some geodesic R with edges $e_1, e_2, \ldots$, there are distinct meridians M, N of T each met by $\tau(e_i)$ for infinitely many i. Then, since the image of random walk (z_i) converges to a point between M and N with positive probability by Corollary 6.9, planarity implis that z_i must meet R infinitely often with positive probability. This however contradicts Theorem 8.1 (iii).

By a similar argument, if R, L are geodesics such that $R \sim L$, then their images under τ on T converge to the same point of $\mathcal{C}$, for random walk converges to a point in $\partial_\sim(G)$ by Theorem 8.1 (ii).

This allows us to define a map $h : \partial_\sim G \to \mathcal{C}$ by mapping any $x \in \partial_\sim G$ to the unique point p of $\mathcal{C}$ such that for every geodesic $x_1, x_2, \ldots$ in x, we have $\tau(x_i) \to p$.

To show that h is a homeomorphism, let us first check that it is injective. For this, suppose there are geodesics R, L such that $h([R]) = h([L]) = p \in \mathcal{C}$, and so $\tau(R_i) \to p$ and $\tau(L_i) \to p$ by the definition of h. Then Corollary 6.9 implies that

$$\text{random walk eventually avoids } (R, L) \text{ almost surely.} \tag{36}$$

We would like to show that $R \sim L$. If this is not the case, then using Lemma 8.2 we can find a geodesics Q in (R, L) and a geodesic P in (Q, L) such that all R, Q, P, L are pairwise non-equivalent.

We distinguish two cases. Suppose first that there are infinitely many, pairwise disjoint paths $B_1, B_2, \ldots$ joining R to Q in $[R, L]$. Then combining these paths with parts of R and Q, we can construct an infinite path X interesecting each of R, Q infinitely often. By (36), this path proves that $R \cong Q$. This however contradicts Lemma 8.3.

Thus this case cannot occur, and so no such family of paths can exist. This means that there is a finite set S_1 of vertices separating R from Q in $[R, L]$ [14, p. 202]. Similarly, there is a finite set S_2 of vertices separating P from L in $[R, L]$. This implies that $S_1 \cup S_2$ is a finite set separating R and L from $[Q, P]$ in $[R, L]$; hence $S_1 \cup S_2$ separates $[Q, P]$ in G. By our choice of Q, P, the subgraph $[Q, P]$ is transient. This easily implies that random walk in G stays eventually in $[Q, P]$ with positive probability, for if it visits a vertex in $[Q, P] \backslash (S_1 \cup S_2)$, it might never visit $(S_1 \cup S_2)$ again. But this contradicts (36) because $[Q, P] \subseteq [R, L]$.

This proves that h is injective. By its definition, h preserves the natural cyclic ordering induced on the circles $\partial_\sim(G)$ and $\mathcal{C}$ by the embedding of G, and so h is continuous. It easily follows by topological arguments that h is onto, hence a homeomorphism. □

Lemma 8.5. *Let G be a accumulation-free plane non-amenable graph with bounded vertex degrees and no infinite faces. Then for every $R, L \in \Gamma$, we have $R \sim L$ if and only if $R \cong L$, and so $\partial_\sim(G)$ coincides with $\partial_\cong(G)$.*

Proof. By Lemma 8.3 $R \cong L$ implies $R \sim L$, so it remains to show that if $R \sim L$ then $R \cong L$. For this, note that since G has no infinite faces, $[R, L]$ contains an infinite set of pairwise disjoint R-L paths. Combining such paths with parts of R and L we hence obtain a path X in $[R, L]$ visiting each of R, L infinitely often. We will show that X is transient, which means that $R \cong L$.

We distinguish three cases. Suppose first that with positive probability, random walk on G stays eventually inside $[R, L]$. Then the subgraph $[R, L]$ is transient, contradicting our assumption that $R \sim L$.

