Discrete probability and the geometry of graphs

by

Thomas Hutchcroft

B.A., University of Cambridge, 2012
M.Math., University of Cambridge, 2013
M.A., University of Cambridge, 2016

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate and Postdoctoral Studies

(Mathematics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

August 2017

Abstract

We prove several theorems concerning random walks, harmonic functions, percolation, uniform spanning forests, and circle packing, often in combination with each other. We study these models primarily on planar graphs, on transitive graphs, and on unimodular random rooted graphs, although some of our results hold for more general classes of graphs. Broadly speaking, we are interested in the interplay between the geometry of a graph and the behaviour of probabilistic processes on that graph. Material taken from a total of nine papers is included. We have also included an extended introduction explaining the background and context to these papers.

Lay Summary

There is a deep and well-developed connection between the geometry of a space and the way probabilistic processes behave on that space. A classical example is given by a heavy particle (e.g. a mote of dust) that moves randomly through space due to its bumping into a large number of light particles (e.g., air molecules): The particle will move much faster in negatively curved spaces than in flat space. In this thesis we continue to develop this connection in a discrete setting, where our spaces are modeled by graphs. We prove several new results connecting the geometry of a graph to the behaviour of random processes on it. We are particularly interested in *random walk*, a discrete version of the random motion mentioned above, *percolation*, which models a random porous material, and the *uniform spanning forest*, a probability model that is closely connected to the theory of electrical networks.

Preface

- Chapter 2 is adapted from the paper *Critical percolation on any quasi-transitive graph of exponential growth has no infinite clusters* [129], of which I am the only author. This paper was published in *Comptes Rendus Mathematique* in 2016.

- Chapter 3 is adapted from the paper *Collisions of random walks in reversible random graphs* [133], by Yuval Peres and myself. Research was conducted as an equal collaboration, while I did most of the writing. This paper was published in *Electronic Communications in Probability* in 2015.

- Chapter 4 is adapted from the paper *Wired cycle-breaking dynamics for uniform spanning forests* [127], of which I am the only author. This paper was accepted for publication in *The Annals of Probability* in 2015.

- Chapter 5 is adapted from the preprint *Interlacements and the wired uniform spanning forest* [128], of which I am the only author. This paper was accepted for publication in *The Annals of Probability* in 2017.

- Chapter 6 is adapted from the paper *Indistinguishability of trees in uniform spanning forests* [131], by Asaf Nachmias and myself. Research was conducted as an equal collaboration, while I did most of the writing. The paper was accepted for publication in *Probability Theory and Related Fields* in 2016.

- Chapter 7 is adapted from the paper *Unimodular hyperbolic triangulations: circle packing and random walk* [21], by Omer Angel, Asaf Nachmias, Gourab Ray, and myself. Research was conducted as an equal collaboration, while I did most of the writing. The paper appeared in *Inventiones Mathematicae* in 2016.

- Chapter 8 is adapted from the preprint *Boundaries of planar graphs: a unified approach* [132], by Yuval Peres and myself. Research was conducted as an equal collaboration, while I did most of the writing. The paper was accepted for publication in *The Electronic Journal of Probability* in 2017.

- Chapter 9 is adapted from the preprint *Hyperbolic and parabolic unimodular random maps* [20], by Omer Angel, Asaf Nachmias, Gourab Ray, and myself. Research was conducted as an equal collaboration, while I did most of the writing. The preprint has been submitted for publication.

- Chapter 10 is adapted from the preprint *Uniform spanning forests of planar graphs* [130], by Asaf Nachmias and myself. Research was conducted as an equal collaboration, while I did most of the writing. The preprint has been submitted for publication.

- Chapter 1, the introduction, contains small amounts of the material from all the other chapters.

Table of Contents

List of Figures

Acknowledgements

First and foremost, I thank my advisors Omer Angel and Asaf Nachmias for guiding me through the early stages of my research, introducing me to the field, sharing problems and ideas, reading drafts, and providing me with many opportunities to meet and collaborate with others. I would also like to thank my Microsoft Research mentors Ander Holroyd and Yuval Peres for several very enjoyable and productive summers in Redmond, which were hugely beneficial to me. I also thank Microsoft Research for supporting me financially in the final year of my doctoral studies by awarding me the Microsoft Research PhD Fellowship. I also thank my collaborators Omer Angel, Nicolas Curien, Ander Holroyd, Avi Levy, Asaf Nachmias, Yuval Peres, and Gourab Ray for all the stimulating conversations and discoveries we have had together. I hope for many more to come. I also thank Russ Lyons and Yuval Peres for writing their wonderful book [173], which I refer to at least once a week, and Itai Benjamini for sending me many interesting questions. I also thank everyone who has been a member of the UBC probability group during my time here, especially my fellow students and visiting students: Owen Daniel, Mallory Flynn, Spencer Frei, Tyler Helmuth, Saraí Hernández-Torres, Jieliang Hong, Thomas Hughes, Tim Jaschek, Matthias Klöckner, Brett Kolesnik, Hongliang Lu, Guillermo Martínez Dibene, Gourab Ray, Saif Syed, Alex Tomberg, Ben Wallace, Zichun Ye, Jing Yu, and Qingsan Zhu; I could not have asked for a better group of colleagues, peers, and friends.

Finally, although I never had the privilege to meet him, I would like to thank Oded Schramm, whose mathematical work, ingenuity, and vision has been a constant inspiration to me. His influence is felt throughout essentially every subject touched on in this thesis, and it is no coincidence that his name appears more than 250 times in this thesis, an average of once every 1.4 pages. Throughout my work, and especially in my joint work with Asaf, whenever we were unsure how best to approach something, it became routine to ask ourselves *"What would Oded do?"* Reflecting on this question always helped us cut right to the heart of the matter at hand, even if we could not always live up to the aspiration of doing things quite as elegantly as he would have.

To Meghan, my family and friends.

Chapter 1

Introduction

1.1 What is this thesis about?

This thesis is a collection of papers about discrete probability on infinite graphs. The central theme can be summarised by the following question.

What does the geometry of a graph tell us about the behaviour of probabilistic processes on that graph? Conversely, what does the behaviour of probabilistic processes on a graph tell us about the underlying geometry?

This question has formed an important strand of thought through modern probability theory since its inception in the early 20th century. In this introduction, we will take a quick tour through the highlights of this tradition. Some of the results of this thesis will be presented as we go, although most of our own contributions will be presented towards the end. Our aim is to present some of the main ideas underlying the field, and our own work, in a way that is hopefully accessible to a general mathematical audience, without getting too much into technical details.

1.2 Random walks on graphs

A **graph** $G = (V, E)$ is a set of **vertices** V and a set of **edges** E. We think of vertices as points, and of edges as curves, each of which is either a loop starting and ending at the same point, or else joins two points together. (More formally, we have an **incidence relation** which assigns to each edge either one or two vertices, which are the endpoints of the edge.) For simplicity, we shall consider in this section only **simple** graphs, that have no self-loops and at most one edge between each two vertices. A graph is **connected** if we can pass from any vertex to any other vertex by moving across a finite sequence of edges. The **degree** of a vertex is the number of edges incident to it. We will assume throughout that our graphs have countably many vertices and edges, and are **locally finite**, meaning that all the degrees are finite.

The **random walk** on a graph is the process that, at each time $n \geq 0$, chooses an edge uniformly at random from among those incident to its current location, independently from everything it has done previously, and then moves to the vertex at the other endpoint of that edge. When we have some graph in mind, we write $p_n(u, v)$ for the probability that a random walk started at the vertex u of the graph is at the vertex v at time n.

Perhaps the first important work in the tradition we describe was George Pólya's 1921 paper [196], in which he proved the following theorem. We say that a graph is **recurrent** if the random walk on the graph returns to its starting location almost surely (i.e., with probability one – we will henceforth abbreviate this to a.s.), and **transient** if it has a positive probability never to return. It is easy to see that the graph is recurrent if and only if the random walk returns to its starting point infinitely many times a.s., and transient if and only if it returns to its starting point only finitely many times a.s. The **hypercubic lattice** $\mathbb{Z}^d$ is the graph whose vertices are d-tuples of integers, and where two vertices are connected by an edge if they differ by exactly one in one coordinate and have the same value in all other coordinates.

Theorem (Pólya 1921)**.** *The d-dimensional hypercubic lattice $\mathbb{Z}^d$ is recurrent if $d \leq 2$, and transient if $d \geq 3$.*

This theorem can be summarised by the following famous aphorism, attributed to Shizuo Kakutani.

A drunk man will find his way home,
but a drunk bird may get lost forever.

Pólya's proof of his theorem was computational in nature. By computing the moment generating function of the location of the random walk at time n, Pólya showed that the random walk on $\mathbb{Z}^d$ satisfies

$$p_{2n}(0,0) \approx n^{-d/2}. \tag{1.2.1}$$

On the other hand, it is easy to see that, for parity reasons, $p_n(0,0) = 0$ for every odd n. This estimate easily yields the theorem, since it is not hard to prove that the random walk on a graph is recurrent if and only if

$$\sum_{n \geq 1} p_n(v,v) = \infty$$

for some (and hence every) vertex v of G – note that this sum is exactly the *expected* number of times the walk returns to v.

A drawback to this approach is that it relies heavily upon the symmetries of the lattice, and is not very robust. For example, the argument breaks down rather seriously if we take some arbitrary set of edges of $\mathbb{Z}^d$ and replace each of the edges in this set with two edges in parallel with the same endpoints. Besides this, it does not give us much insight into the *geometric* reasons that $\mathbb{Z}^2$ is recurrent but $\mathbb{Z}^3$ is transient.

1.2.1 The Nash-Williams Criterion and electrical networks

The first major step towards a geometric understanding of the recurrence/transience problem was the *Nash-Williams Criterion* [186], which gives a sufficient condition for recurrence in terms of the isoperimetry of the graph. For each set of vertices $W \subseteq V$ in a graph G, we write $|W| = \sum_{v \in W} \deg(v)$. We say that a set W is a **cutset** for v if every path from v to infinity must pass through W.

Theorem (Nash-Williams 1959). *Let G be an infinite graph, let v be a vertex of G, and suppose that there exists a disjoint sequence of finite cutsets W_i for v such that*

$$\sum_{i\geq 1} |W_i|^{-1} = \infty.$$

Then G is recurrent.

This theorem easily recovers Pólya's result in dimensions 1 and 2 by, for example, using the cutsets $W_i = \{-i, i\}$ in $\mathbb{Z}$ and the cutsets $W_i = \{(x,y) : |x| + |y| = i\}$ in $\mathbb{Z}^2$.

Nash-Williams's proof was based on the correspondence between random walks and **electrical networks**, a tool which will be of great importance throughout this thesis. A detailed introduction, covering everything we discuss in this subsection, can be found in [173, Chapters 2 and 3]. We can think of a graph as an electrical network in which each edge contains a unit resistor[1]. Suppose that we have such a network arising from a *finite* graph, and we attach the two ends of a battery to two disjoint sets of vertices, A and B. Doing this will cause a **current** to flow through the network. Mathematically, the **current flow** is a function θ from the set of **oriented edges** of the graph (that is, edges e together with a choice of one of the endpoints as the tail of the edge, denoted e^-, and the other as the head of the edge, denoted e^+) that is antisymmetric in the sense that $\theta(e) = -\theta(-e)$ for every oriented edge e of G, where $-e$ denotes the reversal of e. Given an antisymmetric edge function θ and a vertex v, we write $\nabla^*\theta(v)$ for the sum $\sum_{e^-=v} \theta(e)$.

Kirchoff's laws of electrical networks [154] are as follows:

1. (The node law) For every vertex v of G, the sum $\nabla^*\theta(v)$ of the currents along the oriented edges emanating from v is positive if $v \in A$, negative if $v \in B$, and zero if v is not in A or B.

2. (The cycle law) If $e_1, e_2, \ldots, e_n$ form a cycle, a path from A to itself, or a path from B to itself, then $\sum_{i=1}^n \theta(e_i) = 0$.

In general, we call an antisymmetric edge function satisfying the node law a **flow** from A to B. It is not hard to prove that, *in a finite network*, the antisymmetry condition and Kirchoff's laws determine the current flow up to a positive multiplicative constant. The **strength** of a flow is defined to be the sum $\sum_{v\in A} \nabla^*\theta(v)$, and we say that a flow is a **unit flow** if it has strength one. We define the **unit current flow** from A to B to be the unique current flow from A to B that has strength 1.

Given a function $f : V \to \mathbb{R}$, we define the gradient ∇f to be the antisymmetric edge function

$$\nabla f(e) = f(e^+) - f(e^-).$$

[1] We can also consider networks with arbitrary resistances, for which there is a corresponding random walk that chooses which edge to travel along weighted by its conductance, that is, the inverse of its resistance.

The basic connection between electrical network theory and random walks is that *voltages* (i.e., antiderivatives of currents) *are hitting probabilities*. Given a graph G, we write $\mathbf{P}_v$ for the law of the random walk started at v. For each set of vertices A, we write τ_A for the first time that the random walk hits A.

Proposition. *Let $f(v) = \mathbf{P}_v(\tau_B < \tau_A)$. Then the gradient ∇f is a current flow from A to B.*

This is all well and good, but so far the electrical theory has not given us a useful way of calculating anything. However, the theory becomes very useful once we introduce *effective resistances*. The **energy** of an antisymmetric edge function θ, and the **Dirichlet energy** of a function $f : V \to \mathbb{R}$, are defined by

$$\mathcal{E}(\theta) = \frac{1}{2}\sum_{e\in E^{\rightarrow}} \theta(e)^2 \quad \text{and} \quad \mathcal{E}(f) = \frac{1}{2}\sum_{e\in E^{\rightarrow}} \Big(f(e^+) - f(e^-)\Big)^2 = \mathcal{E}(\nabla f).$$

The **effective resistance** between two disjoint sets A and B in a finite graph is defined to be the energy of the unit current flow from A to B, and is denoted $\mathscr{R}_{\text{eff}}(A \leftrightarrow B)$ (or $\mathscr{R}_{\text{eff}}(A \leftrightarrow B; G)$ if the graph we are considering is not clear). Effective resistances are related to hitting probabilities as follows. Write τ_A^+ for the first time that the random walk hits A after time zero.

Proposition. $\mathscr{R}_{\text{eff}}(A \leftrightarrow B)^{-1} = \sum_{v\in A} \deg(v)\mathbf{P}_v(\tau_B < \tau_A^+)$.

In particular, the reciprocal of the effective resistance (a.k.a. the **effective conductance**) between a single vertex a and a set B is exactly the multiple of $\deg(a)$ with the probability that the random walk started at a hits B before it returns to a.

The reason this is so useful is that we have the following two variational principles for the effective resistance.

Theorem (Thomson's Principle)**.**

$$\mathscr{R}_{\text{eff}}(A \leftrightarrow B) = \inf\{\mathcal{E}(\theta) : \theta \textit{ is a flow from } A \textit{ to } B \textit{ with strength at least one}\}.$$

Theorem (Dirichlet's Principle)**.**

$$\mathscr{R}_{\text{eff}}(A \leftrightarrow B)^{-1} = \inf\{\mathcal{E}(f) :\ f : V \to \mathbb{R} \textit{ such that } f(v) \le 0 \textit{ for all } v \in A \textit{ and } f(v) \ge 1 \textit{ for all } v \in B\}.$$

These allow us to get an upper bound on the effective resistance by constructing a low energy flow, or to get lower bounds by exhibiting low energy functions. (There are also several other closely related variational formulas for the effective resistance - the key words are **extremal length** and **the method of random paths**.)

Exercise. Prove both Thomson's and Dirichlet's principles. Hint: Differentiate the energy with respect to θ or f as appropriate and find the critical points.

Now suppose that G is an infinite graph and A is a finite set of vertices in G, and let $\langle V_n \rangle_{n\geq 0}$ be an **exhaustion** of G, that is, an increasing sequence of finite sets of vertices such that $\bigcup_{n\geq 0} V_n$. We define $\partial_V V_n$ to be the set of vertices of V_n that have a neighbour not in V_n, and let G_n be the subgraph of G induced by V_n, that is, the graph that has vertex set V_n and has all the edges of G that have both endpoints in V_n. We define

$$\mathscr{R}_{\text{eff}}(A \to \infty) = \lim_{n\to\infty} \mathscr{R}_{\text{eff}}(A \to \partial V_n; G_n),$$

so that

$$\mathscr{R}_{\text{eff}}(A \to \infty)^{-1} = \lim_{n\to\infty} \sum_{v\in A} \deg(v)\mathbf{P}_v(\tau_{\partial V_n} < \tau_A^+) = \sum_{v\in A} \deg(v)\mathbf{P}_v(\tau_A^+ = \infty).$$

Thus, we have that a graph is recurrent if and only if the effective resistance from some (and hence every) vertex to infinity is infinite. Both Thomson's and Dirichlet's principles extend to yield variational formulas for the effective conductance to infinity, with appropriate modifications: a flow from A to infinity is required to have $\nabla^*\theta(v)$ positive if $v \in A$ and zero otherwise, while in the Dirichlet principle we can take the infima over functions that are finitely supported (i.e. have $f(v) = 0$ for all but finitely many vertices v) and have $f(v) \geq 1$ for all $v \in A$.

An important consequence of Thomson's principle is that the effective resistance is a monotone function of the edge set, a fact known as **Rayleigh's monotonicity principle.** In particular, this implies that every subgraph of a recurrent graph is recurrent. This is not obvious at all by direct consideration of the random walk.

We are now ready to see how the Nash-Williams Criterion is proven. For each of the cutsets W_i, let V_i be the set of vertices u such that W_i is a cutset for u, and, for each $n \geq 1$, define a finitely supported function f_n with $f_n(v) = 1$ by

$$f_n(u) = \frac{\sum_{i=1}^n |W_i|^{-1} \mathbb{1}(u \in V_i)}{\sum_{i=1}^n |W_i|^{-1}}.$$

Since the gradients $\nabla \mathbb{1}(u \in V_i)$ have disjoint support, we can compute that

$$\mathcal{E}(f_n) = \frac{\sum_{i=1}^n |\partial_E V_i| \cdot |W_i|^{-2}}{\left(\sum_{i=1}^n |W_i|^{-1}\right)^2} \leq \left(\sum_{i=1}^n |W_i|^{-1}\right)^{-1}.$$

Thus, we have

$$\mathscr{R}_{\text{eff}}(v \to \infty) \geq \mathcal{E}(f_n)^{-1} \geq \sum_{i=1}^n |W_i|^{-1},$$

and the claim follows since n was arbitrary.

Exercise. Prove that $\mathbb{Z}^d$ is transient for $d \geq 3$ by exhibiting a flow from the origin to infinity that has finite energy.

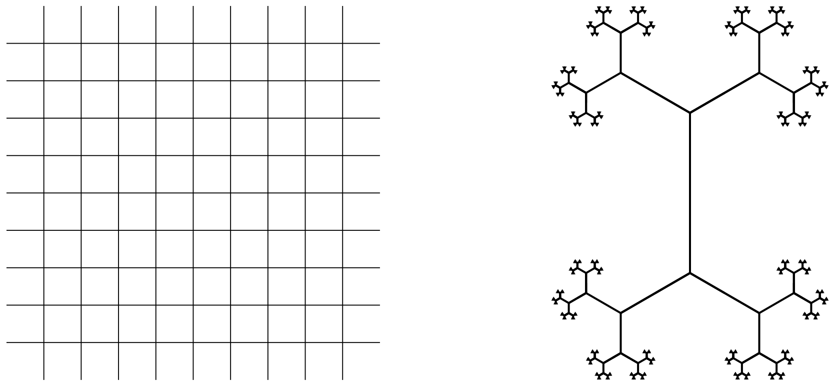

Figure 1.1: $\mathbb{Z}^2$, left, is amenable, while the 3-regular tree, right, is nonamenable.

1.2.2 The isoperimetric approach

The next milestone result in the theory was proven by Kesten, also in 1959. Nash-Williams gives us a good geometric criterion for recurrence, but what about transience? One rather extreme way for a graph to be transient, which turns out to have a simple geometric description, is for the return probabilities $p_n(v,v)$ to decay exponentially in n, meaning that there exists a constant $\alpha < 1$ and constants $C_v < \infty$ such that

$$p_n(v,v) \leq C_v \alpha^n$$

for every $v \in V$ and $n \geq 0$.

Kesten's theorem [151, 152] relates exponential decay of the return probabilities to the **isoperimetry** of the graph, that is, to the relationship between the sizes of sets in the graph to the sizes of their boundaries. We write $|\partial W|$ for the number of edges that have one endpoint in W and the other outside of W. A graph is said to be **nonamenable** if there exists a positive constant $c > 0$ such that $|\partial W|/|W| > c$ for every finite set $W \subset V$, and **amenable** otherwise. For example, $\mathbb{Z}^d$ is amenable for every $d \geq 1$ because

$$\left|[0,n]^d\right| = n^d \quad \text{ but } \quad \left|\partial[0,n]^d\right| \approx n^{d-1},$$

so that $|\partial[0,n]^d|/|[0,n]^d| \to 0$ as $n \to \infty$. On the other hand, the d-regular tree[2] is nonamenable for every $d \geq 3$: The average degree of a vertex in a *finite* tree with n vertices is easily seen to be $2 - 2/n \leq 2$, and it follows by an elementary calculation that for any finite subset W of a d-regular tree, we have

$$|\partial W| \geq (d-2)|W|.$$

Theorem (Kesten 1959)**.** *Let G be a graph, and let v be a vertex of G. Then $p_n(v,v)$ decays exponentially if and only if G is nonamenable.*

[2]A forest is a graph that does not contain any simple cycles. A tree is a connected forest. The d-regular tree is the unique tree in which every vertex has degree d.

Kesten proved his theorem only for random walks on groups. The generalisation to arbitrary graphs, as well as quantitative forms of the theorem, are due to many authors working in various contexts [11, 53, 65, 68, 80, 81, 101, 143, 224]. See [173, Chapter 6.10] for a detailed history. The fact that the exponential decay of the return probabilities implies nonamenability is rather easy, the difficult part is the other direction.

Nonamenability has many guises, and the amenability/nonamenability dichotomy is fundamental not only to the study of random walks, but also of percolation, uniform spanning forests, and many other topics that we will not touch on here, such as the ergodic theory of group actions. Indeed, the Wikipedia page on amenable groups lists nine different properties of groups that are equivalent to amenability.

We might hope for a nice isoperimetric criterion for transience, along the lines of Kesten's Theorem and the Nash-Williams Criterion, which also has an isoperimetric character. Note that attaching an infinite path to a graph will not affect whether or not the graph is transient, but will cause the graph to have very bad isoperimetry, suggesting that the best we can hope for in general is a *sufficient* condition for transience in terms of isoperimetry.

We say that a graph G satisfies an **$\phi(t)$-isoperimetric inequality** if there exists a positive constant $c > 0$ such that $|\partial W|/\phi(|W|) > c$ for every finite set $W \subset V$. In particular, G is nonamenable if and only if it satisfies a t-isoperimetric inequality. Similarly, we say that G satisfies an **anchored $\phi(t)$-isoperimetric inequality** if there exists a positive constant $c > 0$ such that $|\partial W|/\phi(|W|) > c$ for every finite connected set $W \subset V$ containing some fixed vertex v (the choice of which does not matter). The following theorem was proven by Thomassen [220]. The version we state here appears in the textbook of Lyons and Peres [173], and is adapted from a theorem of Lyons, Morris, and Schramm [170]. Similar theorems have also been obtained by He and Schramm [121], and Benjamini and Kozma [42].

Theorem (Thomassen's Criterion)**.** *Suppose that G satisfies an anchored $\phi(t)$-isoperimetric inequality for some increasing function ϕ such that*

$$\sum_{n=1}^{\infty} \phi(n)^{-2} < \infty.$$

Then G is transient.

What if we are interested in the rate of decay of $p_n(u, u)$ (known in the jargon of the field as the **on-diagonal heat kernel**), rather than just the transience/recurrence question? In particular, is there a geometric proof of Pólya's estimate (1.2.1)? For $d \geq 2$ an answer is provided by the following results. We say that a graph satisfies a d-dimensional isoperimetric inequality if it satisfies a $t^{(d-1)/d}$-isoperimetric inequality. $\mathbb{Z}^d$ satisfies a d-dimensional isoperimetric inequality, so that we can recover Pólya's estimate from the following two theorems: See the textbooks [229] and [159] for history, proofs, and related theorems.

Theorem. *Let G be an infinite, bounded degree graph and let $d \geq 2$. Then G satisfies a d-dimensional isoperimetric inequality if and only if there exists a constant $C < \infty$ such that*

$$p_n(u,u) \leq C\, n^{-d/2}$$

for every vertex u and $n \geq 1$.

Theorem. *Let G be an infinite, bounded degree graph that satisfies a d-dimensional isoperimetric inequality for some $d \geq 2$, and suppose that there exists a constant $C < \infty$ such that*

$$|B(u,r)| \leq C\, r^d$$

for every vertex u and every $r \geq 1$. Then there exists a constant $c > 0$ such that

$$p_n(u,u) \geq c n^{-d/2}$$

for every vertex u and every $n \geq 1$.

It is possible to say much more. In particular, it is possible to recover the entire (off-diagonal) behaviour of the heat kernel $p_n(x,y)$ on $\mathbb{Z}^d$ from geometric considerations. This understanding has proven very important for developing a similar understanding for random walk on *fractal-like* graphs such as the (pre-)Sierpinski gasket.

1.2.3 Bounded harmonic functions

Recall that a function h on a graph is said to be **harmonic** if for each vertex u, the value of h at u is equal to the average value of h on the neighbours of u. The existence or nonexistence of various types of non-constant harmonic functions (e.g. bounded, finite energy, positive, polynomial growth) on a graph is closely connected to the behaviour of the random walk on the graph.

Perhaps the most important class of harmonic functions are the *bounded* harmonic functions. In probabilistic terms, bounded harmonic functions encode all possible 'behaviours at infinity' of the random walk that occur with positive probability. The correspondence, which goes back to the work of Blackwell [55], can be stated (more or less) precisely as follows: First, suppose we have a (measurable) set $\mathscr{A}$ of paths in G that is **shift invariant** in the sense that if we delete an initial segment of any path in $\mathscr{A}$, then the truncated path is also in $\mathscr{A}$. We can define a function h on G by setting $h(v)$ to be the probability that if we start a random walk at v, then the resulting path will be in the set $\mathscr{A}$. It is not hard to verify that, since $\mathscr{A}$ is shift invariant, the function $h(v)$ is harmonic.

More generally, we can obtain a bounded harmonic function on G from any (measurable) bounded shift invariant *function* on the set of paths in G, by taking the expectation of the function applied to the random walk. In fact, the function needs only be defined on some set that the random walk path is a.s. in, and two shift invariant functions will yield the same harmonic function if and only if they yield the same value on a.e. random walk path. It can be shown that in fact *every* bounded harmonic function on G arises this way. In particular, G admits non-constant bounded harmonic functions if and only if there is some shift-invariant set of paths $\mathscr{A}$ and a vertex v such that the probability that the random walk started at v is in the set $\mathscr{A}$ is strictly between 0 and 1 (that is, if *the invariant σ-algebra is nontrivial*). We say that a graph is **Liouville** if it does not admit any non-constant bounded harmonic functions, and **non-Liouville** otherwise. See [173] for a detailed treatment of the theory of bounded harmonic functions and the Liouville property.

For example, it can be shown that $\mathbb{Z}^d$ is Liouville for every $d \geq 1$ (see below). On the other hand, if we attach together two copies of $\mathbb{Z}^3$ (with vertex sets $\{1\} \times \mathbb{Z}^3$ and $\{2\} \times \mathbb{Z}^3$) by a single edge connecting their origins, then the resulting graph is non-Liouville, and in fact the vector space of bounded harmonic functions on this graph is equal to the linear span of the constant 1 function and the function

$$h(x) = \mathbb{P}(\text{a random walk } \langle X_n \rangle_{n\geq 0} \text{ started at } v \text{ visits the set } \{1\} \times \mathbb{Z}^3 \text{ infinitely often}).$$

In other words, which of the two copies of $\mathbb{Z}^3$ the random walk is eventually absorbed into is the only non-trivial information there is about the random walk on this graph.

How can we tell if a graph is Liouville or not? For *transitive* graphs, there is a nice characterisation, which follows in the special case of Cayley graphs from the work of Avez [26, 27], Derriennic [77], and Kaimanovich and Vershik [149, 226], and was generalised to transitive graphs by Kaimanovich and Woess [148]. Here, an **automorphism** of a (simple) graph G is a bijection $\phi : V \to V$ such that $\phi(u)$ is adjacent to $\phi(v)$ if and only if u is adjacent to v, and a graph G is **transitive** if for every two vertices u and v of G, there is an automorphism of G sending u to v. Intuitively, a graph is transitive if it 'looks the same from every vertex'.

Theorem. *Let G be a transitive graph. Then the following are equivalent.*

1. *The random walk on G has positive speed, meaning that*

$$\liminf_{n\to\infty} \frac{1}{n} d(X_0, X_n) > 0$$

almost surely when $\langle X_n \rangle_{n\geq 0}$ is a random walk on G.