The second case is where random walk visits both $[R, L]$ and its complement infinitely often with positive probability. This implies, using planarity, that at least one of R, L is visited infinitely often with positive probability. But this contradicts Theorem 8.1 (iii).

The only remaining case is where random walk visits $[R, L]$ only finitely often almost surely. But then our path X is transient indeed, proving that $R \cong L$ as desired. □

8.2 Hyperbolic graphs

In this section we prove Corollary 1.2 and disprove another conjecture of [34] saying that graphs as in Theorem 8.1 are necessarily hyperbolic. We assume that the reader is familiar with the basic notions about hyperbolic graphs and their boundaries; otherwise [22, 37] can be consulted.

One of the homeomorphisms in Corollary 1.2 was already known: that the Martin boundary coincides with the hyperbolic boundary in a consequence of

Theorem 8.6 ([1, 2]). *Let G be an infinite, Gromov-hyperbolic, non-amenable, graph (bounded degree?). Then the hyperbolic boundary of G is canonically homeomorphic to the Martin boundary.*

Moreover, it is proved in [34, LEMMA 5.1] that if G is hyperbolic in addition to satisfying the above conditions, then the hyperbolic boundary $\partial_h(G)$ coincides with $\partial_\cong(G)$ as a set; more precisely, two geodesics are equivalent with respect to to $\cong$ if and only if they converge to the same point of $\partial_h(G)$. It follows that the corresponding topologies also coincide in that case:

Lemma 8.7. *Let G be a hyperbolic, accumulation-free plane, non-amenable graph with bounded vertex degrees. Then $\partial_h(G)$ is canonically homeomorphic to $\partial_\cong(G)$.*

Proof. As already mentioned, there is a canonical bijection between $\partial_h(G)$ and $\partial_\cong(G)$, so it just remains to show that their topologies coincide. The topology of $\partial_\cong(G)$ is, by definition, generated by the 'intervals' of the form $\overline{(R,L)} \cap \partial_\cong(G), R, L \in \Gamma$. We will show that the same sets generate the topology of $\partial_h(G)$.

For this, suppose that the geodesics $L, M, R \in \Gamma$ are non-equivalent and that $M \in [L, R]$. Let π_{LM}, π_{MR} and π_{RL} be two-way infinite geodesics joining the corresponding limit points in $\partial_h(G)$. We may assume without loss of generality that both π_{LM} and π_{MR} lie in the region of the plane bounded by π_{RL} that does not contain o, for otherwise we can replace any subpath crossing to the other region by a subpath of π_{RL} of the same length. For $x, y \in \partial_h(G)$, define the *Gromov product* $\hat{xy}$ as the shortest length of a path from o to any x-y geodesic. It is known [37] that for every hyperbolic graph G the topology of $\partial_h(G)$ can be generated by the sets of the form $O_{x,r} := \{y \in \partial_h(G) \mid \hat{xy} > r\}$, where x ranges in all of $\partial_h(G)$ and r in $\mathbb{R}_{>0}$. Now since π_{LM} and π_{MR} lie in the region of the plane bounded by π_{RL} that does not contain o, it follows that $d_h(\partial L, \partial M), d_h(\partial M, \partial R) \leq d_h(\partial L, \partial R)$. This easily implies our claim that the intervals $\overline{(R,L)} \cap \partial_h(G), R, L \in \Gamma$ generate the topology of $\partial_h(G)$. $\square$

Combining Lemmas 8.7 and 8.5 with Theorems 8.6 and 8.4 yields Corollary 1.2. For this, we use the fact that an 1-ended plane graph is always accumulation-free [35].

We now adress the question of whether all conditions in the assertion of Corollary 1.2 and Theorem 8.4 are needed.

To see that having no infinite faces is essential in Corollary 1.2, let G be a graph satisfying all conditions, pick a 2-way infinite geodesic R, and let G' be obtained from G by deleting one of the two sectors into which R divides G. Then the hyperbolic and Martin boundary becomes an arc, while the Northshield circle and the square tiling boundary remain circles.

Northshield conjectured that hyperbolicity is implied by the other conditions in Corollary 1.2:

Conjecture 8.2 ([34, CONJECTURE 5.3]). *Let G be a plane, accumulation-free, non-amenable graph with bounded vertex degrees. Then G is hyperbolic.*

Figure 4 shows a 1-ended counterexample without infinite faces. It is constructed recursively as follows. Start with a triangle L_0 embedded in $\mathbb{R}^2$. Draw a new edge incident to each vertex of L_0 inside the unbounded region of $\mathbb{R}^2$, and join the new vertices with a cycle L_1. Then perform infinitely many steps of the following type: at step $i = 1, 2, \ldots$, subdivide each edge of L_i by placing i new vertices on it, draw a new edge incident to each vertex (new or old) of L_i inside the unbounded region of $\mathbb{R}^2$, and join their other ends with a cycle L_{i+1}.

We claim that $\inf_K \frac{|\partial K|}{|K|} \geq 1$ for this graph, which means that it is non-amenable. This can be proved by induction on the number of vertices in a smallest hypothetical finite set K such that $\frac{|\partial K|}{|K|} < 1$: consider the highest level L_n of the graph containing a vertex of K. If $L_n \subset K$, then letting $K' = K \backslash L_n$ we easily have $\frac{|\partial K'|}{|K'|} < \frac{|\partial K|}{|K|}$, contradicting the minimality of K. Otherwise,
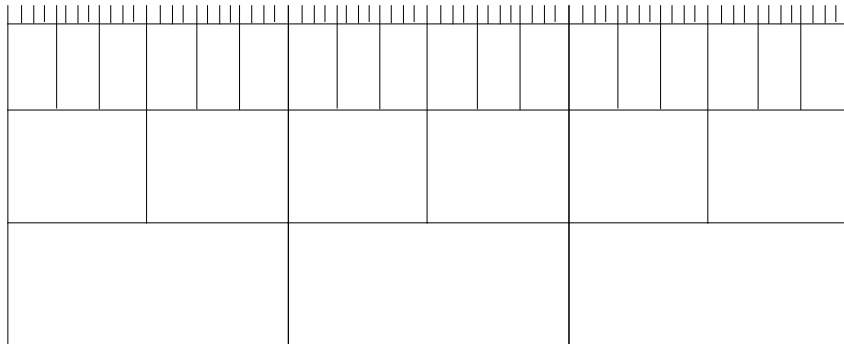
Figure 4: A counterexample to Conjecture 8.2. The righmost and leftmost geodesics are to be thought as being identified. The cycles L_i are drawn horizontally.

there is a vertex $x \in L_n \cap K$ having a neighbour $y \in L_n \backslash K$. Thus x contributes at least 2 to $|\partial K|$, and again letting $K' = K - x$ we have $\frac{|\partial K'|}{|K'|} < \frac{|\partial K|}{|K|}$ because adding the same amount to both the numerator and denumerator to a positive fraction that is smaller than 1 increases the value of the fraction.

The fact that this graph is not hyperbolic can be seen by noticing that every face boundary forms a geodesic triangle —considering the endpoints and the midpoint of its upper path as the corners of the triangle— and that these triangles are not thin because there are arbitrarily long face boundaries.

To see that non-amenability is not implied by the other conditions, let G be a graph satisfying them all, choose an infinite sequence $(v_n)_{n \in \mathbb{N}}$ of vertices of G, and attach a path of length n to each v_n; then the new graph is amenable, but retains all other properties.