2. *The asymptotic entropy*

$$\liminf_{n\to\infty} \frac{1}{n} \sum_{v\in V} -p_n(u,v)\log p_n(u,v) = \liminf_{n\to\infty} \frac{1}{n}\mathbf{E}_u\left[-\log p_n(X_0,X_n)\right]$$

is positive for every vertex u of G.

3. *G is non-Liouville.*

(In fact the limits infimum here can be replaced with limits: This follows from Kingman's sub-additive ergodic theorem for the speed, and Fekete's Lemma for the asymptotic entropy.) This theorem allows us to immediately conclude that every transitive nonamenable graph is non-Liouville, since Kesten's theorem easily implies that both the speed and asymptotic entropy are positive for any bounded degree nonamenable graph. On the other hand, the theorem also implies that every transitive graph that has **subexponential growth**, meaning that

$$\liminf_{n\to\infty} \frac{1}{n} \log |B(u,n)| = 0$$

is Liouville, since an easy computation implies that

$$\sum_{v\in V} -p_n(u,v) \log p_n(u,v) \leq \log |B(u,n)|$$

for every vertex u and $n \geq 0$. Both of these conclusions fail for general bounded degree graphs.

Unfortunately, there is still no good *geometric* understanding of the Liouville property, particularly for the interesting case of amenable transitive graphs of exponential growth, some of which are Liouville and some of which are not. In particular, it is a long-standing open problem to prove the widely believed claim that two Cayley graphs of the same group are either both Liouville or both non-Liouville.

The fact that item (1) implies item (2) in the previous theorem is a corollary of the **Varopoulos-Carne inequality** [66, 225], which is a beautiful theorem in its own right. (Its proof, due to Carne, is an extremely elegant piece of linear algebra involving an expansion of the n-step transition matrix P^n in terms of Chebyshev polynomials of P).

Theorem (Varopoulos-Carne). *Let G be a graph. Then*

$$p_n(u,v) \leq 2\sqrt{\deg(v)/\deg(u)}\ \exp\left(-d(u,v)^2/(2n)\right).$$

Exercise. Use the Varopoulos-Carne Inequality to prove that the positivity of the speed and of the asymptotic entropy are equivalent for any bounded degree graph.

A further nice consequence of the Varopoulos-Carne bound is that for any graph G which has **polynomial growth**, meaning that there exist constants $C, d < \infty$ such that

$$|B(u,n)| \leq C\, n^d$$

for every vertex u and every $n \geq 0$, we have that

$$\limsup_{n\to\infty} \frac{d(X_0, X_n)}{\sqrt{n \log(n)}} < \infty$$

almost surely, so that graphs with polynomial growth are quite far away from having positive speed.

Exercise. Prove this claim.

1.2.4 The Poisson boundary

If we have a graph that we know to be non-Liouville, it is interesting to classify the space of *all* bounded harmonic functions in some geometric way. See [173, Chapter 14] for a detailed treatment of the material covered in this subsection.

The model situation is the classical Poisson integral formula, which states that for every bounded harmonic function u on the unit disc $\mathbb{D}$, there exists a bounded measurable function f on $\partial\mathbb{D}$ (unique up to almost-everywhere equivalence) such that

$$u(z) = \frac{1}{2\pi}\int_0^{2\pi} \frac{1-|z|^2}{\left|z-e^{i\theta}\right|^2} f\left(e^{i\theta}\right) \mathrm{d}\theta, \tag{1.2.2}$$

and moreover that any function u on $\mathbb{D}$ defined in this way is bounded and harmonic. The equation (1.2.2) therefore yields an isomorphism between the Banach spaces of bounded harmonic functions on the unit disc and bounded measurable functions on the unit circle, both with the uniform norm. Moreover, the right hand side of the expression is equal to the expectation of $f(B_T)$, where B is Brownian motion started at z, and T is the first time B hits $\partial\mathbb{D}$.

This example led Furstenberg, in his pioneering works [91–93], to define the notion of a **Poisson boundary** of a graph in the discrete setting. A **compactification** $\overline{G}$ of a graph G is a compact, metrisable topological space containing the vertex set of G as a dense, discrete subset. Given a compactification $\overline{G}$, we write $\partial G = \overline{G} \setminus V$. We say that the (boundary of the) compactification is a **Poisson boundary** of a transient graph G if the random walk on G converges in the compactification a.s., and, for every bounded harmonic function h on G, there exists a bounded measurable function (unique up to almost everywhere equivalence) $f : \partial G \to \mathbb{R}$ such that

$$h(v) = \mathbf{E}_v\, f\left(\lim_{n\to\infty} X_n\right)$$

for every vertex v. (Note that the limit point of the walk is a shift invariant function of the walk, so that if we start with a bounded function f on the boundary then the function h defined as above is certainly bounded and harmonic.) For example, the one-point compactification is a Poisson boundary of a transient graph G if and only if G is Liouville. Probabilistically, a compactification $\overline{G}$ of G is a Poisson boundary of G if the limit of a random walk in the compactification encapsulates all of the walk's limiting behaviour.

Although we have defined it topologically, the Poisson boundary is really a measure-theoretic notion, and two compactification Poisson boundaries of the same graph need not be homeomorphic (this is in contrast to the *Martin boundary*, defined in terms of positive harmonic functions, which is truly a topological object, see e.g. [201]).

For example, suppose that we have a transient tree. An **end** of the tree is an equivalence class of rays (i.e. infinite simple paths) in the tree, where two rays are equivalent if they have finite symmetric difference. The **ends compactification** of the tree is a compact topological space whose points are the vertices and ends of the tree, and where a sequence of vertices $\langle v_i \rangle_{i\geq 0}$ converges to an end ξ if, fixing some root vertex ρ of the tree arbitrarily, the geodesics in the tree connecting ρ to v_i converge to a ray in the equivalence class of the end ξ. In particular, a path converges to an end in the ends compactification of a tree if and only if the path is transient, i.e. visits each vertex of the tree at most finitely often, and so the random walk on a tree converges a.s. in the ends compactification if and only if the tree is transient. Now, if a function converges along a transient path, in a tree, it must also converge along any ray corresponding to the end that that path converges to. If h is a bounded harmonic function on a transient tree and X is a random walk started at ρ, then, by the Martingale convergence theorem, the limit of h along X exists a.s., and it follows that h converges a.s. along the ray starting at ρ that corresponds to the end that X converges to. This allows us to define a bounded function f on the boundary by

$$f(\xi) = \lim_{n\to\infty} h(v_i)$$

where $\langle v_i \rangle_{i\geq 0}$ is a ray corresponding to the end ξ. We clearly have that

$$h(v) = \mathbf{E}_v\, f\left(\lim_{n\to\infty} X_n\right),$$

so that the ends compactification is a Poisson boundary of any transient tree.

In general, it is not hard to show that a compactification Poisson boundary exists for any transient graph. For example, the Martin boundary is also a Poisson boundary (but easier constructions also exist). However, giving a *geometric* construction of the Poisson boundary, or showing that some particular compactification is indeed a Poisson boundary, can be highly non-trivial. In the next section, we shall see that for *bounded degree triangulations*, there is a beautiful geometric construction of the Poisson boundary using circle packing.

1.2.5 Infinite electrical networks and harmonic Dirichlet functions

In Section 1.2.1, we saw how to define the effective resistance between two sets in a finite graph, and also the effective resistance from a finite set to infinity in an infinite graph. What about the effective resistance between two finite sets in an infinite graph? It turns out that there is more than one reasonable way to define this, leading to different quantities that have different significance, since unit currents in the graph might no longer be unique. See [173, Chapter 9] for a detailed treatment of the material covered in this subsection.

Perhaps the most obvious way to define effective resistances in an infinite graph is to use an exhaustion by induced subgraphs as we did before. This yields the **free effective resistance**. Let G be an infinite graph, let $\langle V_n \rangle_{n\geq 1}$ be an exhaustion of G, and let G_n be the subgraph of G induced

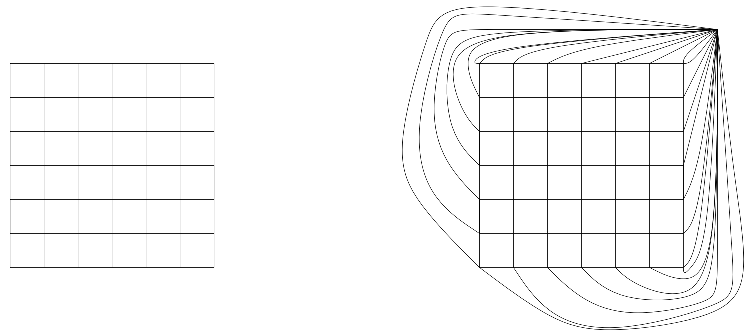

Figure 1.2: Free (left) and wired (right) boundary conditions on a 6×6 square in $\mathbb{Z}^2$.

by V_n for each $n \geq 1$. The free effective resistance between two finite sets A and B in G is defined to be

$$\mathscr{R}^{\mathrm{F}}_{\mathrm{eff}}(A \leftrightarrow B; G) = \lim_{n\to\infty} \mathscr{R}_{\mathrm{eff}}(A \leftrightarrow B; G_n).$$

Exercise. Prove that this limit exists and does not depend on the choice of exhaustion.

We call an antisymmetric edge function on an infinite graph G a **current** if it satisfies Kirchhoff's node and cycle laws, and has finite energy. It turns out that the unit current flows from A to B in the graphs G_n also converge to a current in G, called the **free unit current flow** from A to B, whose energy is exactly the free effective resistance.

Free effective resistances also obey a version of the Thomson and Dirichlet principles:

$$\mathscr{R}^{\mathrm{F}}_{\mathrm{eff}}(A \leftrightarrow B) = \\ \inf\{\mathcal{E}(\theta) : \theta \text{ is a finitely supported flow from } A \text{ to } B \text{ with strength at least one}\}$$

and

$$\mathscr{R}^{\mathrm{F}}_{\mathrm{eff}}(A \leftrightarrow B)^{-1} = \\ \inf\{\mathcal{E}(f) : f : V \to \mathbb{R} \text{ is a function with } f(v) \geq 1 \text{ for all } v \in A \text{ and } f(v) \leq 0 \text{ for all } v \in B\}.$$

What if we instead take the infimal energy over *all* flows, not just those that are finitely supported? This leads to the **wired effective resistance**, defined by

$$\mathscr{R}^{\mathrm{W}}_{\mathrm{eff}}(A \leftrightarrow B) = \inf\{\mathcal{E}(\theta) : \theta \text{ is a flow from } A \text{ to } B \text{ with strength at least one}\}.$$

In fact, it can be shown that this infimum is a minimum, that there is a unique flow obtaining this minimum, and that this flow is a unit current flow from A to B, called the **wired unit current flow** from A to B.

Wired effective resistances and currents can also be obtained from an exhaustion as follows. Let $\langle V_n \rangle_{n\geq 0}$ be an exhaustion of an infinite graph G. For each n, we also construct a graph G^*_n from

G by gluing (wiring) every vertex of $G \setminus V_n$ into a single vertex, denoted ∂_n, and deleting all the self-loops that are created. We identify the set of edges of G_n^* with the set of edges of G that have at least one endpoint in V_n. Then we have

$$\mathscr{R}_{\text{eff}}^{\text{W}}(A \leftrightarrow B) = \lim_{n\to\infty} \mathscr{R}_{\text{eff}}(A \leftrightarrow B; G_n^*),$$

and the unit current flow from A to B in G_n^* converges to the wired unit current flow from A to B.

When are the free and wired effective resistances between any two sets the same? We say that a graph has **unique currents** if for any two finite sets A and B, there is a unique unit current flow from A to B.

Proposition. *Let G be an infinite graph. Then the wired and free effective resistances between any two sets are equal if and only if G has unique currents, if and only if G does not admit non-constant harmonic functions of finite Dirichlet energy.*

This proposition tells us to expect non-constant harmonic functions of finite Dirichlet energy to play an important role in the electrical theory of infinite graphs. When do such functions exist?

Proposition. *Let G be an infinite graph. If G admits non-constant harmonic functions of finite Dirichlet energy, then G admits non-constant bounded harmonic functions of finite Dirichlet energy. In particular, G is non-Liouville.*

Proposition. *Let G be an amenable transitive graph. Then G does not admit harmonic functions of finite Dirichlet energy.*

It is more subtle to determine whether or not a nonamenable transitive graph admits non-constant harmonic functions of finite Dirichlet energy. (For Cayley graphs of groups, it is equivalent to the positivity of the *first ℓ^2-Betti number* of the group, see [173, Chapter 8.10].) Surprisingly, it is not monotone in how 'large' or 'expansive' the graph is. For example, if G is a bounded degree graph rough isometric to d-dimensional hyperbolic space, then G admits non-constant harmonic Dirichlet functions if and only if $d = 2$. Moreover, if G_1 and G_2 are any two infinite graphs, then the direct product of G_1 and G_2 does not admit any non-constant harmonic Dirichlet functions.

1.3 Circle packing

1.3.1 Planar graphs and maps

A **planar graph** is a graph that can be drawn in the plane so that no two edges intersect. Given a graph, a **proper embedding** of the graph into an oriented surface S is a drawing of the graph in the surface such that each compact set in the domain intersects at most finitely many edges of the graph, and each face of the drawing (i.e. connected component of the complement of the drawing) is homeomorphic to a disc. (For embeddings with infinite faces in multiply-connected surfaces there is an additional technical condition required of proper embeddings which will not concern

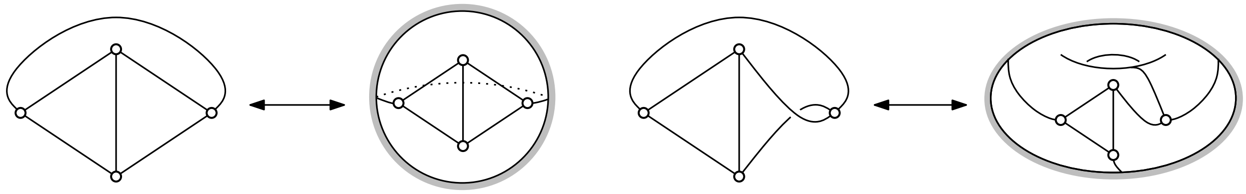

Figure 1.3: Different maps with the same underlying graph. A map is determined by a graph together with a cyclic ordering of the oriented edges emanating from each vertex.

us here.) A **map** is a graph together with an equivalence class of proper embeddings, where two proper embeddings are considered equivalent if there is an orientation preserving homeomorphism mapping one surface to the other that sends one embedding to the other. The map is **planar** if the surface can be taken to be a domain $D \subset \mathbb{C} \cup \{\infty\}$, and is **simply connected** if the surface can be taken to be the plane (or, equivalently, the open disc). Thus, a graph is planar if and only if it is the underlying graph of some planar map.

It turns out (see [184]) that maps can be described combinatorially as graphs together with cyclic permutations $\{\sigma_v : v \in V\}$ of the oriented edges emanating from each vertex, which specify the clockwise order of these edges in an embedding, so one does not need to be a topologist to study them.

The **dual** of a map is the map $M^\dagger$ that has the faces of M as vertices, the vertices of M as faces, and for each edge e of M has an edge $e^\dagger$ connecting the faces of M that are on either side of e. We say that a map has **bounded codegrees** if its faces have a bounded number of sides, or, equivalently, if its dual has bounded degrees.

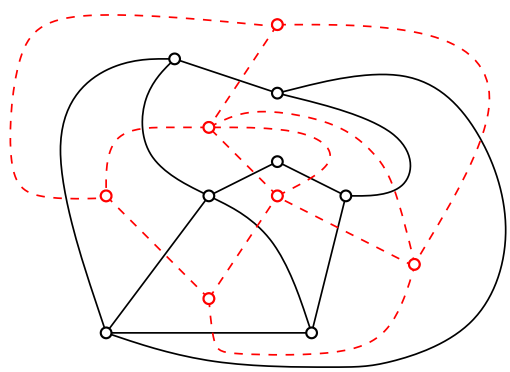

Figure 1.4: A finite planar map (black, solid) and its dual (red, dashed).

1.3.2 Circle packing

There are many ways to draw a planar map, which may or may not be useful for analyzing the behaviour of random walk and other processes on the graph. One of the best ways is given by the *circle packing theorem*. This theorem yields a canonical method of drawing planar graphs, closely connected to conformal mapping, and endows the graph with a geometry that, for many purposes, is better than the usual graph metric. Indeed, for bounded degree triangulations, a comprehensive theory has been developed by Angel, Barlow, Gurel-Gurevich, and Nachmias [17], and Chelkak [69], showing that the random walk on the circle packing behaves very similarly to a quasiconformal image of standard planar Brownian motion: Effective resistances, heat kernels, and harmonic measures on the graph can each be estimated in terms of the corresponding Brownian quantities. See [202] and [215] for introductions to the theory of circle packing. The interested reader may also enjoy making their own circle packings using Ken Stephenson's CirclePack software [214].

In this section, we will review the main theorems of circle packing as they relate to random walks and potential theory. Later, we will apply this theory to study uniform spanning forests of planar graphs, and will also study circle packings of unimodular random triangulations, two of the major topics of original research in this thesis. Moreover, we will develop the dichotomy between parabolic and hyperbolic bounded degree planar maps, which will motivate the development of a similar dichotomy for unimodular random planar maps, a further major topic of original work in this thesis.

A **circle packing** of a planar map G is a set of discs $P = \{P(v) : v \in V\}$ with disjoint interiors in the Riemann sphere $\mathbb{C} \cup \{\infty\}$, one for each vertex of G, such that two discs are tangent if and only if their corresponding vertices are adjacent in G. The existence and uniqueness theorem for circle packings, or simply the Circle Packing Theorem, is in our opinion one of the most beautiful results in all of mathematics. It was first discovered by Koebe in 1936 [156] as a corollary to his work on the Riemann mapping theorem for multiply connected domains, but went largely forgotten until Thurston [221] (who was unaware of Koebe's proof) rediscovered it as a corollary to the work of Andreev [14] on convex polyehdra. Due to this storied history, the Circle Packing Theorem is often called the Koebe-Andreev-Thurston Theorem. Like graphs, maps are said to be **simple** if they do not have any loops or multiple edges.

Theorem. *Every finite, simple planar map has a circle packing in the Riemann sphere. If the map is a triangulation, then its circle packing is unique up to Möbius transformations.*

The circle packing theorem was extended to infinite, simple, simply connected triangulations by He and Schramm [121, 122]. Their theorem can be thought of as a discrete analogue of the Uniformization Theorem for Riemann Surfaces. (Indeed, a celebrated theorem of Rodin and Sullivan [200] states that circle packings can be used to approximate conformal maps.) The **carrier** of the circle packing of a triangulation is the union of the discs in the triangulation together with the curved triangular regions surrounded by three circles that correspond to the triangles of the map. A circle packing is said to be **in** a domain $D \subseteq \mathbb{C} \cup \{\infty\}$ if its carrier is D.

Theorem (Schramm 1991, He and Schramm 1993)**.** *Let T be an infinite, simple, simply connected triangulation. Then T admits a circle packing either in the plane or in the disc, but not both, and this circle packing is unique up to Möbius transformations of the plane or the disc as appropriate.*

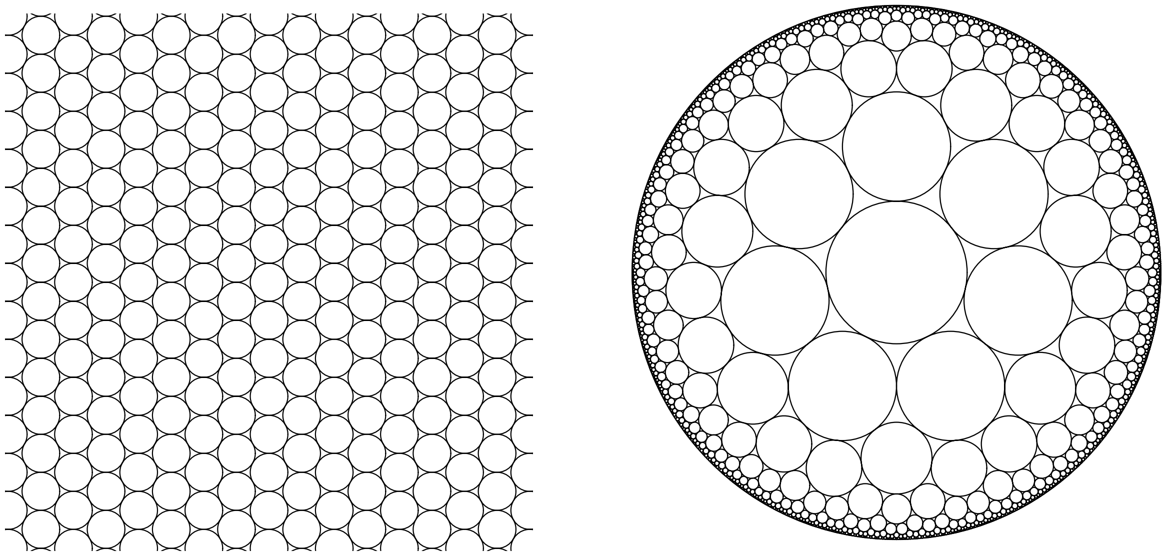

Figure 1.5: Circle packings of the 6-regular and 7-regular triangulations.

In light of this theorem, we call a simply connected triangulation **CP parabolic** if it can be circle packed in the plane, and **CP hyperbolic** otherwise. A rather trivial compactness argument (using the Ring Lemma, below) shows that every simple triangulation can be circle packed in *some* domain $D \subseteq \mathbb{C} \cup \{\infty\}$. It is much harder to show that we can circle pack inside a domain that is geometrically nice. (In fact, He and Schramm proved that, in the CP hyperbolic case, we can circle pack in *any* simply connected domain D strictly contained in the plane.)

Recall that the unit disc can be identified with the hyperbolic plane through the Poincaré disc model, and that, under this identification, circles in the unit disc and circles in the hyperbolic plane are the same, but have different centres and radii. Moreover, the Möbius transformations of the disc are exactly the orientation-preserving isometries of the hyperbolic plane. Thus, the above theorem tells us to expect a strong connection between CP hyperbolic triangulations and hyperbolic geometry.

He and Schramm [121] also pioneered the application of circle packing to probabilistic problems, proving the following remarkable theorem.

Theorem (He and Schramm 1995)**.** *Let T be an infinite, simple, bounded degree, proper plane triangulation. Then T is CP parabolic if and only if it is recurrent.*

The He-Schramm Type Theorem follows as an immediate consequence of the following estimate, which also gives a good flavour of the kind of analysis we can do with circle packings. Let us first introduce some notation. Given a triangulation T and a circle packing P of T, we write $z(v)$ and $r(v)$ for the (Euclidean) centre and radius, respectively, of the disc $P(v)$. Given $z \in \mathbb{C}$ and $R \geq r > 0$, we write $B_z(r)$ for the ball $\{w \in \mathbb{C} : |w - z| \leq r\}$, and $A_z(r, R)$ for the annulus $\{w \in \mathbb{C} : r \leq |w - z| \leq R\}$.

Lemma (Resistances across annuli). *Let T be a plane triangulation, and let P be a circle packing of T in some domain $D \subseteq \mathbb{C} \cup \{\infty\}$.*

1. *(Upper bound.) There exists a universal constant C such that the following holds. For every closed annulus $A_{z_0}(r, \alpha r) \subset D$ with $\alpha \geq 2$ such that $\{v : z(v) \in B_{z_0}(r)\} \neq \emptyset$, we have*

$$\mathscr{R}_{\text{eff}}\Big(\big\{v : z(v) \in B_{z_0}(r)\big\} \leftrightarrow \big\{v : z(v) \in D \setminus B_{z_0}(\alpha r)\big\}\Big) \leq C \log \alpha.$$

2. *(Lower bound.) Suppose that T has bounded degrees. Then there exists a constant C' depending on the maximal degree of T such that the following holds. For every closed annulus $A_{z_0}(r, \alpha r)$ with $\alpha \geq 2$ (not necessarily contained in D) such that $\{v : P(v) \subseteq B_{z_0}(r)\} \neq \emptyset$, we have*

$$\mathscr{R}^F_{\text{eff}}\Big(\big\{v : P(v) \subseteq B_{z_0}(r)\big\} \leftrightarrow \big\{v : z(v) \in D \setminus B_{z_0}(\alpha r)\big\}\Big) \geq C' \log \alpha.$$

The idea is that we can use the geometry of the circle packing to define a flow of sufficiently low energy (to obtain the upper bound via Thomson's principle), and functions of sufficiently low energy (to obtain the lower bound via the Dirichlet principle).

For the upper bound, we can define a flow that dissipates radially outwards in a roughly symmetric fashion. (This is best done with the *method of random paths*, defining a random path by taking the path in the graph that interpolates a radial line segment at a uniformly random angle from the centre of the annulus.)

For the lower bound, the Ring Lemma of Rodin and Sullivan [200] is a crucial ingredient, which gives us geometric control of the circle packing if we can control the degrees of vertices in the triangulation (for instance, in the bounded degree case).

Theorem (The Ring Lemma; Rodin and Sullivan 1987). *There exists a sequence of constants $\langle k_n : n \geq 3\rangle$ such that the following holds. If P is a circle packing of a planar triangulation in some domain $D \subseteq \mathbb{C} \cup \{\infty\}$, and v is a vertex of G such that $P(v)$ does not contain ∞, then*

$$r(v)/r(u) \leq k_{\deg(v)}$$

for every vertex u adjacent to v.

In the setting of item 2. of the resistances across annuli lemma, the Ring Lemma implies that there exists some constant $\beta > 2$ depending on the maximal degree such that if $B_{z_0}(r)$ contains a circle then the set of circles intersected by each of the annuli $A_{z_0}(\beta^n r, 2\beta^n r)$ and $A_{z_0}(\beta^{n+1} r, 2\beta^{n+1} r)$ are disjoint for every n. This allows us to reduce to the case $\alpha = 2$, since we can then add up the resistances across each of the $\log_\beta(\alpha)$ many annuli $A_{z_0}(\beta^n r, 2\beta^n r)$ by the series law. The case $\alpha = 2$ can then be handled by normalizing so that $z = 0$ and $r = 1/2$, and then using the function $|z|$ in Dirichlet's principle to obtain a lower bound on the free effective resistance.

This resistance estimate has many further uses. For example, it can be used to prove our next milestone theorem about circle packing, the Benjamini-Schramm Convergence Theorem [47].

Theorem (Benjamini and Schramm 1996)**.** *Let T be a CP hyperbolic, simply connected, bounded degree triangulation, and let P be a circle packing of T in the disc. Then the random walk on T converges to a point in the boundary of the disc almost surely, and the law of the limit point has full support and no atoms.*

Here, 'the law of the limit point has full support' means that every interval in the boundary of positive length has a positive probability to contain the limit point of the random walk, while 'the law of the limit point has no atoms' means that for any particular point in the boundary, the limit of the walk does not equal that point a.s.

In fact, there is an easy proof of the convergence and nonatomicity parts of this theorem using the resistances across annuli lemma above, first observed in the author's work with Yuval Peres [132], which we now sketch. If v is a vertex and A is a set of vertices in a transient graph G, then we have the estimate

$$\mathbf{P}_v(\tau_A < \infty) \leq \frac{\mathscr{R}_{\mathrm{eff}}(v \to \infty)}{\mathscr{R}^{\mathrm{F}}_{\mathrm{eff}}(v \leftrightarrow A).}$$

Suppose we have a bounded degree triangulation circle packed in the disc, and that v is a fixed vertex of the triangulation. Then we have that there is some constant $C < \infty$ such that

$$\mathbf{P}_v(\text{the random walk started at } v \text{ ever hits } \{v : P(v) \subseteq B_\xi(\varepsilon)\}) \leq C \log(1/\varepsilon)^{-1}$$

for all $\xi \in \partial\mathbb{D}$ and all ε sufficiently small. In particular, for each point $\xi \in \partial\mathbb{D}$, we have

$$\mathbf{P}_v\left(\xi \text{ is in the closure of the set } \{z(X_n) : n \geq 0\}\right) = 0. \tag{1.3.1}$$

We can now deduce convergence and nonatomicity: Observe that, for topological reasons, the intersection of $\partial\mathbb{D}$ with the closure of the set $\{z(X_n) : n \geq 0\}$ is an interval a.s. (The key point powering this is that $|z(X_n)| \to 1$ and $r(X_n) \to 0$ a.s. as $n \to \infty$ by transience.) The estimate (1.3.1) implies that the expected length of this interval is zero, and so the interval must be a point a.s., establishing convergence. The fact that the law of the limit point does not have any atoms is obvious from (1.3.1).

An immediate consequence of the Benjamini-Schramm convergence theorem is that every transient, bounded degree, simple, proper plane triangulation is non-Liouville: For any bounded, measurable function $f : \partial\mathbb{D} \to V$, we can define a harmonic function on T by

$$h(v) = \mathbf{E}_v f\left(\lim_{n\to\infty} z(X_n)\right),$$

and this function will be non-constant provided that f is not (almost everywhere) constant. In fact, it is not too hard to reduce the case of a general bounded degree planar graph to this case, yielding the following corollary.

Corollary (Benjamini and Schramm 1996). *Every transient, bounded degree planar graph is non-Liouville.*

We remark that Benjamini and Schramm also gave a different proof of this theorem using a different embedding, the *square tiling*, for which they also proved a convergence theorem [48].

More recently, Angel, Barlow, Gurel-Gurevich, and Nachmias [17] showed that *every* harmonic function on a bounded degree, simple, simply connected triangulation can be represented by a function on the unit circle in this way.

Theorem (Angel, Barlow, Gurel-Gurevich, and Nachmias 2013). *Let T be a bounded degree, CP hyperbolic, simple, simply connected triangulation, and let P be a circle packing of T in the unit disc $\mathbb{D}$. Then the circle packing compactification is a Poisson boundary of T.*

Recall from earlier that the probabilistic interpretation of this theorem is that the limit point of the random walk on the unit circle encapsulates *all* the information about the limiting behaviour of the random walk. The analogous theorem for square tiling was proven slightly earlier by Georgakopoulos [100].

The original proof of this theorem was highly analytic. In order to prove it, the authors employed sophisticated machinery to prove various estimates for random walks on bounded degree circle packings. Roughly speaking, these estimates allow one to compare the random walk on a bounded degree circle packing in a domain to a (quasiconformal image of) standard planar Brownian motion in that domain. In particular, the authors proved heat kernel estimates, diffusivity-type estimates, and harmonic measure estimates. These estimates were further developed by Chelkak [69], whose work implies, among other things, that the harmonic measure on the unit disc of a bounded degree CP hyperbolic triangulation is Hölder continuous. (He considered finite triangulations with boundary; the theorem below follows from his work by taking limits.) Here, given a bounded degree simple triangulation circle packed in the unit disc, the **harmonic measure** from v, denoted ω_v, is the law of the limit point in the unit circle of the random walk started at v.

Theorem (Chelkak 2013). *Let T be a bounded degree, CP hyperbolic, simple, simply connected triangulation, let P be a circle packing of T in the unit disc, and let v be such that $z(v) = 0$. Then there exist positive constants $\alpha \leq \beta$ and C (depending only on the maximal degree of T) such that*

$$C^{-1}|I|^{\alpha} \leq \omega_v(I) \leq C|I|^{\beta}$$

for every interval $I \subset \partial\mathbb{D}$.

These estimates are interesting in their own right, and have several further applications. Indeed, they also allowed those authors to prove that the unit circle is the Martin boundary of the triangulation, a much stronger result that gives a representation formula for all *positive* harmonic functions on the triangulation. Later, they also will play an important role in our calculation, carried out in joint work with Nachmias [130] (Chapter 10 of this thesis), of the critical exponents for uniform spanning forests of planar graphs.

Theorem (Angel, Barlow, Gurel-Gurevich, and Nachmias 2013)**.** *Let T be a bounded degree, CP hyperbolic, simple, simply connected triangulation, and let P be a circle packing of T in the unit disc $\mathbb{D}$. Then the unit circle is a Martin boundary of T. That is, the following hold.*

1. *For every two vertices u and v of T, the Radon-Nikodym derivative of the harmonic measures*

$$\frac{\mathrm{d}\omega_v}{\mathrm{d}\omega_u} : \partial\mathbb{D} \to \mathbb{R}$$

is continuous.

2. *For every positive harmonic function h on T and every vertex u of T, there exists a unique probability measure μ on $\partial\mathbb{D}$ such that*

$$h(v) = h(u) \int \frac{\mathrm{d}\omega_v}{\mathrm{d}\omega_u}(\xi)\, \mathrm{d}\mu(\xi)$$

for every $v \in V$.

It turns out, however, that the heavy machinery developed by Angel, Barlow, Gurel-Gurevich, and Nachmias is not in fact required for the Poisson boundary result. In joint work with Yuval Peres in 2015 [132] (Chapter 8 of this thesis), we proved the following theorem by a much more elementary argument, which implies the Poisson boundary results for both circle packing and square tiling.

Theorem (H. and Peres 2015)**.** *Let G be a planar map, and let z be an embedding of G into a domain $D \subseteq \mathbb{D}$ such that $\langle z(X_n)\rangle_{n\geq 0}$ converges to a point in $\partial\mathbb{D}$ almost surely, and the law of the limit point is non-atomic. Then the compactification of G given by taking the closure of $z(V)$ in $\overline{\mathbb{D}}$ is a Poisson boundary of G.*

Thus, once one has established a.s. convergence of the random walk to a boundary point in an embedding of a planar graph in the disc, the identification of the Poisson boundary follows for free: No further geometric control on the embedding is needed whatsoever. The proof of this theorem was adapted from our work with Angel, Nachmias, and Ray [21] on unimodular random triangulations of unbounded degree, for which the analytic approach was not available.

1.3.3 Harmonic Dirichlet functions on planar graphs

Besides proving that every transient, bounded degree planar graph is non-Liouville, the Benjamini-Schramm convergence theorem also implies the even stronger result that every transient, bounded degree planar graph admits non-constant harmonic functions of finite Dirichlet energy.

If $f : V \to \mathbb{R}$ is a function and X_t is the *continuous time* random walk[3] on G, and we define f_t by

$$f_t(v) = \mathbf{E}_v \left[f(X_t) \right],$$

then it turns out that $\mathcal{E}(f_t)$ is a decreasing function of t. It follows that if f is a function of finite energy such that the limit

$$\lim_{n\to\infty} f(X_n)$$

exists a.s. and is not concentrated on a point (the existence of this limit, and its law if it exists, does not depend on whether we use the continuous or discrete time random walk), then

$$f_\infty(v) := \mathbf{E}_v \left[\lim_{n\to\infty} f(X_n) \right] = \lim_{t\to\infty} \mathbf{E}_v f(X_t)$$

is a non-constant harmonic function with $\mathcal{E}(f_\infty) \leq \mathcal{E}(f) < \infty$. (In fact, Ancona, Lyons, and Peres [12] proved that the limit of $f(X_n)$ as $n \to \infty$ always exists a.s. whenever f is a Dirichlet function on a transient graph, from which it is possible to deduce the convergence part of the Benjamini-Schramm convergence theorem.)

Now, if P is a circle packing of a bounded degree CP hyperbolic triangulation in the unit disc, then the function z assigning each vertex to its center has finite energy, since

$$\mathcal{E}(z) \leq \sum_{v\in V} \deg(v) r(v)^2 \leq \max_{v\in V} \deg(v) \sum_{v\in V} r(v)^2 \leq \max_{v\in V} \deg(v).$$

Thus, it follows from the Benjamini-Schramm convergence theorem and the above discussion that every transient, bounded degree, simply connected, simple plane triangulation admits non-constant harmonic functions of finite Dirichlet energy, namely $z_\infty(v) = \mathbf{E}_v[\lim_{n\to\infty} z(X_n)]$. With a little extra work to remove some of these hypotheses, we arrive at the following dichotomy.

Corollary (Benjamini and Schramm 1996)**.** *Let G be a bounded degree planar graph. Then the following are equivalent.*

1. *G is transient.*
2. *G admits non-constant positive harmonic functions.*
3. *G is non-Liouville, i.e., admits non-constant bounded harmonic functions.*
4. *G admits non-constant harmonic functions of finite Dirichlet energy.*

The implications (4) ⇒ (3) ⇒ (2) ⇒ (1) hold on any graph. We shall later see that an even more far reaching dichotomy holds for *unimodular random planar maps*, without the assumption of bounded degree.

[3]This is the walk that has a semigroup of transition operators given by $P_t = e^{-t\Delta}$, where $\Delta = \nabla^*\nabla$ is the Laplacian of G

Let us take a moment to mention a relevant work in progress. Now that we know that non-constant harmonic Dirichlet functions exist on any bounded degree, CP hyperbolic, simple, proper plane triangulation, it is natural to wonder whether there is a geometric representation for the space of all such functions, as we had for the spaces of bounded harmonic functions and positive harmonic functions. In upcoming work, we give such a representation. For each function $f : \partial\mathbb{D} \to \mathbb{R}$, we define

$$\mathcal{D}(\phi) := \frac{1}{4\pi^2}\int_{\partial\mathbb{D}}\int_{\partial\mathbb{D}}\left|\frac{\phi(\xi)-\phi(\zeta)}{2\sin\left(\frac{1}{2}|\xi-\zeta|\right)}\right|^2 \mathrm{d}\xi\,\mathrm{d}\zeta < \infty.$$

where the integrals are taken with respect to arc length, and say f is **Douglas-integrable** if $\mathcal{D}(f) < \infty$. The space of Douglas-integrable functions is a Hilbert space with the inner product

$$f(o)g(o) + \mathcal{D}(f,g) = f(o)g(o) + \frac{1}{4\pi^2}\int_{\partial\mathbb{D}}\int_{\partial\mathbb{D}}\frac{(f(\xi)-f(\zeta))(g(\xi)-g(\zeta)))}{2\sin\left(\frac{1}{2}|\xi-\zeta|\right)}\,\mathrm{d}\xi\,\mathrm{d}\zeta,$$

where $o \in V$ is an arbitrary root vertex. In fact, $\mathcal{D}(f)$ is exactly the Dirichlet energy of the harmonic extension of f to the unit disc.

Theorem (H. 2017+)**.** *Let T be a transient, bounded degree, proper plane triangulation, and let P be a circle packing of T in the unit disc. Then the function*

$$f \longmapsto h, \quad \textit{where } h(v) = \mathbf{E}_v\, f\left(\lim_{n\to\infty} X_n\right)$$

is a bounded linear isomorphism from the space of Douglas-integrable functions on $\partial\mathbb{D}$ to the space of harmonic Dirichlet functions on T.

By a bounded linear isomorphism we mean a bounded linear map with a bounded linear inverse, and not necessarily an isometry.

A potentially surprising thing here is that the space of Douglas-integrable functions on the boundary is always the same; the expression for the Douglas energy does not involve the harmonic measure, or otherwise involve the triangulation in any way.

1.3.4 Double circle packing

Before moving on, let us mention a generalisation of the circle packing theorem which, in our opinion, is even more beautiful.

Let M be a planar map with vertex set V and face set F. A **double circle packing** of M is a pair of circle packings $P = \{P(v) : v \in V\}$ and $P^\dagger = \{P^\dagger(f) : f \in F\}$ satisfying the following conditions (see Figure 1.6):

1. (M **is the tangency map of** P.) For each pair of vertices u and v of M, the discs $P(u)$ and $P(v)$ are tangent if and only if u and v are adjacent in M. Moreover, for each vertex u, the discs corresponding to the vertices adjacent u appear around $P(u)$ in the clockwise order specified by the map structure σ_u.

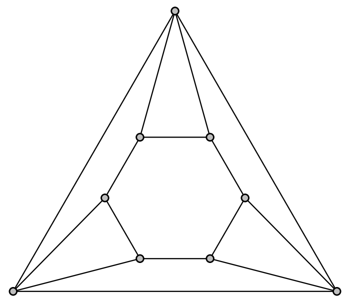

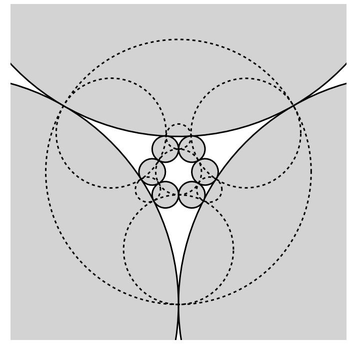

Figure 1.6: A polyhedral planar map and its double circle packing.

2. ($M^\dagger$ **is the tangency map of** $P^\dagger$.) For each pair of faces f and g of G, the discs $P^\dagger(f)$ and $P^\dagger(g)$ are tangent if and only if f and g are adjacent in $G^\dagger$. Moreover, for each face f, the discs corresponding to the vertices adjacent f appear around $P^\dagger(f)$ in the clockwise order specified by the map structure $\sigma^\dagger_f$.

3. (**Primal and dual circles are perpendicular**.) For each vertex v and face f of M, the discs $P^\dagger(f)$ and $P(v)$ have non-empty intersection if and only if f is incident to v, and in this case the boundary circles of $P^\dagger(f)$ and $P(v)$ intersect at right angles.

Recall that a graph is **3-connected** if the removal of any two vertices from the graph does not cause the graph to become disconnected. We call a map that is both simple and 3-connected **polyhedral**. Any embedding of a planar map that is not polyhedral must contain faces that are not convex, and it follows that any finite planar map admitting a double circle packing must be polyhedral. Conversely, Thurston's interpretation of Andreev's Theorem [179, 221] implies that every finite, polyhedral, planar map admits a double circle packing (see also [58]). The corresponding infinite theory was developed by He [120], who proved that every infinite, polyhedral, simply connected map M with locally finite dual admits a double circle packing in either the Euclidean plane or the hyperbolic plane (but not both) and that this packing is unique up to Möbius transformations. As before, we say that M is CP parabolic or CP hyperbolic as appropriate.

In our work with Nachmias [130], we proved a form of the Ring Lemma for double circle packings. We write $v \perp f$ to mean that the vertex v is incident to the face f.

Theorem (H. and Nachmias 2016)**.** *There exists a family of positive constants $\langle k_{n,m} : n \geq 3, m \geq 3\rangle$ such that if $(P, P^\dagger)$ is a double circle packing of a polyhedral planar map M in a domain $D \subseteq \mathbb{C}\cup\{\infty\}$ and v is a vertex of M such that $P(v)$ does not contain ∞, then*

$$r(v)/r(f) \leq k_{\deg(v),\max_{g\perp v}\deg(g)}$$

for all $f \in F$ incident to v.

Once the existence and uniqueness theorems and the Ring Lemma are in place, everything else that we have said about circle packings of simple triangulations extends more or less immediately to double circle packings of polyhedral maps of bounded codegree. In particular, as in the He-Schramm Theorem [121], CP hyperbolicity is equivalent to transience for simply connected polyhedral maps with bounded degrees and codegrees [120].

1.3.5 Circle packing and the degree

An interesting aspect of circle packing is that, in certain situations, we can deduce the CP type of a triangulation from the degrees of its vertices. This was first observed by Beardon and Stephenson [34], who proved the following.

Theorem (Beardon and Stephenson 1991)**.** *Let T be an infinite, simple, simply connected triangulation. If $\deg(v) \leq 6$ for every vertex v of T, then T is CP parabolic. If $\deg(v) \geq 7$ for every vertex v of T, then T is CP hyperbolic.*

Exercise. Show that there exists a constant $C > 1$ such that the following holds. Suppose T is a simple triangulation circle packed in a domain $D \subseteq \mathbb{C}$. Then

1. for every vertex v of T with $\deg(v) \leq 5$, there is a neighbour u of v such that $r(u)/r(v) \geq C$.
2. for every vertex v of T with $\deg(v) = 6$, there is a neighbour u of v such that $r(u)/r(v) \geq 1$ and a neighbour w of v such that $r(w)/r(v) \leq 1$.
3. for every vertex v of T with $\deg(v) \geq 7$, there is a neighbour u of v such that $r(u)/r(v) \leq C^{-1}$.

Let $r_{\mathbb{H}}(v)$ denote the hyperbolic radius of a disc $P(v) \subset \mathbb{D}$.

Exercise. Show that there exists a constant $C > 1$ such that the following holds. Suppose T is a simple triangulation circle packed in a domain $D \subseteq \mathbb{D}$. Then for every vertex v of T with $\deg(v) \leq 6$, there is a neighbour u of v such that $r_{\mathbb{H}}(u)/r_{\mathbb{H}}(v) \geq C$.

Exercise. Deduce the above theorem of Beardon and Stephenson from the previous two exercises.

We shall later see that for unimodular random simply connected triangulations, the circle packing type is determined by the *average* degree.

1.4 Percolation

In **Bernoulli bond percolation**, each edge of a (usually infinite) graph G is chosen randomly to be either deleted with probability $1-p$, or else retained with probability p. The random graph obtained in this way is denoted $G[p]$. Connected components of $G[p]$ are referred to as **clusters**.

Although we will have relatively little new to say about percolation in this thesis, the theory of percolation provides important motivation for our work on uniform spanning forests, and also leads naturally to the theory of unimodular random rooted graphs, another central topic of this thesis. Percolation has become a topic of central importance to modern probability theory, and has a huge literature surrounding it. The reader is referred to the beautiful, if slightly outdated, monograph of Grimmett [106] for a thorough treatment of the Euclidean ($\mathbb{Z}^d$) theory, and to [173] for the theory of percolation on more general graphs.

The first basic result about percolation, without which the model would not be nearly as interesting, is that for most graphs, percolation undergoes a *non-trivial phase transition*, meaning that the **critical probability**,

$$p_c(G) = \inf\{p \in [0,1] : G[p] \text{ has an infinite cluster almost surely}\},$$

is strictly between zero and one. Conjecturally, $p_c \in (0,1)$ for every transitive graph that is not rough isometric to $\mathbb{Z}$: This has been resolved in most cases, and remains open only for transitive graphs of intermediate volume growth (i.e., subexponential but superpolynomial). For example, a simple first moment argument shows that $p_c > 0$ for any bounded degree graph, since the expected number of open simple paths of length n starting at a vertex v is at most

$$p^n \left(\max_{u \in V} \deg(u)\right)^n,$$

which yields the bound $p_c \geq (\max_{u\in V} \deg(u))^{-1}$. For $\mathbb{Z}^2$, a similar first moment argument with simple curves surrounding the origin (known as a **Peierls argument** [189]) shows that $p_c(\mathbb{Z}^2) < 1$, while for $d \geq 2$ we clearly have that $p_c(\mathbb{Z}^d) \leq p_c(\mathbb{Z}^2)$.

The study of percolation can naturally be decomposed into the study of the subcritical ($p < p_c$), critical ($p = p_c$), and supercritical ($p > p_c$) regimes. Here, we shall be interested mostly in the critical and supercritical regimes.

At criticality, the central question concerns the existence or nonexistence of an infinite cluster. Indeed, perhaps the best known open problem in modern probability theory is to prove that there is no infinite cluster at criticality on $\mathbb{Z}^d$ for all $d \geq 2$. This problem was solved in two dimensions by Russo in 1981 [203], and for all $d \geq 19$ by Hara and Slade in 1994 [117]. More recently, Fitzner and van der Hoftstad [89] sharpened the methods of Hara and Slade to solve the problem for all $d \geq 11$. (It is expected that this method can in principle, and with great effort, be pushed to handle all $d \geq 7$. Dimensions $3, 4, 5$, and 6 are expected to require new approaches.)

In 1996, Benjamini and Schramm [51] proposed a systematic study of percolation on general transitive (and quasi-transitive) graphs. They made several conjectures and posed many questions. Here is one of them.

Conjecture (Benjamini and Schramm 1996)**.** *Let G be a transitive graph. If $p_c(G) < 1$, then $G[p_c]$ does not have any infinite clusters almost surely.*

It quickly emerged that it is often substantially easier to study percolation on *unimodular* transitive graphs than on general transitive graphs. A transitive graph $G = (V, E)$ is said to be **unimodular** if for every function $F : V^2 \to [0, \infty]$ that is automorphism equivariant in the sense that $F(\gamma u, \gamma v) = F(u, v)$ for every $u, v \in V$ and every automorphism γ of G, then, letting ρ be a fixed root vertex of G,

$$\sum_{v \in V} F(\rho, v) = \sum_{u \in V} F(u, \rho). \tag{1.4.1}$$

The equality (1.4.1) is referred to as the **Mass-Transport Principle** (MTP)[4]. We think of F as a rule for sending a non-negative amount of mass from each vertex to each other vertex, so that the mass-transport principle says that

$$\textit{mass out of the root} = \textit{mass into the root}.$$

It turns out that most transitive graphs one encounters 'in the wild' are unimodular. In particular, every amenable transitive graph, every Cayley graph of a finitely generated group, and every transitive, simply connected, planar map with locally finite dual has the property of unimodularity.

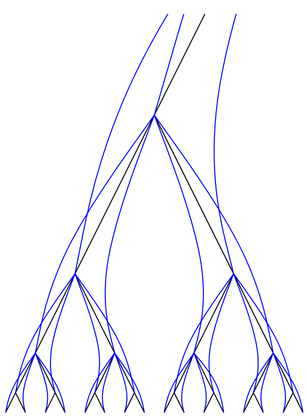

Figure 1.7: The grandparent graph.

An example of a transitive graph that is *not* unimodular is the **grandparent graph**, defined as follows. Take a 3-regular tree, and draw it in the plane so that every vertex has one 'parent' above it and two 'children' below it. The grandparent graph is formed by adding to the 3-regular tree an edge connecting each vertex to its grandparent, that is, the parent of its parent. We can define an automorphism equivariant function on the grandparent graph by

$$F(u, v) = \mathbb{1}(v \text{ is } u\text{'s grandparent}).$$

Since every vertex has one grandparent but four grandchildren, we have

$$\sum_{v \in V} F(\rho, v) = 1 \quad \text{but} \quad \sum_{u \in V} F(u, \rho) = 4,$$

so that the Mass-Transport Principle does not hold for the grandparent graph.

[4]The history of this definition is as follows: A locally compact group is said to be unimodular if its left and right Haar measures coincide. A transitive graph is unimodular in our sense if and only if its group of automorphisms has a transitive unimodular subgroup.

A major achievement in the theory of critical percolation outside of the Euclidean setting was made in 1999 by Benjamini, Lyons, Peres, and Schramm [43]. This theorem came out of work [45] by a subset of the same authors on arbitrary (non-Bernoulli) automorphism invariant percolation processes, which we shall have more to say about later.

Theorem (Benjamini, Lyons, Peres, and Schramm 1999)**.** *Let G be a nonamenable, unimodular, transitive graph. Then $G[p_c]$ has no infinite clusters almost surely.*

Partial progress was made in the nonunimodular case by Peres, Pete, and Scolnicov [192], who proved there are no infinite clusters at criticality a.s. on certain well-known examples of nonunimodular transitive graphs, including decorated trees (like the grandparent graph) and nonamenable Diestel-Leader graphs, and by Timár [222], who proved that there is at most one infinite cluster at criticality on any nonunimodular transitive graph a.s.

In 2016 [129] (Chapter 2 of this thesis), we found an elementary proof that the following estimate holds in any transitive graph:

$$\inf\left\{\mathbb{P}\left(x \text{ is connected to } y \text{ in } G[p_c]\right) : x, y \in V,\, d(x,y) \leq n\right\} \leq \left(\liminf_{r\to\infty} |B(x,r)|^{1/r}\right)^{-n}.$$

This estimate readily implies that there cannot be a *unique* infinite cluster at p_c in any transitive graph of exponential growth. However, the existence of *multiple* infinite clusters at criticality had already been ruled out by previous work (the amenable and unimodular nonamenable cases are discussed in the following subsection, while the nonunimodular case is handled by the work of Timár mentioned above). Thus, we obtained the following theorem.

Theorem (H. 2016)**.** *Let G be a transitive graph with exponential growth. Then $G[p_c]$ has no infinite clusters almost surely.*

1.4.1 The number of infinite clusters

Let us now turn to the supercritical regime. In this regime, we are particularly interested in understanding the geometry of the infinite clusters of $G[p]$. The most basic question concerns the *number* of infinite clusters. For this question, the foundational result was proven by Newman and Schulman in 1981 [187].

Theorem (Newman and Schulman 1981)**.** *Let G be a transitive graph. Then $G[p]$ has either no infinite clusters, a unique infinite cluster, or infinitely many infinite clusters almost surely for every $p \in [0,1]$.*

The key to this theorem is the **insertion tolerance** of Bernoulli percolation, which says that for any set of edges A and $p > 0$, the law of the subgraph $G[p] \cup A$ is absolutely continuous[5] with respect to the law of $G[p]$.

[5]Recall that the law of a random variable X is said to be absolutely continuous with respect to the law of a random variable Y if whenever $\mathscr{A}$ is a set such that Y is a.s. not in $\mathscr{A}$, then X is also a.s. not in $\mathscr{A}$.

The proof is as follows: Given a configuration $\omega \in \{0,1\}^E$ and an automorphism γ of G, define

$$\gamma\omega(e) = \omega(\gamma^{-1}(e))$$

for every edge e. It is an easy fact that Bernoulli bond percolation on any transitive graph is **ergodic**, meaning that if $\mathscr{A} \subseteq \{0,1\}^E$ is a set of configurations that is automorphism invariant in the sense that

$$\omega \in \mathscr{A} \Rightarrow \gamma\omega \in \mathscr{A}$$

for every configuration ω and automorphism γ, then we must have that $\mathbb{P}(G[p] \in \mathscr{A})$ is either zero or one. In particular, the number of infinite components is automorphism invariant, and we deduce that the number of infinite components is not random. That is, for every p, there exists $N_p \in \mathbb{N} \cup \{\infty\}$ such that $G[p]$ has N_p infinite clusters a.s.

Suppose for contradiction that $N_p \notin \{0,1,\infty\}$. If we take a large connected set of edges A, then with high probability A is incident to more than one of the infinite clusters of $G[p]$. Thus, the configuration $G[p] \cup A$ has positive probability to have strictly fewer infinite clusters than $G[p]$. But $G[p] \cup A$ is absolutely continuous with respect to $G[p]$ by insertion-tolerance, hence it must have N_p infinite clusters a.s., yielding a contradiction.

The Newman-Schulman Theorem can in fact be extended to any insertion-tolerant **automorphism invariant percolation** on a transitive graph. Here, an automorphism-invariant percolation is a random subgraph ω of G such that $\gamma\omega$ has the same distribution as ω for every automorphism γ of G. (In the case that ω is also ergodic the Newman-Schulman Theorem follows exactly as above. In general, it is possible to reduce to the ergodic case by taking an *ergodic decomposition*, which preserves insertion tolerance.)

The next major result on the number of clusters was due to Aizenmann, Kesten, and Newman in 1987 [3], who proved that percolation on $\mathbb{Z}^d$ has at most one infinite cluster a.s. for every $d \geq 1$. A much simpler proof of this theorem was found by Burton and Keane in 1989 [64], which was then generalised to all transitive amenable graphs by Gandolfi, Keane, and Newman in 1992 [97]. Like the Newman-Schulman Theorem, this theorem also extends to arbitrary insertion-tolerant, automorphism invariant percolations.

Theorem. *Let G be an amenable transitive graph, and let $p \in [0,1]$. Then $G[p]$ has at most one infinite cluster almost surely.*

It is a major open problem, first stated by Benjamini and Schramm [51], to prove that the converse holds. The **uniqueness threshold** is defined to be

$$p_u(G) = \inf\{p \in [0,1] : G[p] \text{ has a unique infinite cluster almost surely}\}.$$

Conjecture (Benjamini and Schramm 1996)**.** *Let G be a transitive graph. Then $p_c < p_u$ if and only if G is nonamenable.*

Not much progress has been made on this question; see [112] and [173, Chapter 9] for an account of what is known.

Since the existence of a unique infinite cluster is not monotone with respect to the configuration, the following theorem, due to Schonmann [205], and Häggström, Peres, and Schonmann [114], is far from obvious.

Theorem (Häggström, Peres, and Schonmann 1999)**.** *Let G be a transitive graph, and let $p_2 > p_1$. If $G[p_1]$ has a unique infinite cluster almost surely, then $G[p_2]$ has a unique infinite almost surely.*

In the unimodular case, this theorem can also be deduced from the following beautiful theorem of Lyons and Schramm [176], which, informally, says that infinite clusters of Bernoulli bond percolation on a unimodular transitive graph all 'look the same'. The theorem extends to insertion-tolerant automorphism invariant percolations.

Theorem (Lyons and Schramm 1999)**.** *Let G be a unimodular transitive graph, let $p \in [0,1]$, and let $\mathscr{A} \subseteq \{0,1\}^E$ be a measurable, shift invariant set. Then either every infinite cluster of $G[p]$ is in $\mathscr{A}$ or no infinite clusters of $G[p]$ are in $\mathscr{A}$ almost surely.*

For example, the theorem implies that either every infinite cluster of $G[p]$ is transient or every infinite cluster of $G[p]$ is recurrent a.s.

The main idea of the proof of the Lyons-Schramm Indistinguishability Theorem can be summarised as follows. The coexistence of infinite clusters of different types (i.e., of infinite clusters in $\mathscr{A}$ and not in $\mathscr{A}$) can be shown to imply that for some infinite cluster, there exist infinitely many *pivotal edges*, that is, closed edges that change the type of the cluster if they are inserted. Heuristically, this should contradicts the measurability of the property: The existence of pivotal edges far away from the origin – which by insertion tolerance are in some sense indifferent to being inserted and hence changing the type of an infinite cluster – should imply that we cannot approximate the event that the cluster at the origin is in $\mathscr{A}$ by an event depending only on finitely many edges of the configuration. (The ability to approximate a set in this way is the definition of the set being measurable.)