It is not clear whether non-amenability or the bounded degree condition is necessary to make Corollary 1.2 true, but the following example, due to Itai Benjamini (private communication), shows that we cannot drop both conditions simultaneously: start with the infinite binary tree T embedded in $\mathbb{R}^2$. For every $n \in \mathbb{N}$, connect the vertices at distance n from a root o with a cycle C_n visiting them all without violating planarity. Then for every n and every edge e of C_n, insert 2^{2^n} paths of length 2 joining the endpoints of e. It is straightforward to check that the resulting graph is hyperbolic. The paths of length 2 inserted in the last step ensure that this graph has the Liouville property, that is, all bounded harmonic functions are constant, e.g. using the methods of [9] or [6]. It follows from the symmetry of this graph and the Liouville property that its Martin boundary is trivial. The other four boundaries are still homeomorphic to a circle. This example also shows that in Corollary 7.8 we cannot drop the condition that the image of random walk almost surely converges to $\mathcal{C}$.

This discussion motivates the following problems

Problem 8.3. *Is every planar graph with the Liouville property amenable?*

Note that for Cayley graphs this is true even without the planarity condition

[26], while for general graphs it is false even assuming bounded degrees [9].

Problem 8.4. *Is there a planar, Gromov-hyperbolic graph with bounded degrees and the Liouville property?*

Acknowledgement

I am grateful to Omer Angel, Ori Gurel-Gurevich and Asaf Nachmias for pointing out some shortcommings in an earlier draft of this paper.

References

[1] A. Ancona. Negatively Curved Manifolds, Elliptic Operators, and the Martin Boundary. *The Annals of Mathematics*, 125(3):495, May 1987.

[2] Alano Ancona. Negatively Curved Manifolds, Elliptic Operators, and the Martin Boundary. *The Annals of Mathematics*, 125(3):495, May 1987.

[3] Michael T. Anderson and Richard Schoen. Positive harmonic functions on complete manifolds of negative curvature. *The Annals of Mathematics*, 121(2):429, March 1985.

[4] Omer Angel, Martin T. Barlow, Ori Gurel-Gurevich, and Asaf Nachmias. Boundaries of planar graphs, via circle packings. *arXiv:1311.3363 [math]*, November 2013.

[5] J. R. Baxter and R. V. Chacon. The equivalence of diffusions on networks to Brownian motion. *Contemp. Math.*, 26:33–47, 1984.

[6] I. Benjamini, N. Curien, and A. Georgakopoulos. The Liouville and the intersection properties are equivalent for planar graphs. 17(42):1–5, 2012.

[7] I. Benjamini and O. Schramm. Harmonic functions on planar and almost planar graphs and manifolds, via circle packings. *Invent. math.*, 126:565–587, 1996.

[8] I. Benjamini and O. Schramm. Random Walks and Harmonic Functions on Infinite Planar Graphs Using Square Tilings. *Ann. Probab.*, 24(3):1219–1238, 1996.

[9] Itai Benjamini and Gady Kozma. Nonamenable liouville graphs. *arXiv:1010.3365 [math]*, October 2010.

[10] R. L. Brooks, C. A. B. Smith, A. H. Stone, and W. T. Tutte. The dissection of rectangles into squares. *Duke Mathematical Journal*, 7(1):312–340, 1940.

[11] J. W. Cannon, W. J. Floyd, and W. R. Parry. Squaring rectangles: The finite Riemann mapping theorem. Abikoff, William (ed.), The mathematical legacy of Wilhelm Magnus. Contemp. Math. 169, 1994.

[12] Y. Colin de Verdière, I. Gitler, and D. Vertigan. Planar electric networks. II. (Reseaux électriques planaires. II.). *Comment. Math. Helv.*, 71(1):144–167, 1996.

[13] M. Dehn. Über Zerlegung von Rechtecken in Rechtecke. *Math. Ann.*, 57:314–332, 1903.

[14] Reinhard Diestel. *Graph Theory* (3rd edition). Springer-Verlag, 2005. Electronic edition available at: `http://www.math.uni-hamburg.de/home/diestel/books/graph.theory.`

[15] A.J.W. Duijvestijn. Simple perfect squared square of lowest order. *J. Comb. Theory, Ser. B*, 25:240–243, 1978.