To see that the Lyons-Schramm Indistinguishability Theorem implies the Häggström-Peres-Schramm Uniqueness Monotonicty Theorem in the unimodular case, first observe that we can sample $G[p_1]$ by first sampling $G[p_2]$, and then performing Bernoulli p_1/p_2 bond percolation on $G[p_2]$. Consider the set

$$\mathscr{A} = \left\{ \omega \in \{0,1\}^E : \begin{array}{l} \omega \text{ spans a connected subgraph } H \text{ of } G \text{ such that Bernoulli} \\ p_1/p_2 \text{ percolation on } H \text{ has no infinite clusters almost surely} \end{array} \right\}$$

If $G[p_2]$ has infinitely many infinite clusters a.s., then, by the Lyons-Schramm Indistinguishability Theorem, either all these clusters are in $\mathscr{A}$, in which case $G[p_1]$ has infinitely many infinite clusters a.s. also, or else none of the clusters are in $\mathscr{A}$, in which case $G[p_1]$ does not have any infinite clusters a.s., concluding the proof.

Indistinguishability of infinite clusters can fail for nonunimodular transitive graphs. However, Häggström, Peres, and Schonmann [114] proved that indistinguishability still holds in this context provided that we restrict attention to what they called *robust* properties, of which the set $\mathscr{A}$ above is an example.

1.4.2 The geometry of infinite clusters

What can we say about the geometry of individual clusters? In general, we expect that if $G[p]$ has a unique infinite cluster, then this cluster will be similar to the original graph G in many ways. For example, it is conjectured that if G is a transient transitive graph, then every infinite cluster of $G[p]$ is transient a.s. for every $p > p_c$. The first result of this form was due to Grimmett, Kesten, and Zhang [104].

Theorem (Grimmett, Kesten, and Zhang 1993)**.** *Let $p > p_c(\mathbb{Z}^d)$. Then the unique infinite cluster of $\mathbb{Z}^d[p]$ is transient almost surely.*

The original proof of this theorem was rather difficult. A much simpler proof was found by Gabor Pete in 2008 [194], using isoperimetric techniques. An obstacle to applying isoperimetric methods to percolation is that the infinite percolation cluster will necessarily contain 'bad regions'. In particular, if G is a transitive graph and $p < 1$, it is easily seen that every infinite cluster of $G[p]$ must contain arbitrarily large sets of vertices with only one edge in their boundary. This is because such sets have positive probabilities to occur at the origin, and therefore must occur somewhere by ergodicity.

This difficulty led Benjamini, Lyons, and Schramm [45] to introduce the notion of the *anchored* isoperimetric inequalities, whose definition we already saw in Section 1.2.2. These are less sensitive to perturbations of the graph, and we can hope that if G is a sufficiently nice (e.g. transitive) graph that satisfies an anchored $\phi(t)$-isoperimetric inequality, then the infinite clusters of $G[p]$ will also satisfy an anchored $\phi(t)$-isoperimetric inequality for $p > p_c$. Chen and Peres [70] proved that this is indeed the case for graphs with anchored expansion, at least for sufficiently large p. A second proof, due to Pete, appeared as an appendix to the same paper.

Theorem (Chen, Peres, and Pete 2004)**.** *Let G be a graph with anchored expansion. Then there exists $p_0 < 1$ such that every infinite cluster of $G[p]$ has anchored expansion almost surely for every $p \geq p_0$.*

A much more general form of this result was later proven by Pete [194], a special case of which is as follows.

Theorem (Pete 2008)**.** *Let G be a Cayley graph of a finitely presented group (e.g. $\mathbb{Z}^d$) that is not rough isometric to $\mathbb{Z}$, and suppose that G satisfies a $\phi(t)$-isoperimetric inequality for some* *increasing ϕ. Then there exists $p_0 < 1$ such that every infinite cluster of $G[p_0]$ satisfies an anchored $\phi(t)$-isoperimetric inequality almost surely.*

By the Thomassen Criterion, these results imply that infinite clusters of $G[p]$ are transient for sufficiently large p for a large class of transient graphs G. In the case of $\mathbb{Z}^d$, it is possible to recover the full Grimmett-Kesten-Zhang Theorem by a renormalization argument.

Now suppose that G is non-Liouville. Does it follow that every infinite cluster of $G[p]$ is non-Liouville for every $p > p_c$. Conversely, if G is Liouville, does it follow that every infinite cluster of $G[p]$ is Liouville for every $p > p_c$? In general, this question is very poorly understood. (In contrast, the existence of harmonic *Dirichlet* functions on percolation clusters of unimodular transitive graphs is very well understood since the work of Gaboriau [94].) Similarly to the case of transitive graphs, the non-Liouville property for percolation clusters on unimodular transitive graphs is equivalent to the random walk on the cluster having positive speed. Thus, we can answer the question in the nonamenable unimodular case for large p by invoking the following beautiful theorem of Virág [227]. We shall see that much easier proofs, avoiding anchored expansion altogether, are possible once we introduce the notion of *invariant nonamenability.*

Theorem (Virág 2000)**.** *Let G be a bounded degree graph with anchored expansion. Then the random walk on G has positive speed almost surely, and for each vertex v there exist positive constants c and C such that*

$$p_n(v, v) \le Ce^{-cn^{1/3}}$$

for all $n \ge 1$.

1.5 Uniform spanning forests

The **Free Uniform Spanning Forest** (FUSF) and the **Wired Uniform Spanning Forest** (WUSF) of an infinite graph G are defined as weak limits of the uniform spanning trees on large finite subgraphs of G, taken with either free or wired boundary conditions respectively. First studied by Pemantle [190], the USFs are closely related many other areas of probability, including electrical networks [62, 154], Lawler's loop-erased random walk [44, 161, 228], sampling algorithms [197, 228], domino tiling [150], the Abelian sandpile model [137, 138, 178], the rotor-router model [126], and the Fortuin-Kasteleyn random cluster model [105, 109]. The USFs are also of interest in group theory, where the FUSFs of Cayley graphs are related to the ℓ^2-Betti numbers [94, 169] and to the fixed price problem of Gaboriau [95], and have also been used to approach the Dixmier problem [87].

Uniform spanning forests are also the major topic of this thesis, being the subject of four of the included papers and being of central importance in one other. One of our main goals will be to address the same questions for the uniform spanning forests that we asked about percolation in the previous section.

1.5.1 Uniform spanning trees

Let us start at the beginning. A **uniform spanning tree** of a finite graph G is simply a uniform random element from the set of spanning trees of G, that is, connected subgraphs of G that contain every vertex and no cycles. The connection between uniform spanning trees and electrical networks goes back as far as it could, all the way to the 1847 work of Kirchhoff [154] in which the laws of electrical networks were introduced.

Theorem (Kirchhoff's effective resistance formula)**.** *Let G be a finite graph. Then for every edge e of G,*

$$\mathbb{P}(e \textit{ is in a uniform spanning tree of } G) = \mathscr{R}_{\mathrm{eff}}(e^- \leftrightarrow e^+)$$

(Kirchhoff did not state this theorem probabilistically, but rather as a statement about the ratio of the number of spanning trees that include the edge to the total number of spanning trees.)

Kirchhoff's interest in this theorem was computational: He also invented a method to compute the number of spanning trees of a graph, the famous *Matrix-Tree Theorem.*

Theorem (Kirchhoff's Matrix-Tree Theorem)**.** *Let G be a finite connected graph with n vertices, and let $\lambda_1, \ldots, \lambda_{n-1}$ be the non-zero eigenvalues of the Laplacian of G. Then*

$$|\{\textit{spanning trees of } G\}| = \frac{1}{n}\prod_{i=1}^{n-1}\lambda_i.$$

Since the number of spanning trees that do not contain a given edge is equal to the number of spanning trees of the graph in which that edge has been deleted, Kirchhoff's effective resistance formula together with the Matrix-Tree theorem give a method for computing effective resistances in finite graphs. From our perspective, we will be more interested in using the formula the other way around, using the variational principles for the effective resistance to estimate the probability that an edge is in a uniform spanning tree (or forest).

How can we sample the uniform spanning tree of a finite graph? One very simple method is as follows. Pick some edge of the graph. We can compute the probability that the edge is in the uniform spanning tree (e.g. using the Matrix-Tree Theorem), and flip an appropriately biased coin to decide whether or not to include it. If we decide to include the edge, then we next contract it, i.e., identify the two endpoints of the edge. Otherwise, if we decided not to include the edge, we delete it from the graph. We then continue as in the first step, choosing another edge from the graph, choosing whether or not to include it by performing a Matrix-Tree calculation (in the modified graph) and flipping an appropriately biased coin, and then contracting or deleting the edge as appropriate. If we continue doing this until we have decided what to do with every edge, then the set of edges we chose to include will form a uniformly random spanning tree of G.

Exercise. Prove that this algorithm works.

A far-reaching extension of Kirchhoff's effective resistance formula, known as the Transfer-Current Theorem of Burton and Pemantle [62], allows one to calculate the probability that *any* set of edges is included in the forest in terms of electrical quantities. (The case of two edges had previously been studied by Brooks, Smith, Stone, and Tutte [60], who were interested by the connection to *square tiling*, which they invented.) Given two oriented edges e_1 and e_2, we define $Y(e_1, e_2)$ to be the current flowing through e_2 when a unit potential is placed on the endpoints of e_1. In other words,

$$Y(e_1, e_2) = \mathbf{P}_{e_2^+}\left(\tau^+_{e_1^+} < \tau^+_{e_1^-}\right) - \mathbf{P}_{e_2^-}\left(\tau^+_{e_1^+} < \tau^+_{e_1^-}\right).$$

We fix arbitrarily an orientation of each edge e, and think of Y as a matrix indexed by $E \times E$, which we call the **transfer-current matrix**.

Theorem (Burton and Pemantle 1993)**.** *Let G be an infinite graph. Then for any collection of edges $e_1, \ldots, e_n$ of G, we have*

$$\mathsf{UST}(\{e_1, \ldots, e_n\} \subseteq T) = \det\langle Y(e_i, e_j)\rangle_{1 \leq i,j \leq n}.$$

Note that the probabilities of all other events can be computed from the probabilities of the events of the form $\{\{e_1, \ldots, e_n\} \subseteq T\}$ using inclusion-exclusion.

1.5.2 Sampling using random walks

The Aldous-Broder Algorithm

Let G be a finite graph, and let $\langle X_n \rangle_{n \geq 0}$ be a random walk on G. For each vertex v of G, let $e(v, X)$ be the oriented edge pointing into v that is crossed by the walk X as it enters v for the first time, and define

$$\mathsf{AB}(X) = \{-e(v, X) : v \neq X_0\}$$

to be the collection of (reversed) first entry edges. It is not hard to see that $\mathsf{AB}(X)$ is a spanning tree, where the edges of the tree are oriented so that every vertex other than X_0 has exactly one oriented edge in the tree emanating from it. The following theorem, however, is surprising at first.

Theorem (Broder 1989, Aldous 1990)**.** $\mathsf{AB}(X)$ *is a uniform spanning tree of G.*

Note in particular that we can start the walk wherever we want, and this does not change the distribution of the tree. This theorem gives another method of sampling the uniform spanning tree of a finite graph, named the **Aldous-Broder algorithm** after its inventors [8, 59].

Exercise. Use the Aldous-Broder algorithm to prove Kirchhoff's effective resistance formula.

Wilson's algorithm

The **loop-erasure** of a path in a graph is formed by erasing cycles from the path chronologically as they are created. (The loop-erasure is only defined for paths that are either finite or transient in the sense that they visit each vertex of the graph at most finitely often.) The loop-erasure of simple random walk is called **loop-erased random walk**, and was first studied by Lawler [161]. Let $\{v_0, v_1, \ldots, v_n\}$ be an enumeration of the vertices of a finite graph G and define a sequence of trees T_i in G as follows:

1. Let T_0 have vertex set v_0 and no edges.

2. Given T_i, start an independent random walk from v_{i+1} stopped when it hits the set of vertices already included in the tree T_i.

3. Form the loop-erasure of this random walk path and let T_{i+1} be the union of T_i with this loop-erased path.

4. Let $T = T_n = \bigcup_{i=1}^{n} T_i$.

Again, this algorithm clearly produces a spanning tree of G. Once again, however, the following theorem is very surprising. It is even surprising that the distribution of the tree T does not depend on the enumeration we chose.

Theorem (Wilson 1996)**.** *The random tree T is a uniform spanning tree of G.*

Wilson's proof [228] uses an ingenious notion of 'cycle popping', of which the above algorithm is merely one implementation. A major advantage of Wilson's algorithm over the Aldous-Broder algorithm is that it readily extends to generate the wired uniform spanning forest of an infinite graph. This led to a surge of progress on the wired uniform spanning forest following its introduction, in particular the landmark work of Benjamini, Lyons, Peres, and Schramm [44].

1.5.3 Uniform spanning forests

We now define the free and wired uniform spanning forests of an infinite graph. These will be defined as weak limits over exhaustions. These weak limits were both implicitly proven to exist by Pemantle [190] in 1991, although the *wired* uniform spanning forest was not considered explicitly until the work of Häggström [109] in 1995.

We write UST_G for the law of the uniform spanning forest of finite graph G. The **free uniform spanning forest** measure FUSF_G is defined to be the weak limit of the sequence $\langle \mathsf{UST}_{G_n} \rangle_{n\geq 1}$, so that

$$\mathsf{FUSF}_G(S \subset \mathfrak{F}) = \lim_{n\to\infty} \mathsf{UST}_{G_n}(S \subset T).$$

for each finite set $S \subset E$, where $\mathfrak{F}$ is a sample of the FUSF of G and T is a sample of the UST of G_n. The **wired uniform spanning forest** measure WUSF_G is defined to be the weak limit of the sequence $\langle \mathsf{UST}_{G_n^*} \rangle_{n\geq 1}$, so that

$$\mathsf{WUSF}_G(S \subset \mathfrak{F}) = \lim_{n\to\infty} \mathsf{UST}_{G_n^*}(S \subset T)$$

for each finite set $S \subset E$, where $\mathfrak{F}$ is a sample of the WUSF of G and T is a sample of the UST of G_n^*.

The existence of these limits follow immediately from the Transfer-Current Theorem, together with what we know about infinite electrical networks, since, for example, we know that currents on G_n converge to free currents on G and hence that the transfer-current matrix on G_n converges to the free transfer current matrix on G (defined in the obvious way). The wired case is similar. In particular, Kirchhoff's effective resistance formula and the transfer-current theorem extend to both the wired and free uniform spanning forests in the natural way. This point of view was not available to Pemantle at the time of his original work, but was available to Benjamini, Lyons, Peres, and Schramm [44], who reproved convergence in this way and deduced the following theorem.

Theorem (Benjamini, Lyons, Peres, and Schramm 2001). *Let G be an infinite graph. Then the free and wired uniform spanning forests of G coincide if and only if G does not admit any non-constant harmonic functions of finite Dirichlet energy.*

In particular, we have that the free and wired forests coincide in any amenable transitive graph and that, by the theorems of Benjamini and Schramm [47, 48] stated in Section 1.3.3, we have that the free and wired uniform spanning forests of a bounded degree planar graph are different if and only if the graph is transient.

1.5.4 The number of trees in the wired forest

Although they are defined as limits of trees, the uniform spanning forests of a graph need not be connected. Indeed, we have the following theorem of Pemantle [190].

Theorem (Pemantle 1991). *The uniform spanning forest of $\mathbb{Z}^d$ is connected if and only if $d \leq 4$.*

Why is dimension four important? A much clearer picture emerged following the introduction of Wilson's algorithm. In their seminal paper [44], Benjamini, Lyons, Peres, and Schramm showed how to extend Wilson's algorithm to generate the wired uniform spanning forest of any transient graph. (It is obvious how to extend both the Aldous-Broder algorithm and Wilson's algorithm to infinite *recurrent* graphs.) This allowed them to prove many things. Their extension, called **Wilson's algorithm rooted at infinity**, can be described as follows. Let G be a transient graph and let $\langle v_i \rangle_{i \geq 1}$ be an enumeration of its vertices. Define a sequence of trees $\mathfrak{F}_i$ in G as follows:

1. Let $\mathfrak{F}_0$ be the empty forest, that has no vertices or edges.
2. Given $\mathfrak{F}_i$, start an independent random walk from v_{i+1} stopped when it hits the set of vertices already included in the tree $\mathfrak{F}_i$. If the walk never hits this set, then it runs forever.
3. Form the loop-erasure of this random walk path and let $\mathfrak{F}_{i+1}$ be the union of $\mathfrak{F}_i$ with this loop-erased path.
4. Let $\mathfrak{F} = \bigcup_{i \geq 1} \mathfrak{F}_i$.

Given the correctness of Wilson's algorithm, the following is fairly easy to verify from the definitions. (Simply consider running Wilson's algorithm on each G_n^* with $v_0 = \partial_n$ and the other vertices appearing in the order specified, and take the limit as $n \to \infty$.)

Theorem (Benjamini, Lyons, Peres, and Schramm 2001)**.** *Let G be a transient graph. Then the random forest $\mathfrak{F}$ generated by Wilson's algorithm rooted at infinity is a wired uniform spanning forest of G.*

Using this algorithm, Benjamini, Lyons, Peres, and Schramm were able to prove the following theorem. We say that a graph G has the **intersection property** if whenever X and Y are independent random walks on G, their traces $\{X_n : n \geq 0\}$ and $\{Y_n : n \geq 0\}$ have non-empty intersection a.s. (or, equivalently, if the traces have infinite intersection a.s.). We say that G has the **non-intersection property** if the traces instead have only finite intersection a.s.

Theorem (Benjamini, Lyons, Peres, and Schramm 2001)**.** *Let G be an infinite graph. Then the wired uniform spanning forest of G is connected a.s. if and only if G has the intersection property. If G has the non-intersection property, then the wired uniform spanning forest of G has infinitely many components almost surely.*

Perhaps this seems that it should be obvious from Wilson's algorithm, but it is not. What *is* obvious is that the WUSF is connected if and only if a random walk almost surely intersects an independent *loop-erased* random walk. The fact that these two properties are equivalent was proved in a companion paper by Lyons, Peres, and Schramm [174], using a clever second moment argument adapted from the work of Fitzsimmons and Salisbury [90, 204].

It can be shown that every transitive graph either has the intersection property or the non-intersection property, yielding the following theorem, which is an analogue of the Newman-Schulman theorem from percolation.

Corollary (Benjamini, Lyons, Peres, and Schramm 2001)**.** *Let G be a transitive graph. Then the wired uniform spanning forest of G is either connected or has infinitely many connected components almost surely.*

Moreover, it can be shown that a transitive graph has the intersection property if and only if

$$\sum_{n \geq 1} n p_n(v, v) = \infty$$

for some (and hence every) vertex v. Once all this is in place, Pemantle's theorem follows from Pólya's estimate (1.2.1). (The fact that $\mathbb{Z}^d$ has the intersection property if and only if $d \leq 4$ was originally proved by Erdös and Taylor in 1960 [88].) More generally, we deduce that the wired uniform spanning forest is disconnected a.s. in any transitive graph satisfying a 5-dimensional isoperimetric inequality.

Greatly extending Pemantle's theorem, Benjamini, Kesten, Peres and Schramm [41] discovered that the transition between connectivity and disconnectivity in dimension four is merely the first of an infinite family of related transitions occurring every four dimensions. In joint work with Yuval Peres [134] (not included in this thesis), we extended this theorem, and developed a detailed picture of how the adjacency structure of the trees in the USF of $\mathbb{Z}^d$ varies as a function of d. In particular, we showed that the adjacency structure of the forest undergoes a qualitative change *every* time the dimension increases and is above four, rather than just every four dimensions.

1.5.5 Geometry of trees in the wired forest

After connectivity, the most basic property of a forest is the number of **ends** its components have. Recall that we defined the space of ends of a tree in Section 1.2.4. Components of the WUSF are one-ended a.s. in many large classes of graphs. Generally speaking, we expect that components of the WUSF will be one-ended a.s. in any transient graph that we have not constructed to serve as a counterexample. The first generally applicable one-endedness theorem was proven using Wilson's algorithm by Benjamini, Lyons, Peres, and Schramm [44].

Theorem (Benjamini, Lyons, Peres, and Schramm 2001)**.** *Let G be a unimodular transitive graph that is not rough isometric to $\mathbb{Z}$. Then every component of the wired uniform spanning forest of G is one-ended almost surely.*

Their proof analysed the recurrent and transient cases separately. The proof in the transient case involved an innovative use of the mass-transport principle, the discovery of which was recounted in the following memorable anecdote of Russ Lyons [1].

> *To me, Oded's most distinctive mathematical talent was his extraordinary clarity of thought, which led to dazzling proofs and results. Technical difficulties did not obscure his vision. Indeed, they often melted away under his gaze. At one point when the four of us [Oded Schramm, Russ Lyons, Itai Benjamini, and Yuval Peres] were working on uniform spanning forests, Oded came up with a brilliant new application of the Mass-Transport Principle. We were not sure it was kosher, and I still recall Yuval asking me if I believed it, saying that it seemed to be "smoke and mirrors." However, when Oded explained it again, the smoke vanished.*

Let us now briefly explain their proof. They begin with a relatively straightforward application of the mass transport principle to deduce that every component has at most two ends. For this part of the proof, they observe that there is a natural (in particular, automorphism invariant) *oriented* version of the WUSF of a transient graph, in which every vertex has exactly one oriented edge emanating from it in the forest. This can be defined by orienting towards the boundary vertex when taking the weak limit, or equivalently by orienting edges of the forest according to the direction they are traversed by the loop-erased random walks used to generate the forest when running Wilson's algorithm. Define the **core** of a forest to be the set of vertices that lie on some simple bi-infinite path in the forest (which is empty if and only if every tree in the forest is finite or one-ended). The core of a two-ended tree is a bi-infinite path which we call the **trunk**. Let $\mathfrak{F}$ be the OWUSF of a transient unimodular transitive graph. Observe that the unique oriented edge in $\mathfrak{F}$ emanating from a vertex in the core of $\mathfrak{F}$ must have its other endpoint in the core also. Define a mass transport

$$\begin{aligned} &f(u,v) \\ &\quad= \mathbb{P}(u \text{ is in the core}, v \text{ is the endpoint of the unique oriented edge emanating from } u \text{ in } \mathfrak{F}). \end{aligned}$$

Then we have that

$$\sum_v f(o,v) = \mathbb{P}(o \text{ is in the core})$$

and

$$\begin{aligned} \sum_v f(v,o) &= \mathbb{E}|\{v \text{ in the core}: \text{ the oriented edge emanating from } v \text{ points into } o\}| \\ &= \mathbb{E}\left[\Big(|\{v \text{ in the core}: v \text{ is adjacent to } o\}| - 1\Big)\,\mathbb{1}(o \text{ is in the core})\right] \end{aligned}$$

Thus, the mass-transport principle implies that

$$\mathbb{E}\left[|\{v \text{ in the core}: v \text{ is adjacent to } o\}| \mid o \text{ is in the core}\right] = 2.$$

If $\mathfrak{F}$ has a multiply-ended component with positive probability, then on this event its core is nonempty, and every vertex in the core has at least two other vertices in the core adjacent to it. If $\mathfrak{F}$ has a component with more than two ends with positive probability, then on this event there exists a vertex in the core with at least three vertices in the core adjacent to it, and it follows from automorphism invariance that the origin is such a vertex with positive probability. The statement about the expectation above implies that this cannot be the case, and so we deduce that every component of the WUSF has at most two ends as claimed.

The authors then rule out the existence of a two ended component, considering separately the cases that the WUSF is connected or disconnected. The case that the WUSF is connected is handled by a reasonably straightforward but still very elegant argument, a generalised version of which appears in [127] (Chapter 4 of this thesis). For the disconnected case, the authors first show that if $\mathfrak{F}$ contains two-ended components, then it is possible to sample $\mathfrak{F}$ conditioned on the origin being contained in the trunk of its component by first sampling the trunk, and then running Wilson's algorithm 'rooted at the trunk' to sample the rest of $\mathfrak{F}$, i.e., beginning the recursion in Wilson's algorithm by setting $\mathfrak{F}^0$ to be the trunk. The fact that this works, which the authors call the **trunk lemma**, is intuitively plausible but unfortunately rather technical to prove.

Once the trunk lemma is in place, we come to Schramm's "smoke and mirrors" mass-transport argument. Suppose for contradiction that $\mathfrak{F}$ contains a two-ended component with positive probability. Since $\mathfrak{F}$ is disconnected, the trunk lemma implies that, conditional on both the trunk of the tree containing the origin, denoted T, and the event that this trunk contains the origin, there is some vertex w of G such that the random walk started at w does not hit T with positive probability. Let W be the set of such vertices. Then it is not hard to see that W must be adjacent to some vertex of T, and automorphsim-invariance implies that the origin is such a vertex with positive probability. it follows that, with positive probability, a random walk started at the origin does not return to T after time zero. In notation,

$$\mathbb{E}\left[\tau_T^+ \mathbb{1}(o \text{ is in the trunk of a two-ended component})\right] = \infty$$

Now, for each vertex v in the trunk of a two-ended component of $\mathfrak{F}$, let B_v be the set of vertices u, including v itself, such that any simple infinite path starting at u in $\mathfrak{F}$ must pass through v, called the **bush** at v. The trunk lemma implies that the probability that u is in B_o conditional on T and the event $o \in T$ is equal to the probability that a random walk started at u hits T, and does so for the first time at o. By time reversal, this is equal to the expected number of times a random walk started at o hits u before returning to T for the first time, and so, summing over u,

$$\begin{aligned}\mathbb{E}&\left[|B_o|\mathbb{1}(o \text{ in the trunk of a two-ended component})\right]\\ &\qquad= \mathbb{E}\left[\tau_T^+ \mathbb{1}(o \text{ in the trunk of a two-ended component})\right] = \infty\end{aligned}$$

On the other hand, if we transport mass from each vertex of u of G in a two-ended component of $\mathfrak{F}$ to the unique vertex v such that $u \in B_v$, we obtain from the mass-transport principle that

$$\mathbb{E}|B_o|\mathbb{1}(o \text{ is in the trunk of a two-ended component}) = \mathbb{P}(o \text{ is in a two-ended component}) < \infty,$$

giving us the desired contradiction.

A more *geometric* understanding of one-endedness in the transient case was developed by Lyons, Morris, and Schramm [170] in 2008. These authors gave an isoperimetric criterion for one-endedness, very similar to the Thomassen criterion for transience. Their proof was based on electrical techniques; see [173] for an updated exposition.

Theorem (Lyons, Morris, and Schramm 2008)**.** *Let G be a graph that satisfies an $f(t)$-isoperimetric inequality for some increasing function $f : (0,\infty) \to (0,\infty)$ for which there exists a constant α such that $f(t) \leq t$ and $f(2t) \leq \alpha f(t)$ for all $t \in (0,\infty)$. If*

$$\int_1^\infty \frac{1}{f(t)^2}\mathrm{d}t < \infty,$$

then every component of the wired uniform spanning forest of G is one-ended almost surely.

Corollary (Lyons, Morris, and Schramm 2008)**.** *Let G be a transitive transient graph. Then every component of the wired uniform spanning forest of G is one-ended almost surely.*

Trees with one end (or finitely many ends) are clearly recurrent by the Nash-Williams Criterion. Recurrence of components in the WUSF holds even more generally than one-endedness, however, as was proven by Morris in 2003 [185] using electrical methods.

Theorem (Morris 2003)**.** *Let G be an infinite graph. Then every component of the wired uniform spanning forest of G is recurrent almost surely.*

1.5.6 The interlacement Aldous-Broder algorithm

Unlike Wilson's algorithm, it was for a long time not apparent how to extend the Aldous-Broder algorithm to generate the wired uniform spanning forest of an infinite transient graph. In our paper [128] (Chapter 5 of this thesis), we showed that such an extension can be done by replacing the random walk in the classical algorithm with the **random interlacement process**. The interlacement Aldous-Broder algorithm turns out, generally speaking, to be better suited than Wilson's algorithm for studying ends in the WUSF, and allowed us to prove several new results. It is also of central importance in upcoming work in which we compute the critical exponents for the USF in high dimensions.

The interlacement process was originally introduced by Sznitman [216] to study the disconnection of cylinders and tori by a random walk trajectory, and was generalised to arbitrary transient graphs by Teixeira [218]. See the monographs [67, 82] for detailed introductions. Roughly speaking, the interlacement process $\mathscr{I}$ on a transient graph G is a 'Poissonian soup' of bi-infinite random walk trajectories in G. As we increase a real time parameter t, more and more of these trajectories appear. Formally, $\mathscr{I}$ is a random subset of $\mathcal{W}/\sim \times \mathbb{R}$, where $\mathcal{W}$ is the space of doubly infinite paths in G that visit each vertex at most finitely often, and $\sim$ holds two such paths to be equivalent if and only if they are reparameterizations of each other.

Theorem (H. 2015)**.** *Let G be a transient graph, let $\mathscr{I}$ be the interlacement process on G, and let $t \in \mathbb{R}$. For each vertex v of G, let $\tau_t(v)$ be the smallest time greater than t such that there exists a trajectory $(W_{\tau_t(v)}, \tau_t(v)) \in \mathscr{I}$ passing through v, and let $e_t(v)$ be the oriented edge of G that is traversed by the trajectory $W_{\tau_t(v)}$ as it enters v for the first time. Then*

$$\mathsf{AB}_t(\mathscr{I}) := \left\{-e_t(v) : v \in V\right\}$$

has the law of the oriented wired uniform spanning forest of G.

We used this algorithm to prove an *anchored* isoperimetric condition for one-endedness, answering positively a question of Lyons, Morris, and Schramm [170].

Theorem (H. 2015)**.** *Let G be a graph that satisfies an anchored $f(t)$-isoperimetric inequality for some increasing function $f : (0,\infty) \to (0,\infty)$ for which there exists a constant α such that $f(t) \leq t$ and $f(2t) \leq \alpha f(t)$ for all $t \in (0,\infty)$. Suppose that f also satisfies each of the following conditions:*

1.
$$\int_1^\infty \frac{1}{f(t)^2}\mathrm{d}t < \infty$$

and

2.
$$\int_1^\infty \exp\left(-\varepsilon\left(\int_s^\infty \frac{1}{f(t)^2}\mathrm{d}t\right)^{-1}\right)\mathrm{d}s < \infty$$

for every $\varepsilon > 0$.

Then every component of the wired uniform spanning forest of G is one-ended almost surely.

On the other hand, we also proved the following, which implies that it is not possible to tell whether or not the components of the WUSF of a graph are one-ended a.s. from the coarse geometry of the graph alone. This answered negatively a question of Lyons, Morris, and Schramm [170].

Theorem (H. 2015)**.** *There exist two bounded degree, rough isometric graphs G and G' such that every component of the wired uniform spanning forest of G is one-ended almost surely, but the wired uniform spanning forest of G' contains a component with uncountably many ends almost surely.*

What can we say about graphs that *do* have multiply ended components in their WUSF? One example of a transient graph in which the WUSF has multiply-ended components is obtained from $\mathbb{Z}^5$ by attaching an infinite path to each of the vertices $u = (0,0,0,0,0)$ and $v = (2,0,0,0,0)$. The WUSF of this graph has the same distribution as the union of the WUSF of $\mathbb{Z}^5$ with each of the two added paths. If the vertices u and v are in the same component of the forest, then this component has exactly three ends, while all other components have one end. If not, then the component containing u and the component containing v both have two ends, while all other components have one end. In particular, the event that the WUSF contains a two-ended component has probability strictly between 0 and 1. Nevertheless, the **number of excessive ends** of $\mathfrak{F}$, that is, the sum over all components of $\mathfrak{F}$ of the number of ends minus 1, is equal to two a.s.

In [128], we used the interlacement Aldous-Broder algorithm to prove that this is a general phenomenon, answering positively a further question of Lyons, Morris, and Schramm [170].

Theorem (H. 2015)**.** *Let G be an infinite graph. Then there exists a constant $c \in \mathbb{N} \cup \{\infty\}$ such that the wired uniform spanning forest of G has c excessive ends almost surely.*

The key to all our results proved using interlacement Aldous-Broder is that it lets us view the WUSF as the stationary measure of the continuous time Markov process $\langle \mathsf{AB}_t(\mathscr{I})\rangle_{t\in\mathbb{R}}$.

In upcoming work, we use the interlacement Aldous-Broder algorithm to compute the critical exponents for the WUSF of $\mathbb{Z}^d$, $d \geq 4$, or more generally of any bounded degree graph satisfying a 4-dimensional isoperimetric inequality. An example of one of the exponent theorems we obtain is the following. The result is new even in the case of $\mathbb{Z}^d$, although related results have been obtained by Barlow and Járai [30] and Bhupatiraju, Hanson, and Járai [52] (who computed the exponent for the *extrinsic* diameter up to a logarithmic correction). Here, the **past** of a vertex v in a sample of the wired uniform spanning forest $\mathfrak{F}$ is defined to be the union of the finite connected components of $\mathfrak{F} \setminus \{v\}$.

Theorem (H. 2017)**.** *Let G be a bounded degree graph satisfying a d-dimensional isoperimetric inequality for some $d > 4$, and let $\mathfrak{F}$ be the wired uniform spanning forest of G. Then there exists a positive constant C such that*

$$C^{-1}R^{-1} \leq \mathbb{P}(\textit{the intrinsic diameter of } \mathrm{past}_{\mathfrak{F}}(v) \textit{ is at least } R) \leq CR^{-1}$$

and

$$C^{-1}R^{-1/2} \leq \mathbb{P}(|\mathrm{past}_{\mathfrak{F}}(v)| \geq R) \leq CR^{-1/2}$$

for every vertex v of G and every $R \geq 1$.

Using these exponents, we also show that every tree of the WUSF has spectral dimension $4/3$ almost surely under the same conditions, meaning that the n-step return probability of a random walk on each of the trees decays like $n^{-2/3+o(1)}$ almost surely.

1.5.7 Indistinguishability of trees and the geometry of trees in the free forest

As we have seen, most of the basic properties of the WUSF are quite well understood. In comparison, our understanding of the FUSF is relatively poor, and several very basic questions remain open. For example, the following very basic property of the FUSF was proven only recently in joint work with Asaf Nachmias [131] (Chapter 6 of this thesis), and independently by Timár [223]. (The corresponding result for the WUSF is an easy consequence of the work of Benjamini, Lyons, Peres, and Schramm [44]: see Lemma 9.5.4 of this thesis.)

Theorem (H. and Nachmias 2015, Timár 2015)**.** *Let G be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$ and let $\mathfrak{F}$ be a sample of the free uniform spanning forest of G. Then $\mathfrak{F}$ is either connected or has infinitely many components almost surely.*

The result remains open in the nonunimodular transitive case, where it is also expected to hold. Our proof was based on **update-tolerance**, a property of the USFs that we introduced in the earlier work [127]. This property allows us to insert an edge of our choice into the uniform spanning forest, so long as we also delete some edge that depends on the forest and the edge that we chose to insert, in such a way that the resulting forest is absolutely continuous with respect to the uniform spanning forest that we started with. This property plays the role for USFs that insertion tolerance played for percolation. However, because we are forced to insert and delete edges at the same time, update-tolerance is a much weaker property than insertion-tolerance. In particular, we cannot simply add an edge to merge two trees into one, so that we cannot simply adapt the Newman-Schulman proof from percolation to show that there is either one or infinitely many components.

Instead, we first showed that every component has a well-defined **frequency**, meaning that if X is a random walk on G independent of the forest $\mathfrak{F}$, then for each tree T of $\mathfrak{F}$ the limit

$$\mathrm{Freq}(T) := \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}(X_n \in T)$$

exists and is almost surely equal to some constant depending on T but not on the walk X. If there are finitely many components, their frequencies must add to 1, and ergodicity implies that the set of frequencies appearing in the forest is non-random. We then argue that, by performing updates in an appropriate way, we can change the set of frequencies that appear, giving us a contradiction. (It is here that unimodularity is used in a crucial way to show that *a tree with positive frequency cannot have a zero frequency branch.*)

The update-tolerance method allowed us to solve a conjecture of Benjamini, Lyons, Peres, and Schramm [44], which stated that both the WUSF and FUSF satisfy a Lyons-Schramm type indistinguishability theorem.

Theorem (H. and Nachmias 2015)**.** *Let G be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$ and let $\mathfrak{F}$ be a sample of either the free uniform spanning forest or the wired uniform spanning forest of G. Then for each automorphism-invariant Borel-measurable set $\mathscr{A}$ of subgraphs of G, either every connected component of $\mathfrak{F}$ is in $\mathscr{A}$ or every connected component of $\mathfrak{F}$ is not in $\mathscr{A}$ almost surely.*

Partial progress on the conjecture was also made in the independent work of Timár [223], who proved that components of the FUSF are indistinguishable from each other when G is a unimodular transitive graph such that the FUSF and WUSF are different. Our proof in that case was handled by a similar approach to his, whereas the other case (of the WUSF or the FUSF when it is equal to the WUSF) was handled by a completely different method.

Along the way, we also answered a further question of Benjamini, Lyons, Peres, and Schramm in the unimodular case. Those authors had proved that if the WUSF and FUSF of a transitive graph are different, then, in contrast to the WUSF, the FUSF has at least one component that is transient and infinitely ended. We proved that in fact, if this occurs in the unimodular case, then *every* component of the FUSF is transient and infinitely-ended.

Theorem (H. and Nachmias 2015, Timár 2015)**.** *Let G be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$ and let $\mathfrak{F}$ be a sample of FUSF_G. If the measures FUSF_G and WUSF_G are distinct, then every component of $\mathfrak{F}$ is transient and has infinitely many ends almost surely.*

Although this would follow as a special case of the indistinguishability theorem above, we in fact had to prove it separately as part of the proof of the general theorem. The result is also expected to hold in the nonunimodular case.

1.5.8 Uniform spanning forests of planar graphs

Unlike for the WUSF, there is no simple criterion for the connectivity of the FUSF. Indeed, it is not even known whether the number of components of the FUSF is nonrandom for any given fixed graph. For *planar* graphs, the situation is much better due to the following duality between the FUSF and WUSF.

Proposition (Benjamini, Lyons, Peres, and Schramm 2001)**.** *Let M be a simply connected planar map with locally finite dual, and let $\mathfrak{F}$ be the free uniform spanning forest of M. Then $\mathfrak{F}^\dagger$ is distributed as the wired uniform spanning forest of $M^\dagger$. In particular, $\mathfrak{F}$ is connected almost surely if and only if every component of $\mathfrak{F}^\dagger$ is one-ended almost surely.*

Using this duality, it follows easily from the theorems we have discussed previously that, for example, the FUSF of any transitive, simply-connected planar map with locally finite dual is connected almost surely. Benjamini, Lyons, Peres, and Schramm [44] asked whether the FUSF of *any* bounded degree, simply connected map is connected. In joint work with Nachmias [130] (Chapter 10 of this thesis), we answered this question positively.

Theorem. *The free uniform spanning forest is almost surely connected in any bounded degree, simply-connected planar map.*

In light of the above duality, this theorem is equivalent in the case of a locally finite dual to the following dual theorem. The general case is then easily deduced from the locally finite dual case via an approximation argument.

Theorem. *Every component of the wired uniform spanning forest is one-ended almost surely in any bounded codegree, simply-connected planar map.*

Our proof uses circle packing in an essential way. Suppose T is a simple, simply-connected triangulation, circle packed in either the plane or the disc, and let e be a fixed edge of T. Let $\mathfrak{F}$ be the wired uniform spanning forest of T. We show that if the components of e^- and e^+ in $\mathfrak{F} \setminus \{e\}$ both have large diameter as measured in either the Euclidean or hyperbolic metric as appropriate on the circle packing, then with high probability e is not in $\mathfrak{F}$. This is done by exploring $\mathfrak{F}$ in a Markovian way, analyzing the circle packing to show that certain effective resistances are small, and then using Kirchoff's effective resistance formula.

In the bounded degree case, the proof can be made quantitative, and led us to a computation of the *critical exponents* of the uniform spanning forest for transient, simply-connected, polyhedral planar maps with bounded degrees and codegrees. Surprisingly, we found that, if one measures distances and areas using the hyperbolic geometry of the circle packings rather than the usual combinatorial geometry, the exponents are universal over the entire class of graphs. For example, the probability that the past of a vertex has diameter at least r in the hyperbolic metric always scales like $1/r$, even for triangulations with rather unusual circle packings such as those in the figure above.

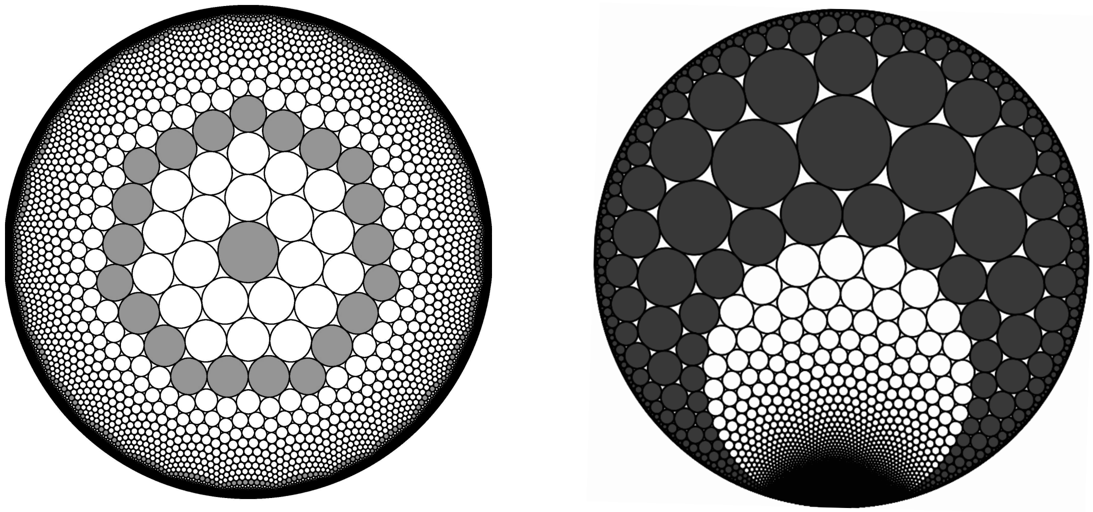

Figure 1.8: Two bounded degree, simple, proper plane triangulations for which the graph distance is not comparable to the hyperbolic distance. Similar examples are given in [215, Figure 17.7]. Left: In this example, rings of degree seven vertices (grey) are separated by growing bands of degree six vertices (white), causing the hyperbolic radii of circles to decay. The bands of degree six vertices can grow surprisingly quickly without the triangulation becoming recurrent [211]. Right: In this example, half-spaces of the 8-regular (grey) and 6-regular (white) triangulations have been glued together along their boundaries; the circles corresponding to the 6-regular half-space are contained inside a horodisc and have decaying hyperbolic radii.

1.5.9 USFs of multiply-connected planar maps

It is natural to ask whether our result concerning the connectivity of the FUSF extends to *all* bounded degree planar maps, without the assumption that they are simply connected. Indeed, if such an extension were true, it would have interesting consequences for finite planar graphs. In [130] (Chapter 10) we present an example, which was analyzed in collaboration with Gady Kozma, to show that the theorem does *not* admit such an extension.

Theorem. *There exists a bounded degree planar graph G such that the free uniform spanning forest of G is disconnected with positive probability.*

In light of this example, it is an interesting problem to determine for which bounded degree planar graphs the FUSF is connected. As the final result of [130], we initiated progress on this question by showing that the FUSF is connected almost surely on any bounded degree, *countably-connected* planar map. Here, countably connected means that the surface associated to the map is homeomorphic to a domain in $\mathbb{C} \cup \{\infty\}$ whose complement has at most countably many connected components.

Theorem. *Let M be a countably-connected, bounded degree planar map. Then the free uniform spanning forest of M is connected almost surely.*

The proof of this theorem requires two additional ingredients besides those appearing in the simply-connected case. The first is the introduction of the *transboundary uniform spanning forest*

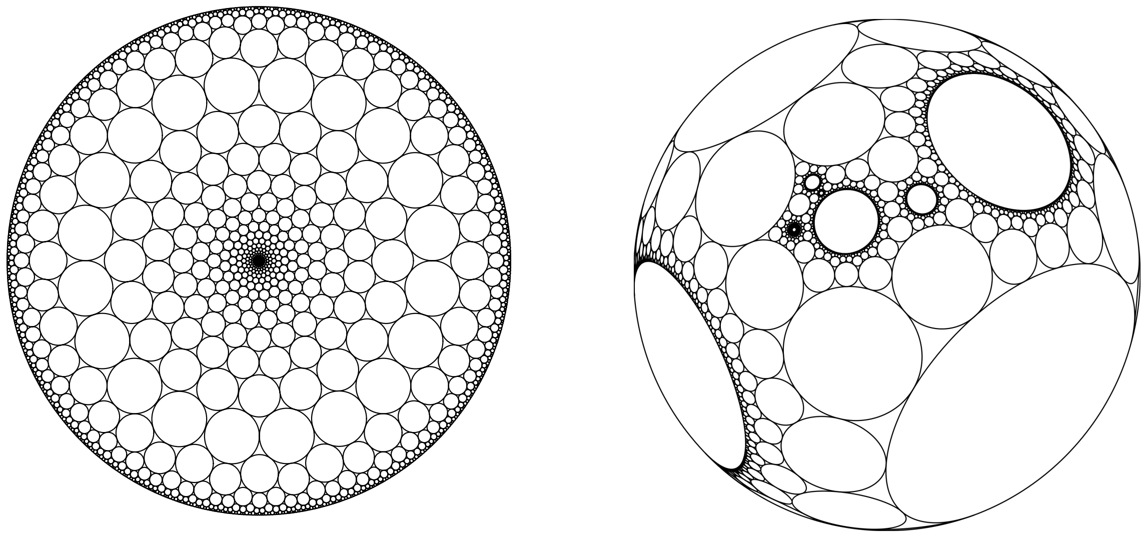

Figure 1.9: Left: A circle packing in the multiply-connected circle domain $\mathbb{D} \setminus \{0\}$. Right: A circle packing in a circle domain domain with several boundary components.

(TUSF), which is dual to the FUSF of a (not necessarily proper) plane graph in the same way that the WUSF is dual to the FUSF of a proper plane graph. The second additional ingredient, which we use to prove this dual statement, is a transfinite induction inspired by He and Schramm's work on the Koebe Conjecture [122, 207].

1.6 Unimodular random graphs

A **rooted graph** (G,ρ) is a connected, locally finite graph $G=(V,E)$ together with a distinguished vertex ρ, called the **root**. A graph isomorphism $\phi : G \to G'$ is an isomorphism of rooted graphs if it maps the root to the root. The **local topology** (see [50]) is the topology on the set $\mathcal{G}_\bullet$ of isomorphism classes of rooted graphs is defined so that two rooted graphs are close if they have large isomorphic balls around their roots. Formally, it is the topology induced by the metric

$$d_{\mathrm{loc}}\left((G,\rho),(G',\rho')\right) = e^{-R},$$

where

$$R = R\left((G,\rho),(G',\rho')\right) = \sup\left\{R \geq 0 : B_R(G,\rho) \cong B_R(G',\rho')\right\},$$

i.e., the maximal radius such that the balls $B_R(G,\rho)$ and $B_R(G',\rho')$ are isomorphic as rooted graphs. A **random rooted graph** is a random variable taking values in the space $\mathcal{G}_\bullet$ endowed with the local topology. Similarly, a **doubly-rooted graph** is a graph together with an ordered pair of distinguished (not necessarily distinct) vertices. Denote the space of isomorphism classes of doubly-rooted graphs equipped with this topology by $\mathcal{G}_{\bullet\bullet}$.

A **mass-transport** is a Borel function $f : \mathcal{G}_{\bullet\bullet} \to [0,\infty]$. A random rooted graph (G,ρ) is said to be **unimodular** if it satisfies the **Mass Transport Principle**: for every mass transport f,

$$\mathbb{E}\left[\sum_{v\in V} f(G,\rho,v)\right] = \mathbb{E}\left[\sum_{u \in V} f(G,u,\rho)\right]. \tag{MTP}$$

Unimodularity of random rooted graphs was first defined by Benjamini and Schramm in their seminal work [50], and was developed thoroughly by Aldous and Lyons [7], who showed that much of what was known about, for example, percolation and uniform spanning forests on unimodular transitive graphs generalised to the unimodular random rooted graph setting, often under the additional assumption of bounded degree or of finite expected degree. Moreover, restricting oneself to use only unimodularity rather than more delicate properties of a specific example often leads to proofs that are not only more general, but also much more simple and elegant. Our work with Yuval Peres on the collision property [133] (Chapter 3) is a good example of this.

It is also possible to define, with similar definitions, the spaces of rooted and doubly-rooted maps $\mathcal{M}_\bullet$ and $\mathcal{M}_{\bullet\bullet}$, and the corresponding notions of unimodularity and the Mass-Transport Principle. Similarly, we have unimodular random rooted **marked** graphs and maps, in which every vertex and/or edge has a random *mark* taking values in some Polish space. For example, given a unimodular transitive graph G and some automorphism-invariant percolation ω on G (such as Bernoulli bond percolation, the FUSF, or the WUSF), we have that (G,ρ,ω) is a unimodular random rooted marked graph, with edge marks in $\{0,1\}$.

Every unimodular transitive graph can be made into a unimodular random rooted graph by choosing a root vertex arbitrarily. Furthermore, if G is a (possibly random) finite graph and ρ

is chosen uniformly from the vertex set of G, then the random rooted graph (G,ρ) is unimodular. This leads to many further examples: It can be shown that the set of distributions of the unimodular random rooted graphs is closed under the topology of weak convergence, and hence that if (G_n,ρ_n) are a sequence of finite random graphs with uniform random roots converging to some infinite random rooted graph (G,ρ) in distribution, then the limit (G,ρ) is itself unimodular. This procedure is known as a **Benjamini-Schramm limit**. A unimodular random rooted graph that can be obtained as a Benjamini-Schramm limit of finite graphs is said to be **sofic**. It is a major open problem, with important ramifications in group theory, to determine whether *every* unimodular random rooted graph is sofic.

Conjecture (Aldous and Lyons 2007)**.** *Every unimodular random rooted graph is sofic.*

It is the widespread belief, certainly among group theorists, that the conjecture is false. We choose to withhold judgment.

1.6.1 Reversibility

Unimodularity is closely related to *reversibility.* A random rooted graph (G,ρ) with at least two vertices is said to be **stationary** if, when $\langle X_n\rangle_{n\geq 0}$ is a simple random walk on G started at the root,

$$(G,\rho) \overset{d}{=} (G,X_n)$$

for all n and is said to be **reversible** if

$$(G,\rho,X_n) \overset{d}{=} (G,X_n,\rho)$$

for all n. Note that this is not the same as the reversibility of the random walk on G, which holds for any graph[6]. Every reversible random rooted graph is clearly stationary, but the converse need not hold in general: transitive graphs that are not unimodular are an important example.

Aldous and Lyons proved the following correspondence between reversible random rooted graphs and unimodular random rooted graphs of finite expected degree: If (G,ρ) is a unimodular random rooted network with $\mathbb{E}[\deg\rho]<\infty$ (that is almost surely not equal to a single degree zero vertex), then biasing the law of (G,ρ) by $\deg\rho$ (that is, re-weighting the law of (G,ρ) by the Radon-Nikodym derivative $\deg\rho/\mathbb{E}[\deg\rho]$) yields the law of a reversible random rooted network. Conversely, if (G,ρ) is a reversible random rooted network, then biasing the law of (G,ρ) by $\deg^{-1}\rho$ yields the law of a unimodular random rooted network. This is summarized by the following bijection:

$$\big\{(G,\rho) \text{ unimodular with } \mathbb{E}[\deg\rho]<\infty,\ \deg(\rho)\geq 1 \text{ a.s.}\big\} \underset{\text{bias by } \deg^{-1}\rho}{\overset{\text{bias by } \deg\rho}{\rightleftarrows}} \big\{(G,\rho) \text{ reversible}\big\}.$$

[6]Rather, a random rooted graph (G,ρ) is reversible if and only if the $\mathcal{G}_{\bullet\bullet}$-valued Markov process $\langle (G,X_n,X_{n+1})\rangle_{n\geq 0}$ is reversible.

1.6.2 Invariant amenability and nonamenability

Given a unimodular random rooted graph (G, ρ) and a random subgraph ω of G, we say that ω is a **percolation** on (G, ρ) if (G, ρ, ω) is unimodular when considered as a random rooted marked graph, with edge marks in $\{0, 1\}$. A primary example is given by (G, ρ) a unimodular transitive graph and ω an automorphism-invariant random subgraph of G. The following property of nonamenable unimodular transitive graphs was discovered for regular trees by Häggström [111], and was extended to all unimodular transitive graphs by Benjamini, Lyons, Peres, and Schramm [36].

Theorem. *Let G be a nonamenable unimodular transitive graph. Then there exists $\alpha > 0$ such that every automorphism-invariant percolation ω on G with $\mathbb{E} \deg_\omega(\rho) \geq \deg \rho - \alpha$ has an infinite connected component almost surely.*

This theorem gives a great deal of insight about percolation on nonamenable proofs, and is proven very easily. We call an automorphism-invariant percolation ω on G **finitary** if all its clusters are finite almost surely. We write $K_\omega(v)$ for the connected component of ω containing v. Suppose that ω is a finitary percolation on G, and define a mass-transport

$$F(G, u, v, \omega) = \frac{\deg u - \deg_\omega u}{|K_\omega(u)|} \mathbb{1} \big(v \in K_\omega(u)\big)$$

Then

$$\mathbb{E} \sum_v F(G, \rho, v, \omega) = \mathbb{E} \left[\deg \rho - \deg_\omega \rho\right]$$

and

$$\mathbb{E} \sum_u F(G, u, \rho, \omega) = \mathbb{E} \left[\frac{|\partial_E K_\omega(\rho)|}{|K_\omega(\rho)|}\right],$$

so that the Mass-Transport Principle yields that

$$\mathbb{E} \left[\deg \rho - \deg_\omega \rho\right] = \mathbb{E} \left[\frac{|\partial_E K_\omega(\rho)|}{|K_\omega(\rho)|}\right].$$

Since G is nonamenable, the right hand side is bounded by a positive constant, and the result follows.

Aldous and Lyons [7] observed that Häggström's theorem can be used as a *definition* for a notion of nonamenability for unimodular random graphs, which we call *invariant nonamenability.* If (G, ρ) is a unimodular random rooted graph and ω is a random subgraph of G, we say that ω is a **percolation** on G if (G, ρ, ω) is a unimodular random rooted marked graph. We say that a unimodular random rooted graph (G, ρ) is **invariantly nonamenable** if there exists a positive constant ε such that

$$\mathbb{E} \left[\frac{|\partial_E K_\omega(\rho)|}{|K_\omega(\rho)|}\right] > \varepsilon$$

for every finitary percolation ω. By the mass-transport argument above, if $\mathbb{E} \deg \rho < \infty$ this is equivalent to there existing a positive constant α such that

$$\mathbb{E}\deg_\omega \rho \leq \mathbb{E}\deg \rho - \alpha$$

for every finitary percolation ω. In other words, (G, ρ) is invariantly nonamenable if the conclusion of Häggström's theorem holds for it. If (G, ρ) is not invariantly nonamenable we say that it is **invariantly amenable**.

Invariant nonamenability turns out to be the right notion of nonamenability for unimodular random graphs in many ways. For example, it is much easier to prove that supercritical percolation clusters on a transitive unimodular random graph are invariantly nonamenable than to prove they have anchored expansion. Moreover, it is often easier to use invariant nonamenability to deduce properties of the graph than it is to use anchored expansion. The following theorem, originally stated for percolations on unimodular random graphs by Benjamini, Lyons, and Schramm [45] and observed to generalize to unimodular random rooted graphs by Aldous and Lyons [7], allows us to easily deduce many of the properties of invariantly nonamenable graphs from their classically nonamenable analogues. It plays for invariant nonamenability a role analogous to that Virag's 'ocean and islands' construction [227] plays in the context of anchored expansion.

Theorem (Benjamini, Lyons, and Schramm 1999)**.** *Let (G, ρ) be an invariantly nonamenable unimodular random rooted graph. Then there exists a percolation ω on (G, ρ) such that $\deg(v)$ is bounded on ω, ω is (classically) nonamenable, and ω is almost surely a forest.*

This theorem easily implies, for example, that the random walk on a bounded degree, invariantly nonamenable random graph has positive speed.

1.6.3 Unimodular random planar maps

Benjamini and Schramm were motivated to introduce the notion of unimodular random graphs by their work on *random planar maps.* They proved the following remarkable theorem.

Theorem (Benjamini and Schramm 2001)**.** *Let (T, ρ) be an infinite unimodular random rooted simple triangulation that is a Benjamini-Schramm limit of finite simple triangulations of the plane. Then T can be circle packed in either the plane $\mathbb{C}$ or the punctured plane $\mathbb{C} \setminus \{0\}$ almost surely.*

In light of the work of He and Schramm [121] (who also prove that bounded degree circle packings in $\mathbb{C} \setminus \{0\}$ are recurrent), the following is an immediate corollary of this theorem for simple triangulations. The general case follows by a straightforward reduction.

Corollary (Benjamini and Schramm 2001)**.** *Let (M, ρ) be a bounded degree unimodular random rooted planar map that is a Benjamini-Schramm limit of finite planar maps. Then M is recurrent almost surely.*

In joint work with Angel, Nachmias, and Ray [21] (Chapter 7), we established a stronger version of this theorem via a much simpler proof, based on a simple mass-transport argument. Our theorem should be compared with the results of Beardon and Stephenson [35] and He and Schramm that we mentioned earlier, which also pertained to the connection between the circle packing type and the average degree in some sense.

Theorem. *Let (G, ρ) be an infinite, simply-connected, ergodic unimodular random rooted planar triangulation. Then either*

> $\mathbb{E}[\deg(\rho)] = 6$*, in which case (G, ρ) is invariantly amenable and almost surely CP parabolic,*

or else

> $\mathbb{E}[\deg(\rho)] > 6$*, in which case (G, ρ) is invariantly nonamenable and almost surely CP hyperbolic.*

Benjamini-Schramm limits of finite planar maps are easily seen to have average degree at most six by Euler's formula. The case of the Benjamini-Schramm theorem in which the limit is not simply-connected can be reduced to the simply-connected case by taking universal covers.