[16] A. Erschler. "poisson-furstenberg boundaries, large-scale geometry and growth of groups". In *Proceedings of the ICM, Hyderabad, India*, pages 681–704, 2010.

[17] A. Erschler. Poisson–Furstenberg boundary of random walks on wreath products and free metabelian groups. *Comment. Math. Helv.*, 86(1):113–143, 2011.

[18] H. Furstenberg. A Poisson formula for semi-simple Lie groups. *The Annals of Mathematics*, 77(2):335–386, March 1963.

[19] H. Furstenberg. Random walks and discrete subgroups of Lie groups. *Advances Probab. related Topics*, 1:3–63, 1971.

[20] A. Georgakopoulos and V. Kaimanovich. In preparation.

[21] A. Georgakopoulos and V. Kaimanovich. Electrical network reduction with a probabilistic interpretation of effective conductance. Preprint 2011.

[22] E. Ghys and P. de la Harpe. *Sur les Groupes Hyperboliques d' après Mikhael Gromov*. Birkhäuser, 1990. English translation available at http://www.umpa.ens-lyon.fr/ ghys/Publis-english.html.

[23] O. Gurel-Gurevich and A. Nachmias. Recurrence of planar graph limits. The Annals of Mathematics (to appear).

[24] V. Kaimanovich. The poisson formula for groups with hyperbolic properties. *The Annals of Mathematics*, 152(3):pp. 659–692, 2000.

[25] V. A. Kaimanovich and H. Masur. The poisson boundary of the mapping class group. *Inventiones Mathematicae*, 125(2):221–264, 1996.

[26] V. A. Kaimanovich and A. M. Vershik. Random walks on discrete groups: Boundary and entropy. *The Annals of Probability*, 11(3):457–490, August 1983.

[27] Ioannis Karatzas and Steven E. Shreve. *Brownian Motion and Stochastic Calculus.* Springer, 1991.

[28] A. Karlsson and W. Woess. The poisson boundary of lamplighter random walks on trees. *Geometriae Dedicata*, 124(1):95–107, November 2006.

[29] A. F. Karr. *Probability.* Springer Texts in Statistics, 1993.

[30] R. Kenyon. Tilings and discrete dirichlet problems. *Israel Journal of Mathematics*, 105(1):61–84, 1998.

[31] R. W. Kenyon and S. Sheffield. Dimers, tilings and trees. *Journal of Combinatorial Theory, Series B*, 92(2):295–317, November 2004.

[32] R. Lyons and Y. Peres. *Probability on Trees and Networks.* Cambridge University Press. In preparation, current version available at `http://mypage.iu.edu/~rdlyons/prbtree/prbtree.html.`

[33] J. W. Pitman M. T. Barlow and M. Yor. On Walsh's Brownian motions. Séminaire de Probabilités (Strasbourg), 23(1):275–293, 1989.

[34] S. Northshield. Circle boundaries of planar graphs. *Potential Analysis*, 2(4):299–314, December 1993.

[35] R.B. Richter and C. Thomassen. 3-connected planar spaces uniquely embed in the sphere. *Trans. Am. Math. Soc.*, 354:4585–4595, 2002.

[36] S. Rohde. Oded schramm: From circle packing to SLE. In I. Benjamini and O. Häggström, editors, *Selected Works of Oded Schramm*, Selected Works in Probability and Statistics, pages 3–45. Springer New York, 2011.

[37] H. Short (editor). Notes on Hyperbolic Groups. `http://www.cmi.univ-mrs.fr/~hamish/`.

[38] C. Thomassen. Planarity and duality of finite and infinite graphs. *J. Combin. Theory (Series B)*, 29(2):244 – 271, 1980.

[39] Wolfgang Woess. *Denumerable Markov chains. Generating functions, boundary theory, random walks on trees.* EMS Textbooks in Mathematics. European Mathematical Society (EMS), Zürich, 2009.