The hypothesis that (M, ρ) is **ergodic** means that its law is an extreme point of the set of laws of unimodular random rooted maps. Equivalently, it means that for any event $\mathscr{A} \subseteq \mathcal{M}_\bullet$ such that does not depend on the choice of root, (M, ρ) has probability either zero or one to be in $\mathscr{A}$. Ergodicity rules out, for example, taking the unimodular random rooted triangulation that is the 6-regular triangular lattice with probability 1/2 and the 7-regular hyperbolic triangulation with probability 1/2.

Random maps play an important role in the theory of *Liouville quantum gravity*, which is far too big a subject to go into on seriously here. The reader is referred to the survey [98] for a summary of the state of that field as it stood in 2013, although they should note that substantial progress has been made since then. A central player in this theory is the *uniform infinite planar triangulation* of Angel and Schramm [24].

Theorem (Angel and Schramm 2003)**.** *Let T_n be chosen uniformly from the set of triangulations of the sphere with n vertices, and let ρ_n be a uniformly chosen random vertex of T_n. Then there exists a simply-connected random rooted triangulation (T, ρ) such that*

$$(T_n, \rho_n) \xrightarrow[n\to\infty]{d} (T, \rho).$$

The random rooted triangulation (T, ρ) is known as the **uniform infinite planar triangulation** (UIPT). A similar result for quadrangulations (i.e., maps in which every face has degree four) is due to Krikun [157].

Since the UIPT does *not* have bounded degrees, its recurrence does not follow from the He-Schramm theorem. Indeed, it remained an open problem to prove that it was recurrent for some time until it was finally solved by Gurel-Gurevich and Nachmias in 2013 [107], using a strengthened, quantitative version of the original Benjamini-Schramm proof from the bounded degree case.

Theorem (Gurel-Gurevich and Nachmias 2013)**.** *Let (M, ρ) be a unimodular random rooted planar map that is a Benjamini-Schramm limit of finite planar maps, and suppose that there exists a positive constant c such that*

$$\mathbb{P}(\deg \rho \geq n) \leq \exp(-cn)$$

for all $n \geq 1$. Then M is recurrent almost surely.

Corollary (Gurel-Gurevich and Nachmias 2013)**.** *The UIPT and UIPQ are recurrent almost surely.*

A completely different new proof of this theorem, not using circle packing, has appeared in a recent work of Lee [164].

A distinguishing feature of the UIPT is that it enjoys a certain form of *spatial Markov property.* In [75], following similar work on half-planar models by Angel and Ray [23], Curien showed that set of all random rooted triangulations with this property form a one parameter family (T_κ, ρ), with $\kappa \in (0, 2/27]$. Each of these random triangulations is unimodular, and the UIPT is given by the extremal value $\kappa = 2/27$. Curien also showed that the other ($\kappa < 2/27$) triangulations in the family are *hyperbolic* in various senses. These triangulations were a major motivation for our work on unimodular random planar maps.

1.6.4 The dichotomy theorem

What about maps that are not triangulations? There, the appropriate quantity to look at is not the average degree but the **average curvature**. Recall that the internal angles of a regular k-gon are given by $(k-2)\pi/k$. We define the **angle sum** at a vertex v of a map M to be

$$\theta(v) = \theta_M(v) = \sum_{f \perp v} \frac{\deg(f) - 2}{\deg(f)} \pi,$$

where we write $f \perp v$ if the vertex v is incident to the face f, and interpret the sum with multiplicities if v is multiply incident to f. This definition extends to maps with infinite faces, with the convention that $(\infty - 2)/\infty = 1$. In the case that every face of M has degree at least 3, we interpret $\theta(v)$ as the total angle of the corners at v if we form M by gluing together regular polygons, where we consider the upper half-space $\{x + iy \in \mathbb{C} : y > 0\}$ with edges $\{[n, n+1] : n \in \mathbb{Z}\}$ to be a regular ∞-gon. Of course, M cannot necessarily be drawn in the plane with regular polygons, and the angle sum at a vertex of M need not be 2π. We define the **curvature** of M at the vertex v to be the angle sum deficit

$$\kappa(v) = \kappa_M(v) = 2\pi - \theta(v)$$

and define the **average curvature** of a unimodular random rooted map (M,ρ), denoted $\mathbb{K}(M,\rho)$, to be the expected curvature at the root

$$\mathbb{K}(M,\rho) = \mathbb{E}[\kappa(\rho)] = 2\pi - \mathbb{E}\left[\sum_{f \perp \rho} \frac{\deg(f)-2}{\deg(f)}\pi\right].$$

Note that if $\mathbb{E}[\deg(\rho)]$ is finite then $\mathbb{K}(M,\rho)$ is also finite.

In joint work with Angel, Nachmias, and Ray [20] (Chapter 9), we used this notion to greatly expand and generalise the dichotomy from the previous subsection, showing that many diverse properties of a unimodular random planar map are determined by the average curvature. The theorem complements, but does not exactly parallel, the dichotomy for deterministic bounded degree planar maps that we saw earlier.

Theorem (Angel, H., Nachmias, and Ray 2016: The Dichotomy Theorem)**.** *Let (M,ρ) be an infinite, ergodic, unimodular random rooted planar map and suppose that $\mathbb{E}[\deg(\rho)] < \infty$. Then the average curvature of (M,ρ) is non-positive and the following are equivalent:*

1. *(M,ρ) has average curvature zero.*
2. *(M,ρ) is invariantly amenable.*
3. *Every bounded degree subgraph of M is amenable almost surely.*
4. *Every subtree of M is amenable almost surely.*
5. *Every bounded degree subgraph of M is recurrent almost surely.*
6. *Every subtree of M is recurrent almost surely.*
7. *(M,ρ) is a Benjamini-Schramm limit of finite planar maps.*
8. *(M,ρ) is a Benjamini-Schramm limit of a sequence $\langle M_n \rangle_{n\geq 0}$ of finite maps such that*

$$\frac{\mathit{genus}(M_n)}{\#\{\mathit{vertices\ of\ } M_n\}} \xrightarrow[n\to\infty]{} 0.$$

9. *The Riemann surface associated to M is conformally equivalent to either the plane $\mathbb{C}$ or the cylinder $\mathbb{C}/\mathbb{Z}$ almost surely.*
10. *M does not admit any non-constant bounded harmonic functions almost surely.*
11. *M does not admit any non-constant harmonic functions of finite Dirichlet energy almost surely.*
12. *The laws of the free and wired uniform spanning forests of M coincide almost surely.*
13. *The wired uniform spanning forest of M is connected almost surely.*
14. *Two independent random walks on M intersect infinitely often almost surely.*
15. *The laws of the free and wired minimal spanning forests of M coincide almost surely.*
16. *Bernoulli(p) bond percolation on M has at most one infinite connected component for every $p \in [0,1]$ almost surely (in particular, $p_c = p_u$).*
17. *M is vertex extremal length parabolic almost surely.*

In light of this theorem, we call a unimodular random rooted map (M,ρ) with $\mathbb{E}[\deg(\rho)] < \infty$

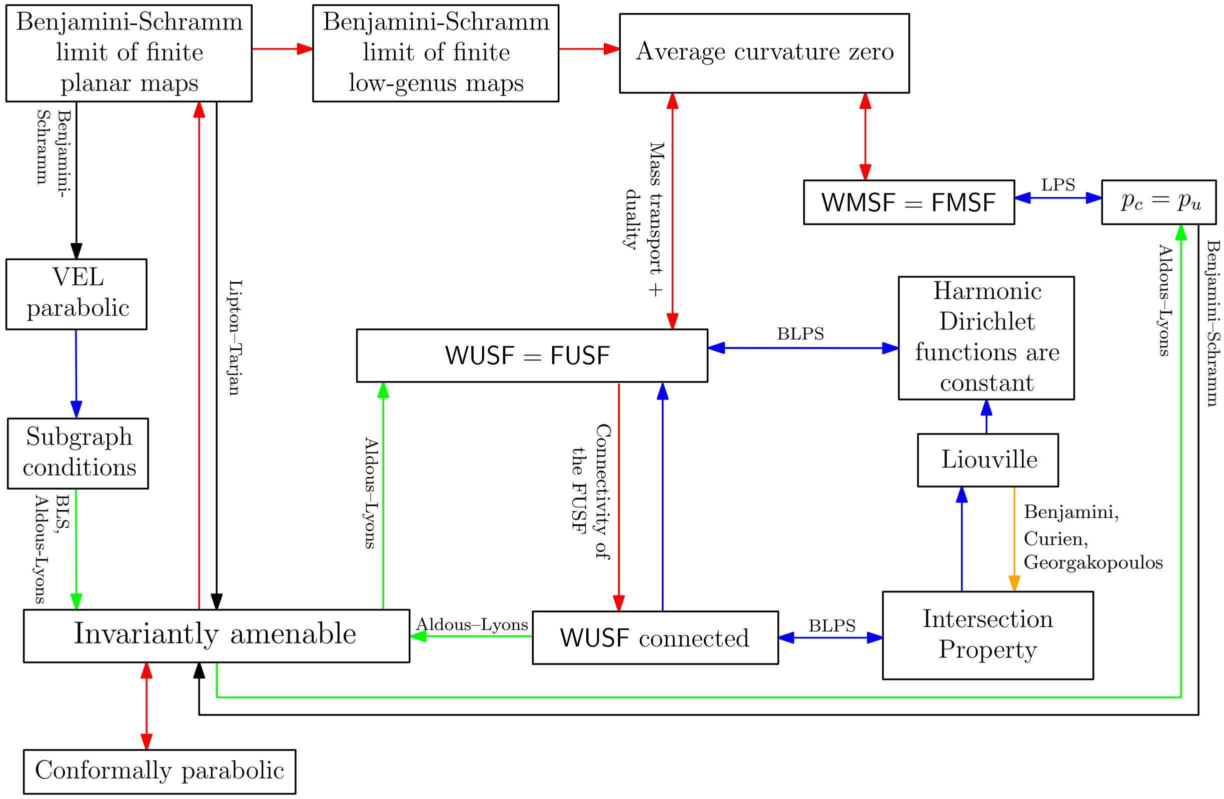

Figure 1.10: The logical structure for the proof of Theorem 9.1.1 in the simply connected case. Implications new to this paper are in red. Blue implications hold for arbitrary graphs; the orange implication holds for arbitrary planar graphs, and green implications hold for unimodular random rooted graphs even without planarity. A few implications between items that are known but not used in the proof are omitted.

parabolic if its average curvature is zero (and, in the planar case, clauses (1)–(17) all hold), and **hyperbolic** if its average curvature is negative (and, in the planar case, the clauses all fail).

A central connection between the average curvature and the behaviour of random processes on the map is given by the following formula for the expected degree of the FUSF. It is known that the expected degree at the root of the WUSF of a unimodular random rooted graph is always equal to 2, and since the FUSF stochastically dominates the WUSF it follows that they are equal if and only the expected degree of the FUSF at the root is 2 also.

Theorem (Angel, H., Nachmias, and Ray 2016)**.** *Let (M, ρ) be an infinite, simply connected unimodular random map, and suppose that $\mathbb{E}[\deg(\rho)] < \infty$, and let $\mathfrak{F}$ be the free uniform spanning forest of M. Then*

$$\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] = \frac{1}{\pi}\mathbb{E}[\theta(\rho)] = 2 - \frac{1}{\pi}\mathbb{K}(M, \rho). \tag{1.6.1}$$

In particular, the FUSF and WUSF of M coincide if and only if (M, ρ) has average curvature zero.

Secondly, we have the following theorem, which is also a major piece of the Dichotomy Theorem above, and complements our results with Nachmias that held for deterministic bounded degree planar graphs [130].

Theorem (Angel, H., Nachmias, and Ray 2016)**.** *Let (M, ρ) be a simply connected unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$. Then the free uniform spanning forest of M is connected almost surely.*

The case in which the dual of M is *recurrent* and locally finite followed easily from previous results. In the case that the dual of M is *transient* and locally finite, this theorem follows from our earlier work [127, 128], the main purpose of which was to remove the bounded degree hypothesis which was required by the earlier result of Aldous and Lyons [7].

Theorem (H. 2015)**.** *Let (G, ρ) be a transient unimodular random rooted graph, and let $\mathfrak{F}$ be the wired uniform spanning forest of G. Then every component of $\mathfrak{F}$ is one-ended almost surely.*

The case that the dual of $M^\dagger$ is *not* locally finite was handled by a separate argument based on a variation of Wilson's algorithm.

Using the fact that the FUSF is connected almost surely, we were able to deduce the following theorem, which verifies the Aldous-Lyons conjecture in the case of simply-connected planar maps.

Theorem (Angel, H., Nachmias, and Ray 2016)**.** *Every simply-connected unimodular random rooted map is sofic, that is, a Benjamini-Schramm limit of finite maps.*

Note that, by the Dichotomy Theorem, the sequence of approximating maps are necessarily non-planar in the hyperbolic case.

1.6.5 Boundary theory of unimodular random triangulations

In our joint work with Angel, Nachmias, and Ray [21], we also developed the *boundary theory* of circle packings of unimodular random triangulations. When P is a circle packing of a graph $G = (V, E)$ in the unit disc, we define $z_h(v)$ to be the hyperbolic centre of the circle of P corresponding to the vertex v.

Theorem. *Let (G, ρ) be a simple, one-ended, CP hyperbolic unimodular random planar triangulation with $\mathbb{E}[\deg^2(\rho)] < \infty$. Let P be a circle packing of G in the unit disc, and let $\langle X_n \rangle_{n \geq 0}$ be a simple random walk on G. The following hold conditional on (G, ρ) almost surely:*

1. *$z(X_n)$ and $z_h(X_n)$ both converge to a (random) point denoted $\Xi \in \partial\mathbb{D}$,*
2. *The law of Ξ has full support $\partial\mathbb{D}$ and no atoms.*
3. *$\partial\mathbb{D}$ is a realisation of the Poisson boundary of G. That is, for every bounded harmonic function h on G there exists a bounded measurable function $g : \partial\mathbb{D} \to \mathbb{R}$ such that*

$$h(v) = \mathbf{E}_v[g(\Xi)].$$

Although the results are similar to those proved in the bounded degree case by Benjamini and Schramm [47] and by Angel, Barlow, Gurel-Gurevich, and Nachmias [17], the proofs are very different. Indeed, the analytic methods of [17, 47] fail badly in the absence of the bounded degree assumption, and were unavailable to us. This required us to develop an entirely different, more probabilistic, approach to the boundary theory of planar graphs. (Georgakopoulos's square tiling analysis [100] is less reliant on bounded degrees, but the assumption is still crucial to Benjamini and Schramm's proof that the random walk converges to a point in the boundary of the square tiling.) In joint work with Peres [132], we later adapted these methods to give simpler proofs in the bounded degree case, as we mentioned earlier. It is worth noting, however, that the analytic methods give stronger results when applicable: for instance, we do not know that the boundary of the unit disc is a realization of the *Martin boundary* for circle packings of unimodular random hyperbolic triangulations.

Finally, our methods also yielded results concerning the speed of the random walk as measured in the hyperbolic metric.

Theorem. *Let (G,ρ) be a simple, one-ended, CP hyperbolic unimodular random rooted planar triangulation with $\mathbb{E}[\deg^2(\rho)] < \infty$ and let P be a circle packing of G in the unit disc. Then almost surely*

$$\lim_{n\to\infty} \frac{d_{hyp}(z_h(\rho), z_h(X_n))}{n} = \lim_{n\to\infty} \frac{-\log r(X_n)}{n} > 0.$$

In particular, both limits exist. Moreover, the limits do not depend on the choice of packing, and if (G,ρ) is ergodic then this limit is an almost sure constant.

Unlike the previous theorem, this result does not have an analogue in the deterministic bounded degree case.

1.7 Open problems

We conclude the introduction with a collection of our favourite open problems that appear in the rest of the thesis. Many other interesting open problems on similar topics, several of which we have already mentioned, appear in [173].

Question. *Let G be a bounded degree proper plane graph.*

1. *Let H be a finite graph. Is the free uniform spanning forest of the product graph $G \times H$ connected almost surely?*

2. *Let G' be a bounded degree graph that is rough isometric to G. Is the the free uniform spanning forest of G connected almost surely?*

Conjecture. *Let D be a domain. Then either every simple, bounded degree triangulation admitting a circle packing in D has an almost surely connected free uniform spanning forest, or every simple, bounded degree triangulation admitting a circle packing in D has an almost surely disconnected free uniform spanning forest.*

Question. *Let G be a unimodular transitive graph, and let $\mathfrak{F}$ be the transboundary uniform spanning forest of G. Is $\mathfrak{F}$ almost surely a topological spanning tree of G?*

Question. *Let T be a bounded degree triangulation that admits a circle packing in the complement of a domain D which is quasi-homogeneous in the sense that the group of conformal automorphisms of D acts cocompactly on D. Is the free uniform spanning forest of T connected almost surely?*

Question. *Let T be a bounded degree triangulation that admits a circle packing in the complement of a positive-length Cantor set. Is the free uniform spanning forest of T disconnected almost surely?*

Question. *Let G be a uniformly transient network with $\inf_e c(e) > 0$. Does it follow that every component of the wired uniform spanning forest of G is one-ended almost surely?*

Question. *Let $d \geq 3$, let $\mathscr{I}$ be the interlacement process on $\mathbb{Z}^d$, and let $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}} = \langle \mathsf{AB}_t(\mathscr{I}) \rangle_{t\in\mathbb{R}}$. If $d = 3, 4$, do there exist times at which $\mathfrak{F}_t$ is disconnected? If $d \geq 5$, do there exist times at which $\mathfrak{F}_t$ is connected?*

Conjecture. *Let (T, ρ) be a simple, CP hyperbolic, unimodular random rooted triangulation with $\mathbb{E}[\deg^2 \rho] < \infty$. Then the boundary of the circle packing of T in $\mathbb{D}$ is a realization of the Martin boundary of T almost surely.*

Part I

Two Short Papers

Chapter 2

Critical percolation on any quasi-transitive graph of exponential growth has no infinite clusters

Summary. We prove that critical percolation on any quasi-transitive graph of exponential volume growth does not have a unique infinite cluster. This allows us to deduce from earlier results that critical percolation on any graph in this class does not have any infinite clusters. The result is new when the graph in question is either amenable or nonunimodular.

2.1 Introduction

In **Bernoulli bond percolation**, each edge of a graph $G = (V, E)$ (which we will always assume to be connected and locally finite) is either deleted or retained at random with retention probability $p \in [0, 1]$, independently of all other edges. We denote the random graph obtained this way by $G[p]$. Connected components of $G[p]$ are referred to as **clusters**. Given a graph G, the **critical probability**, denoted $p_c(G)$ or simply p_c, is defined to be

$$p_c(G) = \sup\left\{p \in [0, 1] : G[p] \text{ has no infinite clusters almost surely}\right\}.$$

A central question concerns the existence or non-existence of infinite clusters at the critical probability $p = p_c$. Indeed, proving that critical percolation on the hypercubic lattice $\mathbb{Z}^d$ has no infinite clusters for every $d \geq 2$ is perhaps the best known open problem in modern probability theory. Russo [203] proved that critical percolation on the square lattice $\mathbb{Z}^2$ has no infinite clusters, while Hara and Slade [117] proved that critical percolation on $\mathbb{Z}^d$ has no infinite clusters for all $d \geq 19$. More recently, Fitzner and van der Hofstad [89] improved upon the Hara-Slade method, proving that critical percolation on $\mathbb{Z}^d$ has no infinite clusters for every $d \geq 11$. See e.g. [106] for further background.

In their highly influential paper [51], Benjamini and Schramm proposed a systematic study of percolation on general **quasi-transitive** graphs; that is, graphs $G = (V, E)$ such that the action of the automorphism group $\mathrm{Aut}(G)$ on V has only finitely many orbits (see e.g. [173] for more detail). They made the following conjecture.

Conjecture 2.1.1 (Benjamini and Schramm)**.** *Let G be a quasi-transitive graph. If $p_c(G) < 1$, then $G[p_c]$ has no infinite clusters almost surely.*

Benjamini, Lyons, Peres, and Schramm [43, 45] verified the conjecture for *nonamenable*, *unimodular*, quasi-transitive graphs, while partial progress has been made for *nonunimodular*, quasi-transitive graphs (which are always nonamenable [173, Exercise 8.30]) by Timár [222] and by Peres, Pete, and Scolnicov [192]. In this note, we verify the conjecture for all quasi-transitive graphs of exponential growth.

Theorem 2.1.2. *Let G be a quasi-transitive graph with exponential growth. Then $G[p_c]$ has no infinite clusters almost surely.*

A corollary of Theorem 2.1.2 is that $p_c < 1$ for all quasi-transitive graphs of exponential growth, a result originally due to Lyons [168].

We prove Theorem 2.1.2 by combining the works of Benjamini, Lyons, Peres, and Schramm [43] and Timár [222] with the following simple connectivity decay estimate. Given a graph G, we write $B(x, r)$ to denote the graph distance ball of radius r around a vertex x of G. Recall that a graph G is said to have **exponential growth** if

$$\mathrm{gr}(G) := \liminf_{r\to\infty} |B(x, r)|^{1/r}$$

is strictly greater than 1 whenever x is a vertex of G. It is easily seen that $\mathrm{gr}(G)$ does not depend on the choice of x. Let $\tau_p(x, y)$ be the probability that x and y are connected in $G[p]$, and let $\kappa_p(n) := \inf\left\{\tau_p(x, y) : x, y \in V,\, d(x, y) \le n\right\}$.

Theorem 2.1.3. *Let G be a quasi-transitive graph with exponential growth. Then*

$$\kappa_{p_c}(n) := \inf\left\{\tau_{p_c}(x, y) : x, y \in V,\, d(x, y) \le n\right\} \le \mathrm{gr}(G)^{-n}$$

for all $n \ge 1$.

Remark 2.1.4. The upper bound on $\kappa_{p_c}(n)$ in Theorem 2.1.3 is attained when G is a regular tree.

There are many amenable groups of exponential growth, and, to our knowledge, the conclusion of Theorem 2.1.2 was not previously known for any of their Cayley graphs. Among probabilists, the best known examples are the *lamplighter groups* [173, 192]. See e.g. [33, 72, 140, 172, 183] for further interesting examples.

Following the work of Lyons, Peres, and Schramm [175, Theorem 1.1], Theorem 2.1.2 has the following immediate corollary, which is new in the amenable case. The reader is referred to [175] and [173] for background on minimal spanning forests.

Corollary 2.1.5. *Let G be a unimodular quasi-transitive graph of exponential growth. Then every component of the wired minimal spanning forest of G is one-ended almost surely.*

2.2 Proof

Proof of Theorem 2.1.2 given Theorem 2.1.3. Let us recall the following results:

Theorem (Newman and Schulman [187]). *Let G be a quasi-transitive graph. Then $G[p]$ has either no infinite clusters, a unique infinite cluster, or infinitely many infinite clusters almost surely for every $p \in [0,1]$.*

Theorem (Burton and Keane [64]; Gandolfi, Keane, and Newman [97]). *Let G be an amenable quasi-transitive graph. Then $G[p]$ has at most one infinite cluster almost surely for every $p \in [0,1]$.*

Theorem (Benjamini, Lyons, Peres, and Schramm [36, 43]). *Let G be a nonamenable, unimodular, quasi-transitive graph. Then $G[p_c]$ has no infinite clusters almost surely.*

Theorem (Timár [222]). *Let G be a nonunimodular, quasi-transitive graph. Then $G[p_c]$ has at most one infinite cluster almost surely.*

The statements given for the first two theorems above are not those given in the original papers; the reader is referred to [173] for a modern account of these theorems and for the definitions of unimodularity and amenability. Similarly, Timár's result is stated for transitive graphs, but his proof easily extends to the quasi-transitive case. For our purposes, the significance of the above theorems is that, to prove Theorem 2.1.2, it suffices to prove that if G is a quasi-transitive graph of exponential growth, then $G[p_c]$ does not have a unique infinite cluster almost surely. This follows immediately from Theorem 2.1.3, since if $G[p_c]$ contains a unique infinite cluster then

$$\begin{aligned}\tau_{p_c}(x,y) &\geq \mathbb{P}_{p_c}(x \text{ and } y \text{ are both in the unique infinite cluster})\\ &\geq \mathbb{P}_{p_c}(x \text{ is in the unique infinite cluster})\mathbb{P}_{p_c}(y \text{ is in the unique infinite cluster})\end{aligned}$$

for all $x, y \in V$ by Harris's inequality [118]. Quasi-transitivity implies that the right hand side is bounded away from zero if $G[p_c]$ contains a unique infinite cluster almost surely, and it follows that $\lim_{n\to\infty} \kappa_{p_c}(n) > 0$ in this case. □

Lemma 2.2.1. *Let G be any graph. Then $\kappa_p(n)$ is a supermultiplicative function of n. That is, for every p, n and m, we have that $\kappa_p(m+n) \geq \kappa_p(m)\kappa_p(n)$.*

Proof. Let u and v be two vertices with $d(u,v) \leq m+n$. Then there exists a vertex w such that $d(u,w) \leq m$ and $d(w,v) \leq n$. Since the events $\{u \leftrightarrow w\}$ and $\{w \leftrightarrow v\}$ are increasing and $\{u \leftrightarrow v\} \supseteq \{u \leftrightarrow w\} \cap \{w \leftrightarrow v\}$, Harris's inequality [118] implies that

$$\tau_p(u,v) \geq \tau_p(u,w)\tau_p(w,v) \geq \kappa_p(m)\kappa_p(n).$$

The claim follows by taking the infimum. □

Lemma 2.2.2. *Let G be a quasi-transitive graph. Then $\sup_{n\geq 1}(\kappa_p(n))^{1/n}$ is left continuous in p. That is,*

$$\lim_{\varepsilon\to 0+}\sup_{n\geq 1}\left(\kappa_{p-\varepsilon}(n)\right)^{1/n} = \sup_{n\geq 1}\left(\kappa_p(n)\right)^{1/n} \qquad \text{for every } p\in(0,1]. \tag{2.2.1}$$

Proof. Recall that an increasing function is left continuous if and only if it is lower semi-continuous, and that lower semi-continuity is preserved by taking minima (of finitely many functions) and suprema (of arbitrary collections of functions). Now, observe that $\tau_p(x,y)$ is lower semi-continuous in p for each pair of fixed vertices x and y: This follows from the fact that $\tau_p(x,y)$ can be written as the supremum of the continuous functions $\tau_p^r(x,y)$, which give the probabilities that x and y are connected in $G[p]$ by a path of length at most r. (See [106, Section 8.3].) Since G is quasi-transitive, there are only finitely many isomorphism classes of pairs of vertices at distance at most n in G, and we deduce that $\kappa_p(n)$ is also lower semi-continuous in p for each fixed n. Thus, $\sup_{n\geq 1}(\kappa_p(n))^{1/n}$ is a supremum of lower semi-continuous functions and is therefore lower semi-continuous itself. □

We will require the following well known theorem.

Theorem 2.2.3. *Let G be a quasi-transitive graph, and let ρ be a fixed vertex of G. Then the expected cluster size is finite for every $p < p_c$. That is,*

$$\sum_x \tau_p(\rho,x) < \infty \qquad \text{for every } p < p_c.$$

This theorem was proven in the transitive case by Aizenmann and Barsky [4], and in the quasi-transitive case by Antunović and Veselić [25]; see also the recent work of Duminil-Copin and Tassion [83] for a beautiful new proof in the transitive case.

Proof of Theorem 2.1.3. Let ρ be a fixed root vertex of G. For every $p\in[0,1]$ and every $n\geq 1$, we have

$$\kappa_p(n)\cdot\left|B(\rho,n)\right| \leq \sum_{x\in B(\rho,n)}\tau_p(\rho,x) \leq \sum_x \tau_p(\rho,x).$$

Thus, it follows from Theorem 2.2.3, Lemma 2.2.1, and Fekete's Lemma that

$$\sup_{n\geq 1}(\kappa_p(n))^{1/n} = \lim_{n\to\infty}(\kappa_p(n))^{1/n} \leq \limsup_{n\to\infty}\left(\frac{\sum_x \tau_p(\rho,x)}{|B(\rho,n)|}\right)^{1/n} = \operatorname{gr}(G)^{-1}$$

for every $p < p_c$. We conclude by applying Lemma 2.2.2. □

Chapter 3

Collisions of random walks in reversible random graphs

3.1 Introduction

Let G be an infinite, connected, locally finite graph. G is said to have the **infinite collision property** if for every vertex v of G, two independent random walks $\langle X_n \rangle_{n \geq 0}$ and $\langle Y_n \rangle_{n \geq 0}$ started from v collide (i.e. occupy the same vertex at the same time) infinitely often almost surely (a.s.). Although transitive recurrent graphs such as $\mathbb{Z}$ and $\mathbb{Z}^2$ are easily seen to have the infinite collision property, Krishnapur and Peres [158] showed that the infinite collision property does not hold for the *comb graph*, a subgraph of $\mathbb{Z}^2$.

Chen and Chen [71] proved that the infinite cluster of supercritical Bernoulli bond percolation in $\mathbb{Z}^2$ a.s. has the infinite collision property. Barlow, Peres and Sousi [32] gave a sufficient condition for the infinite collision property in terms of the Green function. They deduced that several classical random recurrent graphs have the infinite collision property, including the incipient infinite percolation cluster in dimensions $d \geq 19$.

However, the methods of [71] and [32] both require precise estimates on the graphs under consideration, and the infinite collision property was still not known to hold for several important random recurrent graphs – see Corollary 3.1.2.

In this note we prove that the infinite collision property holds a.s. for a large class of random recurrent graphs. Recall that a **rooted graph** (G, ρ) is a graph G together with a distinguished root vertex ρ, and that a random rooted graph (G, ρ) is said to be **reversible** if (G, ρ, X_1) and (G, X_1, ρ) have the same distribution, where X_1 is the first step of a simple random walk on G started at ρ (see Section 3.1.1 for more details).

Theorem 3.1.1. *Let (G, ρ) be a recurrent reversible random rooted graph. Then G has the infinite collision property almost surely.*

The assumption of reversibility can be replaced either by the assumption that (G, ρ) is *stationary* or by the assumption that (G, ρ) is *unimodular* and the root ρ has finite expected degree (see Section 3.1.1 for definitions of these terms).

Corollary 3.1.2. *Each of the following graphs has the infinite collision property almost surely. (This is by no means an exhaustive list.)*

1. *The Uniform Infinite Planar Triangulation (UIPT) and Quadrangulation (UIPQ).*

2. *The Incipient Infinite Cluster (IIC) of Bernoulli bond percolation in $\mathbb{Z}^2$.*

3. *Every component of each of the Wired Uniform and Minimal Spanning Forests (WUSF and WMSF) of any Cayley graph.*

The UIPT was introduced by Angel and Schramm [24] and the UIPQ was introduced by Krikun [157]. They are unimodular by construction and were shown to be recurrent by Gurel-Gurevich and Nachmias [107]. The IIC is an infinite random subgraph of $\mathbb{Z}^2$ introduced by Kesten [153]. For each n, let C_n be the largest cluster of a critical Bernoulli bond percolation on the box $[0,n]^2$ and let ρ_n be a uniformly random vertex of C_n. Járai [136] showed that the IIC can be defined as the weak limit of the random rooted graphs (C_n, ρ_n), and is therefore unimodular [7, §2]. For background on the Wired Uniform and Minimal Spanning Forests, see Chapters 10 and 11 of [173] and Section 7 of [7].

In Section 3.3 we provide extensions of Theorem 3.1.1 to networks and to the continuous-time random walk.

Remark 3.1.3. If G is a non-bipartite graph with the infinite collision property, it is easy to see that two independent random walks started from *any* two vertices of G will collide infinitely often a.s. On the other hand, if G is a bipartite graph with the infinite collision property, then two independent random walks on G will collide infinitely often if and only if their starting points are at an even distance from each other.

3.1.1 Definitions

In this section we give concise definitions of stationary, reversible and unimodular random rooted graphs. We refer the reader to Aldous and Lyons [7] for more details.

A **rooted graph** (G, ρ) is a connected, locally finite (multi)graph $G = (V, E)$ together with a distinguished vertex ρ, the **root**. An isomorphism of graphs $\phi : G \to G'$ is an isomorphism of rooted graphs $\phi : (G, \rho) \to (G', \rho')$ if $\phi(\rho) = \rho'$. The set of isomorphism classes of rooted graphs is endowed with the **local topology** [50], in which, roughly speaking, two (isomorphism classes of) rooted graphs are close to each other if and only if they have large isomorphic balls around the root. A **random rooted graph** is a random variable taking values in the space of isomorphism classes of rooted graphs endowed with the local topology. Similarly, a **doubly-rooted graph** is a graph together with an ordered pair of distinguished (not necessarily distinct) vertices. Denote the space of isomorphism classes of doubly-rooted graphs equipped with this topology by $\mathcal{G}_{\bullet\bullet}$.

Recall that the **simple random walk** on a locally finite (multi)graph $G = (V, E)$ is the Markov process $\langle X_n \rangle_{n \geq 0}$ on the state space V with transition probabilities $p(u, v)$ defined to be the fraction of edges emanating from u that end in v. A random rooted graph (G, ρ) is said to be **stationary** if, when $\langle X_n \rangle_{n\geq 0}$ is a simple random walk on G started at the root,

$$(G, \rho) \stackrel{d}{=} (G, X_n)$$

for all n and is said to be **reversible** if

$$(G, \rho, X_n) \stackrel{d}{=} (G, X_n, \rho)$$

for all n. Note that this is not the same as the reversibility of the random walk on G, which holds for any graph[7]. Every reversible random rooted graph is clearly stationary, but the converse need not hold in general [36, Examples 3.1 and 3.2]. However, Benjamini and Curien [38, Theorem 4.3] showed that every recurrent stationary random rooted graph is necessarily reversible, so that Theorem 3.1.1 also applies under the apparently weaker assumption of stationarity.

Reversibility is closely related to the property of unimodularity. A **mass transport** is a function $f : \mathcal{G}_{\bullet\bullet} \to [0, \infty]$. A random rooted graph (G, ρ) is said to be **unimodular** if it satisfies the **Mass-Transport Principle**: for every mass transport f,

$$\mathbb{E}\left[\sum_v f(G, \rho, v)\right] = \mathbb{E}\left[\sum_u f(G, u, \rho)\right]. \tag{MTP}$$

That is,

Expected mass out equals expected mass in.

The Mass-Transport Principle was first introduced by Häggström [110] to study dependent percolation on Cayley graphs. The current formulation of the Mass-Transport Principle was suggested by Benjamini and Schramm [50] and developed systematically by Aldous and Lyons in [7].

As noted in [38], if (G, ρ) is a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$, then biasing the law of (G, ρ) by $\deg(\rho)$ gives an equivalent law of a reversible random rooted graph. Conversely, if (G, ρ) is a reversible random rooted graph, then biasing the law of (G, ρ) by $\deg(\rho)^{-1}$ gives an equivalent law of a reversible random rooted graph. For example, if (G, ρ) is a finite random rooted graph then it is unimodular if and only if ρ is uniformly distributed on G, and is reversible if and only if ρ is distributed according to the stationary measure of simple random walk on G.

In light of the above correspondence, Theorem 3.1.1 may be stated equivalently as follows.

Theorem 3.1.4. *Let (G, ρ) be a recurrent unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then G has the infinite collision property almost surely.*

We provide two variations on the proof of Theorem 3.1.1. The first uses the Mass-Transport Principle. The second, given in Section 3.3, uses reversibility, and applies in the network setting also.

[7] Rather, a random rooted graph (G, ρ) is reversible if and only if the $\mathcal{G}_{\bullet\bullet}$-valued Markov process $\langle (G, X_n, X_{n+1}) \rangle_{n\geq 0}$ is reversible.

3.2 Proof of Theorem 3.1.1

Proof. Let G be a graph and let $p_n(\cdot\,,\cdot)$ denote the n-step transition probabilities for simple random walk on G. For each vertex u of G, let $q_{\text{fin}}(u)$ denote the probability that two independent random walks started at u collide only finitely often, and let $q_0(u)$ denote the probability that two independent random walks started at u do not collide at all after time zero. Finally, for each pair of vertices u and v let $q_{\text{last}}(u,v)$ be the probability that two independent random walks started at u collide for the last time at v, so that

$$q_{\text{fin}}(u) = \sum_v q_{\text{last}}(u,v).$$

Decomposing according to the time of the last collision gives

$$q_{\text{last}}(u,v) = \sum_{n\geq 0} p_n(u,v)^2 q_0(v) \tag{3.2.1}$$

and hence

$$q_{\text{fin}}(u) = \sum_v \sum_{n\geq 0} p_n(u,v)^2 q_0(v). \tag{3.2.2}$$

Suppose that (G,ρ) is a recurrent unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Consider the mass transport

$$f(G,u,v) = \deg(u) q_{\text{last}}(u,v).$$

Each vertex u sends a total mass of $\deg(u)q_{\text{fin}}(u)$, while, by (3.2.1), each vertex v receives a total mass of

$$\begin{aligned}\sum_u f(G,u,v) &= \sum_{n\geq 0}\sum_u \deg(u) p_n(u,v)^2 q_0(v)\\ &= q_0(v)\sum_{n\geq 0}\sum_u \deg(v) p_n(v,u) p_n(u,v)\\ &= q_0(v)\deg(v)\sum_{n\geq 0} p_{2n}(v,v).\end{aligned}$$

Since G is recurrent, the sum $\sum_{n\geq 0} p_{2n}(v,v)$ is infinite a.s. for every vertex v of G. By the Mass-Transport Principle,

$$\mathbb{E}[\deg(\rho)q_{\text{fin}}(\rho)] = \mathbb{E}\left[q_0(\rho)\deg(\rho)\sum_{n\geq 0} p_{2n}(\rho,\rho)\right].$$

Since the left-hand expectation is finite by assumption, we must have that $q_0(\rho) = 0$ a.s., and consequently that $q_{\text{fin}}(v) = 0$ for every vertex v in G a.s. $\square$

3.3 Extensions

3.3.1 Networks

Recall that a **network** (G, c) is a connected locally finite graph $G = (V, E)$ together with a function $c : E \to (0, \infty)$ assigning to each edge e of G a positive **conductance** $c(e)$. Graphs may be considered to be networks by setting $c(e) \equiv 1$. Write $c(u)$ for the sum of the conductances of the edges emanating from u and $c(u, v)$ for the sum of the conductances of the edges joining u and v. The random walk $\langle X_n \rangle_{n \geq 0}$ on a network (G, c) is the Markov chain on V with transition probabilities $p(u, v) = c(u, v)/c(u)$. Unimodular and reversible random rooted networks are defined similarly to the unweighted case [7]. In particular, a random rooted network (G, c, ρ) is defined to be reversible if, letting $\langle X_n \rangle_{n \geq 0}$ be a random walk on (G, c) started from ρ,

$$(G, c, \rho, X_n) \stackrel{d}{=} (G, c, X_n, \rho) \text{ for all } n \geq 1.$$

We now extend Theorem 3.1.1 to the setting of reversible random rooted networks. The proof given also yields an alternative proof of Theorem 3.1.1.

Theorem 3.3.1. *Let (G, c, ρ) be a recurrent reversible random rooted network. Then (G, c) has the infinite collision property almost surely.*

Proof. Let q_{fin} and q_0 be defined as in the proof of Theorem 3.1.1. Taking expectations on both sides of (3.2.2) with $u = \rho$,

$$\mathbb{E}\left[q_{\text{fin}}(\rho)\right] = \mathbb{E}\left[\sum_v \sum_{n \geq 0} p_n(\rho, v)^2 q_0(v)\right] = \sum_{n \geq 0} \mathbb{E}\big[p_n(\rho, X_n) q_0(X_n)\big].$$

Applying reversibility,

$$\mathbb{E}[q_{\text{fin}}(\rho)] = \sum_{n \geq 0} \mathbb{E}\big[p_n(X_n, \rho) q_0(\rho)\big] = \mathbb{E}\left[q_0(\rho) \sum_{n \geq 0} p_{2n}(\rho, \rho)\right].$$

Since (G, c) is recurrent, $\sum_{n \geq 0} p_{2n}(v, v) = \infty$ a.s. for every vertex v of G. Thus, since the left-hand expectation is finite, we must have that $q_0(\rho) = 0$ a.s. and hence also that $q_{\text{fin}}(\rho) = 0$ a.s. $\square$

3.3.2 Continuous-time random walk

Let (G, c, ρ) be a recurrent unimodular random rooted network, and let $\langle X_t \rangle_{t \geq 0}$ denote the continuous-time random walk on (G, c) started from ρ, which jumps across each edge e with rate $c(e)$ (see e.g. [188]). Corollary 4.3 of [7] states that if (G, c, ρ) is a unimodular random rooted network such that the continuous-time random walk on (G, c) a.s. does not explode (i.e. make infinitely many jumps in a finite amount of time), then (G, c, ρ) is reversible for the continuous-time random walk in the sense that

$$(G, c, \rho, X_t) \stackrel{d}{=} (G, c, X_t, \rho) \tag{3.3.1}$$

for all $t \geq 0$. It is clear that the continuous-time walk on any graph can make at most finitely many visits to any fixed vertex in a finite amount of time, and it follows that the continuous-time walk on a recurrent network does not explode a.s. Thus, we deduce that eq. (3.3.1) holds whenever (G, c, ρ) is a recurrent reversible random rooted network (in particular, we do not require the assumption that $\mathbb{E}[c(\rho)] < \infty$).

The natural analogue of the infinite collision property also holds a.s. for the continuous-time random walk on recurrent unimodular random rooted networks.

Theorem 3.3.2. *let (G, c, ρ) be a recurrent unimodular random rooted network and let $\langle X_t \rangle_{t \geq 0}$ and $\langle Y_t \rangle_{t \geq 0}$ be two independent continuous-time random walks on (G, c) started from any two vertices. Then the set of times $\{t : X_t = Y_t\}$ has infinite Lebesgue measure almost surely.*

Proof. Let $\langle X_t \rangle_{t \geq 0}$ and $\langle Y_t \rangle_{t \geq 0}$ be independent continuous-time random walks starting at the same vertex u of G. By considering the walks only at integer times $t = n$, the proof of Theorem 3.3.1 readily shows that the set of integer collision times $\{n \in \mathbb{N} : X_n = Y_n\}$ is infinite a.s.

For every $s \geq 0$ there is a positive probability that neither $\langle X_t \rangle_{t \geq 0}$ nor $\langle Y_t \rangle_{t \geq 0}$ has made any jumps by time s. Thus, the law of the sequence $\langle (X_{n+s}, Y_{n+s}) \rangle_{n \geq 1}$ is absolutely continuous with respect to the law of $\langle (X_n, Y_n) \rangle_{n \geq 1}$. It follows that for every $s \geq 0$, the walks $\langle X_t \rangle_{t \geq 0}$ and $\langle Y_t \rangle_{t \geq 0}$ collide at an infinite set of times of the form $\{n + s : n \in \mathbb{N}\}$ a.s., and consequently that

$$\mathcal{L}eb\big(\{t : X_t = Y_t\}\big) = \int_0^1 \big|\{n : X_{n+s} = Y_{n+s}\}\big| \, \mathrm{d}s = \infty \quad \text{a.s.}$$

For every other vertex v, there is a positive probability that $X_1 = u$ and $Y_1 = v$, so that the law of two independent continuous-time random walks started from u and v is absolutely continuous with respect to the law of $\langle (X_{t+1}, Y_{t+1}) \rangle_{t \geq 0}$. Thus, the set of collision times of two independent continuous-time random walks started from u and v has infinite Lebesgue measure a.s. □

Remark 3.3.3. Theorem 3.3.2 has consequences for the *voter model* on unimodular random recurrent networks. For every network (G, c) in which two independent continuous-time random walks collide a.s., duality between the voter model and continuous-time coalescing random walk ([166, §5] and [6, §14]) implies that the only ergodic stationary measures for the voter model on (G, c) are the constant (a.k.a. consensus) measures. Thus, a consequence of Theorem 3.3.2 is that this holds for the voter model on recurrent unimodular random rooted networks.

Part II

Uniform Spanning Forests

Chapter 4

Wired cycle-breaking dynamics for uniform spanning forests

Summary. We prove that every component of the wired uniform spanning forest (WUSF) is one-ended almost surely in every transient reversible random graph, removing the bounded degree hypothesis required by earlier results. We deduce that every component of the WUSF is one-ended almost surely in every supercritical Galton-Watson tree, answering a question of Benjamini, Lyons, Peres and Schramm.

Our proof introduces and exploits a family of Markov chains under which the oriented WUSF is stationary, which we call the *wired cycle-breaking dynamics*.

4.1 Introduction

The **uniform spanning forests** (USFs) of an infinite, locally finite, connected graph G are defined as infinite-volume limits of uniformly chosen random spanning trees of large finite subgraphs of G. These limits can be taken with respect to two extremal boundary conditions, **free** and **wired**, giving the **free uniform spanning forest** (FUSF) and **wired uniform spanning forest** (WUSF) respectively (see Section 4.2 for detailed definitions). The study of uniform spanning forests was initiated by Pemantle [190], who, in addition to showing that both limits exist, proved that the wired and free forests coincide in $\mathbb{Z}^d$ for all d and that they are almost surely a single tree if and only if $d \leq 4$. The question of connectivity of the WUSF was later given a complete answer by Benjamini, Lyons, Peres and Schramm (henceforth referred to as BLPS) in their seminal work [44], in which they proved that the WUSF of a graph is connected if and only if two independent random walks on the graph intersect almost surely [44, Theorem 9.2].

After connectivity, the most basic topological property of a forest is the number of ends its components have. An infinite connected graph G is said to be k**-ended** if, over all finite sets of vertices W, the graph $G \setminus W$ formed by deleting W from G has a maximum of k distinct infinite connected components. In particular, an infinite tree is one-ended if and only if it does not contain any simple bi-infinite paths and is two-ended if and only if it contains a unique simple bi-infinite path.

Components of the WUSF are known to be one-ended for several large classes of graphs. Again, this problem was first studied by Pemantle [190], who proved that the USF on $\mathbb{Z}^d$ has one end for $2 \leq d \leq 4$ and that every component has at most two ends for $d \geq 5$. (For $d = 1$ the forest is all of $\mathbb{Z}$ and is therefore two-ended.) A decade later, BLPS [44, Theorem 10.1] completed and extended Pemantle's result, proving in particular that every component of the WUSF of a Cayley graph is one-ended almost surely if and only if the graph is not itself two-ended. Their proof was then adapted to random graphs by Aldous and Lyons [7, Theorem 7.2], who showed that all WUSF components are one-ended almost surely in every transient reversible random rooted graph with bounded vertex degrees. Taking a different approach, Lyons, Morris and Schramm [170] gave an isoperimetric condition for one-endedness, from which they deduced that all WUSF components are one-ended almost surely in every transient transitive graph and every non-amenable graph.

In this paper, we remove the bounded degree assumption from the result of Aldous and Lyons [7]. We state our result in the natural generality of reversible random rooted networks. Recall that a **network** is a locally finite, connected (multi)graph $G = (\mathsf{V}, \mathsf{E})$ together with a function $c : \mathsf{E} \to (0, \infty)$ assigning a positive **conductance** $c(e)$ to each unoriented edge e of G. For each vertex v, the conductance $c(v)$ of v is defined to be the sum of the conductances of the edges adjacent to v, where self-loops are counted twice. Locally finite, connected graphs without specified conductances are considered to be networks by setting $c \equiv 1$. The WUSF of a network is defined in Section 4.2 and reversible random rooted networks are defined in Section 4.5.

Theorem 4.1.1. *Let (G, ρ) be a transient reversible random rooted network and suppose that $\mathbb{E}[c(\rho)^{-1}] < \infty$. Then every component of the wired uniform spanning forest of G is one-ended almost surely.*

The condition that the expected inverse conductance of the root is finite is always satisfied by graphs, for which $c(\rho) = \deg(\rho) \geq 1$. In Example 4.5.1 we show that the theorem can fail in the absence of this condition.

Theorem 4.1.1 applies (indirectly) to supercritical Galton-Watson trees conditioned to survive, answering positively Question 15.4 of BLPS [44].

Corollary 4.1.2. *Let T be a supercritical Galton-Watson tree conditioned to survive. Then every component of the wired uniform spanning forest of T is one-ended almost surely.*

Previously, this was known only for supercritical Galton-Watson trees with offspring distribution either bounded, in which case the result follows as a corollary to the theorem of Aldous and Lyons [7], or supported on a subset of $[2, \infty)$, in which case the tree is non-amenable and we may apply the theorem of Lyons, Morris and Schramm [170].

Our proof introduces a new and simple method, outlined as follows. For every transient network, we define a procedure to 'update an oriented forest at an edge', in which the edge is added to the forest while another edge is deleted. Updating oriented forests at randomly chosen edges defines a family of Markov chains on oriented spanning forests, which we call the *wired cycle-breaking dynamics*, for which the oriented wired uniform spanning forest measure is stationary (Proposition 4.3.2). This stationarity allows us to prove the following theorem, from which we show Theorem 4.1.1 to follow by known methods.

Theorem 4.1.3. *Let G be any network. If the wired uniform spanning forest of G contains more than one two-ended component with positive probability, then it contains a component with three or more ends with positive probability.*

The case of *recurrent* reversible random rooted graphs remains open, even under the assumption of bounded degree. In this case, it should be that the single tree of the WUSF has the same number of ends as the graph (this prediction appears in [7]). BLPS proved this for transitive recurrent graphs [44, Theorem 10.6].

4.1.1 Consequences

The one-endedness of WUSF components has consequences of fundamental importance for the *Abelian sandpile model*. Járai and Werning [138] proved that the infinite-volume limit of the sandpile measures exists on every graph for which every component of the WUSF is one-ended almost surely. Furthermore, Járai and Redig [137] proved that, for any graph which is both transient and has one-ended WUSF components, the sandpile configuration obtained by adding a single grain of sand to the infinite-volume random sandpile can be stabilized by finitely many topplings (their proof is given for $\mathbb{Z}^d$ but extends to this setting, see [135]). Thus, a consequence of Theorem 4.1.1 is that these properties hold for the Abelian sandpile model on transient reversible random graphs of unbounded degree.

Theorem 4.1.1 also has several interesting consequences for random plane graphs, which we address in upcoming work with Angel, Nachmias and Ray [20]. In particular, we deduce from Theorem 4.1.1 that every Benjamini-Schramm limit of finite planar graphs is almost surely Liouville, i.e. does not admit non-constant bounded harmonic functions.

4.2 The wired uniform spanning forest

In this section we briefly define the wired uniform spanning forest and introduce the properties that we will need. For a comprehensive treatment of uniform spanning trees and forests, as well as a detailed history of the subject, we refer the reader to Chapters 4 and 10 of [173].

Notation and orientation Throughout this paper, the graphs on which the USFs and USTs are defined will be connected and locally finite unless stated otherwise. We do not distinguish notationally between oriented and unoriented trees, forests or edges. Whether or not a tree, forest or edge is oriented will be clear from context. Edges e are oriented from their tail e^- to their head e^+, and have reversal $-e$. An oriented tree or forest is a tree or forest together with an orientation of its edges. Given an oriented tree or forest in a graph, we define the **past** of each vertex v to be the set of vertices u for which there is a directed path from u to v in the oriented tree or forest.

For a finite connected graph G, we write UST_G for the uniform measure on the set of spanning trees (i.e. connected cycle-free subgraphs containing every vertex) of G, considered for measure-theoretic purposes to be functions from E to $\{0,1\}$. More generally, if G is a finite network, we define UST_G to be the probability measure on spanning trees of G for which the measure of a tree t is proportional to the product of the conductances of its edges.

There are two extremal (with respect to stochastic ordering) ways to define infinite volume limits of the uniform spanning tree measures. Let G be an infinite network and let V_n be an increasing sequence of finite connected subsets of V such that $\bigcup V_n = V$, which we call an **exhaustion** of G. For each n, let the network G_n be the subgraph of G induced by V_n together with the conductances inherited from G. The weak limit of the measures UST_{G_n} is known as the **free uniform spanning forest**: for each finite subset $S \subset \mathsf{E}$,

$$\mathsf{FUSF}_G(S \subseteq \mathfrak{F}) := \lim_{n\to\infty} \mathsf{UST}_{G_n}(S \subseteq T),$$

where $\mathfrak{F}$ is a sample of the FUSF of G and T is a sample of the UST of G_n. Alternatively, at each step of the exhaustion we define a network G_n^* by identifying ('wiring') $V \setminus V_n$ into a single vertex ∂_n and deleting all the self-loops that are created, and define the **wired uniform spanning forest** to be the weak limit

$$\mathsf{WUSF}_G(S \subseteq \mathfrak{F}) := \lim_{n\to\infty} \mathsf{UST}_{G_n^*}(S \subseteq T),$$

where $\mathfrak{F}$ is a sample of the WUSF of G and T is a sample of the UST of G_n^*.

Both limits were shown (implicitly) to exist for every network and every choice of exhaustion by Pemantle [190], although the WUSF was not defined explicitly until the work of Häggström [109]. As a consequence, the limits do not depend on the choice of exhaustion. Both measures are supported on spanning forests (i.e. cycle-free subgraphs containing every vertex) of G for which every connected component is infinite. The WUSF is usually much more tractable, thanks in part to Wilson's algorithm rooted at infinity, which both connects the WUSF to loop-erased random walk and allows us to sample the WUSF of an infinite network directly rather than by passing to an exhaustion.

Wilson's algorithm [228] is a remarkable method of generating the UST on a finite or recurrent network by joining together loop-erased random walks. It was extended to generate the WUSF of transient networks by BLPS [44]. Let G be a network, and let γ be a path in G that is either finite or transient, i.e. visits each vertex of G at most finitely many times. The **loop-erasure** $\mathsf{LE}(\gamma)$ is formed by erasing cycles from γ chronologically as they are created. Formally, $\mathsf{LE}(\gamma)_i = \gamma_{t_i}$ where the times t_i are defined recursively by $t_0 = 0$ and $t_i = 1 + \max\{t \geq t_{i-1} : \gamma_t = \gamma_{t_{i-1}}\}$. (In the presence of multiple edges, a path is not determined by its vertex-trajectory. However, the definition of the loop-erasure extends to this setting in the obvious way. Similarly, when performing Wilson's algorithm in the presence of multiple edges, we consider the random walks and their loop-erasures to be random paths in the graph.) Let $\{v_j : j \in \mathbb{N}\}$ be an enumeration of the vertices of G and define a sequence of forests in G as follows:

1. If G is finite or recurrent, choose a root vertex v_0 and let $\mathfrak{F}_0$ include v_0 and no edges (in which case we call the algorithm **Wilson's algorithm rooted at** v_0). If G is transient, let $\mathfrak{F}_0 = \emptyset$

(in which case we call the algorithm **Wilson's algorithm rooted at infinity**).

2. Given $\mathfrak{F}_i$, start an independent random walk from v_{i+1} stopped if and when it hits the set of vertices already included in $\mathfrak{F}_i$.

3. Form the loop-erasure of this random walk path and let $\mathfrak{F}_{i+1}$ be the union of $\mathfrak{F}_i$ with this loop-erased path.

4. Let $F = \bigcup F_i$.

This is Wilson's algorithm: the resulting forest $\mathfrak{F}$ has law UST_G in the finite case [228] and WUSF_G in the infinite case [44], and is independent of the choice of enumeration.

We also consider oriented spanning trees and forests. Let $\mathsf{OUST}_{G_n^*}$ denote the law of the uniform spanning tree of G_n^* oriented towards the boundary vertex ∂_n, so that every vertex of G_n^* other than ∂_n has exactly one oriented edge emanating from it in the tree, while ∂_n does not have any oriented edges emanating from it. Wilson's algorithm on G_n^* rooted at ∂_n may be modified to produce an oriented tree with law $\mathsf{OUST}_{G_n^*}$ by considering the loop-erased paths in step (2) to be oriented chronologically. If G is transient, making the same modification to Wilson's algorithm rooted at infinity yields a random oriented forest, known as the **oriented wired uniform spanning forest** [44] of G and denoted OWUSF_G. The proof of the correctness of Wilson's algorithm rooted at infinity [44, Theorem 5.1] also shows that, when G_n is an exhaustion of a transient network G, the measures $\mathsf{OUST}_{G_n^*}$ converge weakly to OWUSF_G.

4.3 Wired cycle-breaking dynamics

Let G be an infinite transient network and let $\mathcal{F}(G)$ denote the set of oriented spanning forests f of G such that every vertex has exactly one oriented edge emanating from it in f. For each $f \in \mathcal{F}(G)$ and oriented edge e of G, the update $U(f,e) \in \mathcal{F}(G)$ of f is defined by the following procedure:

Definition 4.3.1 (Updating f at e)**.** If e or its reversal $-e$ is already included in f, or is a self-loop, let $U(f,e) = f$. Otherwise,

- If e^+ is in the past of e^- in f, so that there is a directed path $\langle e_1, \ldots, e_k, d\rangle$ from e^+ to e^- in f, let
$$U(f,e) = f \cup \{-e, -e_1, \ldots, -e_k\} \setminus \{d, e_k, \ldots, e_1\}.$$

- Otherwise, if e^+ is not in the past of e^- in f, let d be the unique oriented edge of f with $d^- = e^-$ and let $U(f,e) = f \cup \{e\} \setminus \{d\}$.

See Figure 1 for examples. Note that in either case, *as unoriented forests*, we have simply that $U(f,e) = f \cup \{e\} \setminus \{d\}$; the change in orientation in the first case ensures that every vertex has exactly one oriented edge emanating from it in $U(f,e)$, so that $U(f,e) \in \mathcal{F}(G)$.

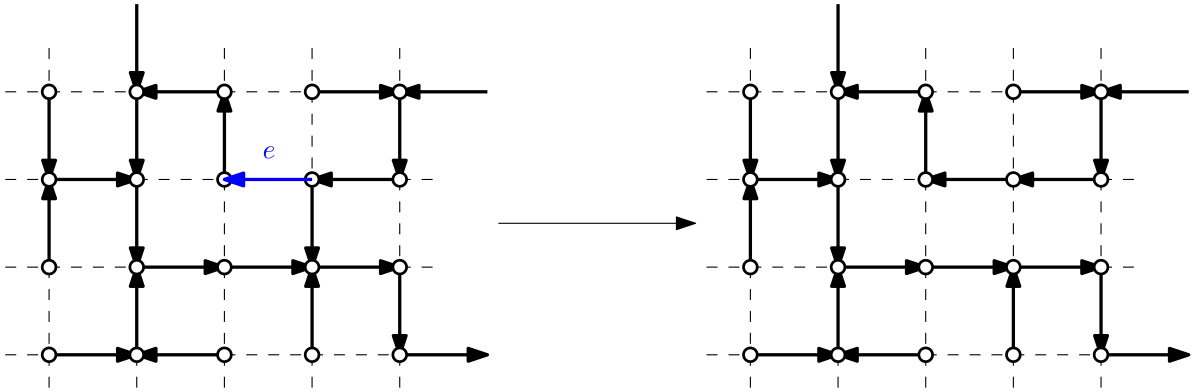

(a) In this example, e^+ is not in the past of e^- in the forest.

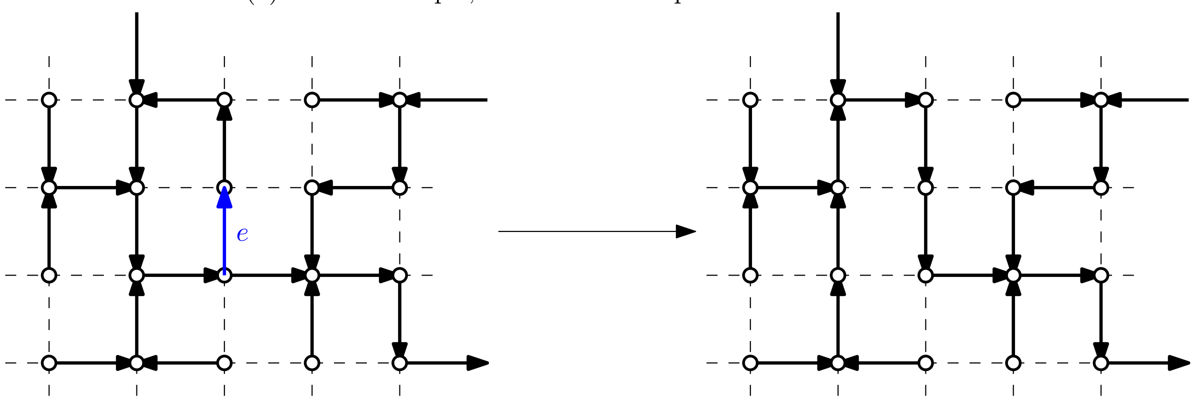

(b) In this example, e^+ is in the past of e^- in the forest.

Figure 4.1: Updating an oriented spanning forest (left, solid black) of $\mathbb{Z}^2$ (dashed black) at an oriented edge e (left, blue) to obtain a new oriented spanning forest (right, solid black). Arrow heads represent orientations of edges.

Let v be a vertex of G. We define the **wired cycle-breaking dynamics rooted at** v to be the Markov chain on $\mathcal{F}(G)$ with transition probabilities

$$p^v(f_0, f_1) = \frac{1}{c(v)} c(\{e : e^- = v \text{ and } U(f_0, e) = f_1\}).$$

That is, we perform a step of the dynamics by choosing an oriented edge randomly from the set $\{e : e^- = v\}$ with probability proportional to its conductance, and then updating at this edge. Dynamics of this form for the UST on *finite* graphs are well-known, see [173, §4.4].

To explain our choice of name for these dynamics, as well as our choice to consider oriented forests, let us give a second, equivalent, description of the update rule.

If e or its reversal $-e$ is already included in f, or is a self-loop, let $U(f, e) = f$. Otherwise,

- If e^+ and e^- are in the same component of f, then $f \cup e$ contains a (not necessarily oriented) cycle. Break this cycle by deleting the unique edge d of f that is both contained in this cycle and adjacent to e^-, letting $\tilde{U}(f,e) = f \cup \{e\} \setminus \{d\}$.

- If e^+ was not in the past of e^- in f, let $U(f,e) = \tilde{U}(f,e)$.
- Otherwise, if e^+ was in the past of e^- in f, then there exists an oriented path from e^- to d^+ in $\tilde{U}(f,e)$. Let $U(f,e)$ be the oriented forest obtained by by reversing each edge in this path.

- If e^+ and e^- are not in the same component of f, we consider e together with the two infinite directed paths in f beginning at e^- and e^+ to constitute a **wired cycle**, or 'cycle through infinity'. Break this wired cycle by deleting the unique edge d in f such that $d^- = e^-$, letting $U(f,e) = f \cup \{e\} \setminus \{d\}$.

The benefit of taking our forests to be oriented is that it allows us to define these wired cycles unambiguously. If every component of the WUSF of G is one-ended almost surely, then there is a unique infinite simple path from each of e^- and e^+ to infinity, so that wired cycles are already defined unambiguously and the update rule may be defined without reference to an orientation.

Proposition 4.3.2. *Let G be an infinite transient network. Then for each vertex v of G, OWUSF_G is a stationary measure for the wired cycle-breaking dynamics rooted at v, i.e. for $p^v(\,\cdot\,,\cdot\,)$.*

Proof. Let $\langle V_n \rangle_{n\geq 1}$ be an exhaustion of G. We may assume that V_n contains v and all of its neighbours for all $n \geq 1$.

Let $\mathcal{T}(G_n^*)$ denote the set of spanning trees of G_n^* oriented towards the boundary vertex ∂_n. For each $t \in \mathcal{T}(G_n^*)$ and oriented edge e with $e^- = v$, we define the update $U(t,e)$ of t at e by the same procedure (Definition 4.3.1) as for $f \in \mathcal{F}(G)$.

Proposition 4.3.3. $U(T_n, E) \overset{d}{=} T_n$ *for every* $n \geq 1$.

Proposition 4.3.3 is a slight variation on the classical *Markov Chain-Tree Theorem* [10, 165, 173]: Define a Markov chain on $\mathcal{T}(G_n^*)$, as we did on $\mathcal{F}(G)$, by

$$p^v(t_0, t_1) = \frac{1}{c(v)} c(\{e : e^- = v \text{ and } U(t_0,e) = t_1\}).$$

The claimed equality in distribution is equivalent to $\mathsf{OUST}_{G_n^*}$ being a stationary measure for $p^v(\,\cdot\,,\cdot\,)$, and so it suffices to verify that $\mathsf{OUST}_{G_n^*}$ satisfies the detailed balance equations for $p^v(\,\cdot\,,\cdot\,)$. This verification, which is both straightforward and similar to that of the classical Markov Chain-Tree Theorem, is omitted.

To complete the proof, we show that $U(T_n, E)$ converges to $U(\mathfrak{F}, E)$ in distribution. It might at first seem that this convergence holds trivially, but in fact some work is required: Updating $\mathfrak{F}$ or T_n at E requires knowledge of whether or not E^+ is in the past of E^-, which cannot necessarily be obtained by observing the tree or forest only within a finite set. A priori, it is therefore possible that E^+ is in the past of E^- in T_n due to the existence of a very long oriented path from E^+ to E^- in T_n that disappears in the limit, obstructing the claimed convergence in distribution. This behaviour will be ruled out by Lemma 4.3.4.

By the Skorokhod representation theorem, there exist random variables $\langle T_n \rangle_{n\geq 1}$ and $\mathfrak{F}$, defined on some common probability space, such that T_n has law $\mathsf{OUST}_{G_n^*}$ for each n, $\mathfrak{F}$ has law OWUSF_G, and T_n converges to $\mathfrak{F}$ almost surely as n tends to infinity. Let E be an oriented edge chosen randomly from the set $\{e : e^- = v\}$ with probability proportional to its conductance, independently of $\langle T_n \rangle_{n\geq 1}$ and $\mathfrak{F}$. We write $\mathbb{P}$ for the probability measure under which $\langle T_n \rangle_{n\geq 1}$, F and E are sampled as indicated. It suffices to prove that $U(T_n, E)$ converges to $U(F, E)$ in probability with respect to $\mathbb{P}$.

Given $\mathfrak{F}$, let R be the length of the longest finite simple path in $\mathfrak{F}$ connecting v to one of its neighbours in G that is in the same component as v in $\mathfrak{F}$. Since T_n converges to $\mathfrak{F}$ almost surely, there exists a random N such that T_n and $\mathfrak{F}$ coincide on the ball $B_R(v)$ of radius R about v in G for all $n \geq N$.

We claim that, with probability tending to one, $\mathfrak{F}$ and T_n agree about whether or not E^+ is in the past of v.

Lemma 4.3.4. *Consider the events*

$$\mathscr{P} = \{E^+ \text{ is in the past of } v \text{ in } \mathfrak{F}\} \text{ and } \mathscr{P}_n = \{E^+ \text{ is in the past of } v \text{ in } T_n\}.$$

The probability of the symmetric difference $\mathscr{P} \triangle \mathscr{P}_n$ converges to zero as $n \to \infty$.

Proof of lemma. Given E, the probability that E^+ is in the past of v in T_n is, by Wilson's algorithm, the probability that v is contained in the loop-erasure of a random walk from E^+ to ∂_n in G_n^*. Since G is transient, this probability converges to the probability that v is contained in the loop-erased random walk from E^+ in G. This probability is exactly the probability that E^+ is in the past of v in $\mathfrak{F}$, and so

$$\mathbb{P}(\mathscr{P}_n) \xrightarrow[n\to\infty]{} \mathbb{P}(\mathscr{P}).$$

If $\mathbb{P}(\mathscr{P}) \in \{0, 1\}$, we are done. Otherwise, on the event $\mathscr{P}$, there is by definition a finite directed path from E^+ to v in $\mathfrak{F}$. This directed path is also contained in T_n for all $n \geq N$ and so

$$\mathbb{P}(\mathscr{P}_n \,|\, \mathscr{P}) \xrightarrow[n\to\infty]{} 1.$$

Combining these two above limits gives

$$\mathbb{P}(\mathscr{P}_n \,|\, \neg\mathscr{P}) = \frac{\mathbb{P}(\mathscr{P}_n) - \mathbb{P}(\mathscr{P}_n \,|\, \mathscr{P})\mathbb{P}(\mathscr{P})}{\mathbb{P}(\neg\mathscr{P})} \xrightarrow[n\to\infty]{} 0.$$

and hence

$$\begin{aligned}\mathbb{P}(\mathscr{P} \triangle \mathscr{P}_n) &= \mathbb{P}(\mathscr{P}) - \mathbb{P}(\mathscr{P} \cap \mathscr{P}_n) + \mathbb{P}(\mathscr{P}_n \cap \neg\mathscr{P}) \\ &= \mathbb{P}(\mathscr{P}) - \mathbb{P}(\mathscr{P}_n \,|\, \mathscr{P})\mathbb{P}(\mathscr{P}) + \mathbb{P}(\mathscr{P}_n \,|\, \neg\mathscr{P})\mathbb{P}(\neg\mathscr{P}) \\ &\xrightarrow[n\to\infty]{} \mathbb{P}(\mathscr{P}) - \mathbb{P}(\mathscr{P}) + 0 = 0.\end{aligned}$$

□

Let $r \geq 1$. Observe that on the event

$$\{T_n \text{ and } \mathfrak{F} \text{ coincide on the ball of radius } \max\{R, r\} \text{ about } v\} \setminus (\mathscr{P} \triangle \mathscr{P}_n),$$

$U(\mathfrak{F}, E)$ and $U(T_n, E)$ coincide on the ball of radius r about v. By Lemma 4.3.4 and the definition of $\mathbb{P}$, the probability of this event converges to 1 as $n \to \infty$, and consequently $U(T_n, E)$ converges to $U(\mathfrak{F}, E)$ in probability with respect to $\mathbb{P}$. □

4.3.1 Update-tolerance

Let G be a transient network and let $\mathfrak{F}$ be a sample of OWUSF_G. An immediate consequence of Proposition 4.3.2 is that for each oriented edge e of G, the law of $U(\mathfrak{F}, e)$ is absolutely continuous with respect to the law of $\mathfrak{F}$.

Corollary 4.3.5. *Let G be a transient network and let e be an oriented edge of G. Then for every event $\mathscr{A} \subset \mathcal{F}(G)$,*

$$\mathsf{OWUSF}_G(\mathfrak{F} \in \mathscr{A}) \geq \frac{c(e)}{c(e^-)} \mathsf{OWUSF}_G(U(\mathfrak{F}, e) \in \mathscr{A}).$$

Proof. By Proposition 4.3.2,

$$\begin{aligned} \mathsf{OWUSF}_G(\mathfrak{F} \in \mathscr{A}) &= \sum_{\hat{e}^- = e^-} \frac{c(\hat{e})}{c(e^-)} \mathsf{OWUSF}_G(U(\mathfrak{F}, \hat{e}) \in \mathscr{A}) \\ &\geq \frac{c(e)}{c(e^-)} \mathsf{OWUSF}_G(U(\mathfrak{F}, e) \in \mathscr{A}). \end{aligned}$$ □

We refer to this property as **update-tolerance** by analogy to the well-established theories of insertion- and deletion-tolerant invariant percolation processes [173, Chapters 7 and 8].

4.4 Proof of Theorem 4.1.3

Proof. Let G be a network such that the WUSF of G contains at least two two-ended connected components with positive probability. Since G's WUSF is therefore disconnected with positive probability, Wilson's algorithm implies that G is necessarily transient. The **trunk** of a two-ended tree is defined to be the unique bi-infinite simple path contained in the tree, or equivalently the set of vertices and edges in the tree whose removal disconnects the tree into two infinite connected components.

Let $\mathfrak{F}_0$ be a sample of OWUSF_G. By assumption, there exists a (non-random) path $\langle \gamma_i \rangle_{i=0}^n$ in G such that, with positive probability, γ_0 and γ_n are in distinct two-ended components of $\mathfrak{F}_0$, γ_n is in the trunk of its component, and γ_i is not in the trunk of γ_n's component for $i < n$. Write $\mathscr{A}_\gamma$ for this event.

For each $1 \leq i \leq n$, let e_i be an edge with $e_i^- = \gamma_i$ and $e_i^+ = \gamma_{i-1}$, and let $\mathfrak{F}_i \in \mathcal{F}(G)$ be defined recursively by

$$F_i = U(F_{i-1}, e_i) \text{ for } 1 \leq i \leq n.$$

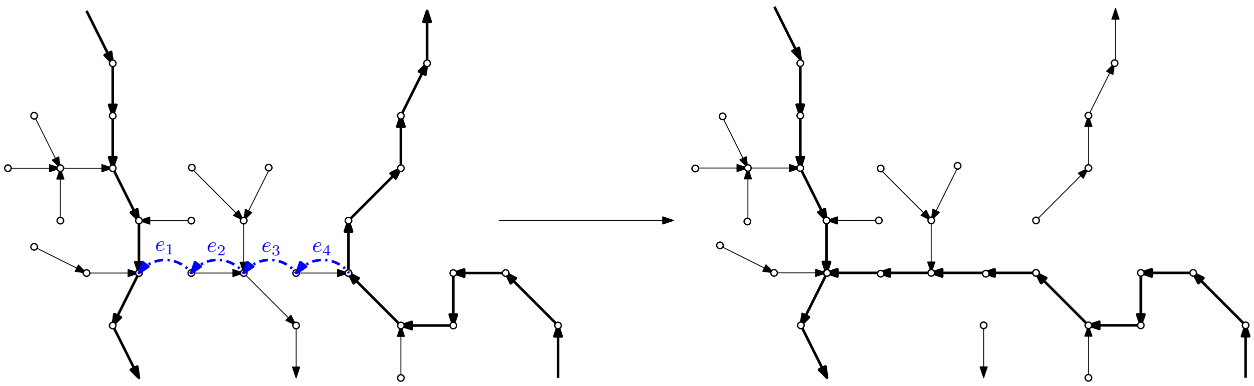

Figure 4.2: When we update along a path (blue arcs) connecting a two-ended component to the trunk of another two-ended component (with each edge oriented backwards), a three-ended component is created. Edges whose removal disconnects their component into two infinite connected components are bold.

We claim that on the event $\mathscr{A}_\gamma$, the component containing γ_n in the updated forest $\mathfrak{F}_n$ has at least three ends. Applying update-tolerance (Corollary 4.3.5) iteratively will then imply that the probability of the WUSF containing a component with three or more ends is at least

$$\mathsf{OWUSF}_G(\mathscr{A}_\gamma)\prod_{i=1}^{n}\frac{c(e_i)}{c(\gamma_i)}$$

which is positive as claimed.

First, notice that γ_i's component in $\mathfrak{F}_i$ has at least two ends for each $0 \leq i \leq n$. This may be seen by induction on i. The component of γ_0 in $\mathfrak{F}_0$ is two-ended by assumption, while for each $0 \leq i < n$:

- If γ_{i+1} is in the same component as γ_i in $\mathfrak{F}_i$, then the component containing γ_{i+1} in the updated forest $\mathfrak{F}_{i+1}$ has the same number of ends and the same vertex set as the component of γ_i in $\mathfrak{F}_i$.

- If γ_{i+1} is in a different component to γ_i in $\mathfrak{F}_i$, then the component containing γ_{i+1} in $\mathfrak{F}_{i+1}$ is equal to the union of the component of γ_i in $\mathfrak{F}_i$, the edge e_i, and the past of γ_{i+1} in $\mathfrak{F}_i$. Thus, the component of γ_{i+1} in $\mathfrak{F}_{i+1}$ has at least as many ends as the component of γ_i in $\mathfrak{F}_i$.

This induction also shows that for every $0 \leq i \leq n$, the component of F_i containing γ_i has vertex set equal to the union of the vertices in the component of F_0 containing γ_0, and the pasts of the vertices γ_j in F_j for $0 \leq j < i$. By definition of the event $\mathscr{A}_\gamma$, the vertex γ_i is not in the trunk of γ_n's component in $\mathfrak{F}_0$ for any $i < n$, and so in particular γ_n is not in the past of γ_i in $\mathfrak{F}_{i-1}$ for any $i < n$, so that γ_{n-1} and γ_n are in different components of $\mathfrak{F}_{n-1}$. Furthermore, since neither endpoint of e_i is contained in the trunk of γ_n's component in $\mathfrak{F}_0$ for any $0 \leq i \leq n-1$, the trunk of γ_n's component in F_0 is still contained in $\mathfrak{F}_{n-1}$. From this, we see that γ_n's component in $\mathfrak{F}_n$ has at least three ends as claimed. See Figure 1 for an illustration. $\square$

4.5 Reversible random networks and the proof of Theorem 4.1.1

A **rooted network** (G, ρ) is a network G together with a distinguished vertex ρ, the root. An isomorphism of graphs is an isomorphism of rooted networks if it preserves the conductances and the root. A **random rooted network** (G, ρ) is a random variable taking values in the space of isomorphism classes of random rooted networks (see [7] for precise definitions, including that of the topology on this space). Similarly, we define **doubly-rooted networks** to be networks together with an ordered pair of distinguished vertices. Let (G, ρ) be a random rooted network and let $\langle X_n \rangle_{n\geq 0}$ be simple random walk on G started at ρ. We say that (G, ρ) is **reversible** if the random doubly-rooted networks (G, ρ, X_n) and (G, X_n, ρ) have the same distribution

$$(G, \rho, X_n) \stackrel{d}{=} (G, X_n, \rho)$$

for every n, or equivalently for $n = 1$. Be careful to note that this is not the same as the reversibility of the random walk on G, which holds for any network. Reversibility is essentialy equivalent to the related property of **unimodularity**. We refer the reader to [7] for a systematic development and overview of the beautiful theory of reversible and unimodular random rooted graphs and networks, as well as many examples.

We now deduce Theorem 4.1.1 from Theorem 4.1.3. Our proof that the WUSF cannot have a unique two-ended component is adapted closely from Theorem 10.3 of [44].

Proof of Theorem 4.1.1. Let (G, ρ) be a reversible random rooted network such that $\mathbb{E}[c(\rho)^{-1}] < \infty$. Biasing the law of (G, ρ) by the inverse conductance $c(\rho)^{-1}$ (that is, reweighting the law of (G, ρ) by the Radon-Nikodym derivative $c(\rho)^{-1}/\mathbb{E}[c(\rho)^{-1}]$) gives an equivalent unimodular random rooted network, as can be seen by checking involution invariance of the biased measure [7, Proposition 2.2]. This allows us to apply Theorem 6.2 and Proposition 7.1 of [7] to deduce that every component of the WUSF of G has at most two ends almost surely. Theorem 4.1.3 then implies that the WUSF of G contains at most one two-ended component almost surely.

Suppose for contradiction that the WUSF contains a single two-ended component with positive probability. Recall that the trunk of this component is defined to be the unique bi-infinite path in the component, which consists exactly of those edges and vertices whose removal disconnects the component into two infinite connected components.

Let $\langle X_n \rangle_{n\geq 0}$ be a random walk on G started at ρ, and let $\mathfrak{F}$ be an independent random spanning forest of G with law WUSF_G, so that (since WUSF_G does not depend on the choice of exhaustion of G) the sequence $\langle (G, X_n, F) \rangle_{n\geq 0}$ is stationary. If the trunk of F is at some distance r from ρ, then X_r is in the trunk with positive probability, and it follows by stationarity that ρ is in the trunk of $\mathfrak{F}$ with positive probability. We will show for contradiction that in fact the probability that the root is in the trunk must be zero.

Recall that, for each n, the forest $\mathfrak{F}$ may be sampled by running Wilson's algorithm rooted at infinity, starting with the vertices ρ and X_n. If we sample $\mathfrak{F}$ in this way and find that both ρ and

X_n are contained in $\mathfrak{F}$'s unique trunk, we must have had either that the random walk started from ρ hit X_n, or that the random walk started from X_n hit ρ. Taking a union bound,

$$\mathbb{P}(\rho \text{ and } X_n \text{ in trunk}) \leq \begin{array}{c} \mathbb{P}(\text{random walk started at } X_n \text{ hits } \rho) \\ +\mathbb{P}(\text{random walk started at } \rho \text{ hits } X_n). \end{array}$$

By reversibility, the two terms on the right hand side are equal and hence

$$\mathbb{P}(\rho \text{ and } X_n \text{ in trunk}) \leq 2\mathbb{P}(\text{random walk started at } X_n \text{ hits } \rho).$$

The probability on the right hand side is now exactly the probability that simple random walk started at ρ returns to ρ at time n or greater, and by transience this converges to zero. Thus,

$$\mathbb{P}(\rho \text{ and } X_n \text{ in trunk}) = \mathbb{E}\left[\mathbb{1}(\rho \text{ in trunk})\mathbb{1}(X_n \text{ in trunk})\right] \xrightarrow[n\to\infty]{} 0$$

and so

$$\mathbb{E}\left[\mathbb{1}(\rho \text{ in trunk})\frac{1}{n}\sum_{1}^{n}\mathbb{1}(X_i \text{ in trunk})\right] \xrightarrow[n\to\infty]{} 0. \qquad (\star)$$

Let $\mathcal{I}$ be the invariant σ-algebra of the stationary sequence $\langle (G, X_n, F)\rangle_{n\geq 0}$. The Ergodic Theorem implies that

$$\frac{1}{n}\sum_{1}^{n}\mathbb{1}(X_i \text{ in trunk}) \xrightarrow[n\to\infty]{\text{a.s.}} \mathbb{P}(\rho \text{ in trunk} \,|\, \mathcal{I}).$$

Finally, combining this with $(\star)$ and the Dominated Convergence Theorem gives

$$\mathbb{E}\left[\mathbb{1}(\rho \text{ in trunk}) \cdot \mathbb{P}(\rho \text{ in trunk} \,|\, \mathcal{I})\right] = \mathbb{E}\left[\mathbb{P}(\rho \text{ in trunk} \,|\, \mathcal{I})^2\right] = 0.$$

It follows that $\mathbb{P}(\rho \text{ in trunk}) = 0$, contradicting our assumption that $\mathfrak{F}$ had a unique two-ended component with positive probability. □

Proof of Corollary 4.1.2. Given a probability distribution $\langle p_k;\ k \geq 0\rangle$ on $\mathbb{N}$, the **augmented Galton-Watson tree** T with offspring distribution $\langle p_k\rangle$ is defined by taking two independent Galton-Watson trees T_1 and T_2, both with offspring distribution $\langle p_k\rangle$, and then joining them by a single edge between their roots. Lyons, Pemantle and Peres [171] proved that T is reversible when rooted at the root of the first tree T_1; See also [7, Example 1.1].

If the distribution $\langle p_k\rangle$ is supercritical (i.e. has expectation greater than 1), then the associated Galton-Watson tree is infinite with positive probability and on this event is almost surely transient [173, Chapter 16]. Thus, Theorem 4.1.1 implies that every component of T's WUSF is one-ended almost surely on the event that either T_1 or T_2 is infinite.

Recall that for every connected graph G and every edge e of G which has a positive probability of not being included in G's WUSF, the law of G's WUSF conditioned not to contain e is equal to $\mathsf{WUSF}_{G\setminus\{e\}}$ [44, Proposition 4.2], where, if $G \setminus \{e\}$ is disconnected, $\mathsf{WUSF}_{G\setminus\{e\}}$ is defined to be the union of independent samples of WUSFs of the two connected components of $G \setminus \{e\}$. Let e be the edge between the roots of T_1 and T_2 that was added to form the augmented tree T. On the positive probability event that T_1 and T_2 are both infinite, running Wilson's algorithm on T started from the roots of T_1 and T_2 shows, by transience of T_1 and T_2, that e has positive probability not to be included in T's WUSF. On this event, T's WUSF is distributed as the union of independent samples of WUSF_{T_1} and WUSF_{T_2}. It follows that every component of T_1's WUSF is one-ended almost surely on the event that T_1 is infinite. □

Example 4.5.1 ($\mathbb{E}[c(\rho)^{-1}] < \infty$ is necessary)**.** Let (T, o) be a 3-regular tree with unit conductances rooted at an arbitrary vertex o. Form a network G by adjoining to each vertex v of T an infinite path, and setting the conductance of the nth edge in each of these paths to be 2^{-n-1}. Let o_n be the nth vertex in the added path at o. Define a random vertex ρ of G which is equal to o with probability $4/7$ and equal to the nth vertex in the path at o with probability $3/(7 \cdot 2^n)$ for each $n \geq 1$. The only possible isomorphism classes of (G, ρ, X_1) are of the form (G, o_n, o_{n+1}), (G, o_{n+1}, o_n), (G, o, o_1), (G, o_1, o), or (G, o, o'), where o' is a neighbour of o in T. This allows us to easily verify that (G, ρ) is a reversible random rooted network:

$$\mathbb{P}((G, \rho, X_1) = (G, o_n, o_{n+1})) = \mathbb{P}((G, \rho, X_1) = (G, o_{n+1}, o_n)) = \frac{1}{7 \cdot 2^n}$$

for all $n \geq 1$ and

$$\mathbb{P}((G, \rho, X_1) = (G, o, o_1)) = \mathbb{P}((G, X_1, \rho) = (G, o, o_1)) = \frac{1}{7}.$$

When we run Wilson's algorithm on G started from a vertex of T, every excursion of the random walk into one of the added paths is erased almost surely. It follows that the WUSF of G is simply the union of the WUSF of T with each of the added paths, and hence every component has infinitely many ends almost surely.

Chapter 5

Interlacements and the wired uniform spanning forest

Summary. We extend the Aldous-Broder algorithm to generate the wired uniform spanning forests (WUSFs) of infinite, transient graphs. We do this by replacing the simple random walk in the classical algorithm with Sznitman's random interlacement process. We then apply this algorithm to study the WUSF, showing that every component of the WUSF is one-ended almost surely in any graph satisfying a certain weak anchored isoperimetric condition, that the number of 'excessive ends' in the WUSF is non-random in any graph, and also that every component of the WUSF is one-ended almost surely in any transient unimodular random rooted graph. The first two of these results answer positively two questions of Lyons, Morris and Schramm, while the third extends a recent result of the author.

Finally, we construct a counterexample showing that almost sure one-endedness of WUSF components is not preserved by rough isometries of the underlying graph, answering negatively a further question of Lyons, Morris and Schramm.

5.1 Introduction

The **uniform spanning forests** (USFs) of an infinite, locally finite, connected graph G are defined as weak limits of uniform spanning trees (USTs) of large finite subgraphs of G. These weak limits can be taken with either free or wired boundary conditions (see Section 5.3.1), yielding the **free uniform spanning forest** (FUSF) and **wired uniform spanning forest** (WUSF) respectively. The USFs are closely related to several other topics in probability theory, including loop-erased random walks [161, 228], potential theory [44, 62], conformally invariant scaling limits [163, 208], domino tiling [150] and the Abelian sandpile model [78, 135]. In this paper, we develop a new connection between the wired uniform spanning forest and Sznitman's **interlacement process** [216, 218].

A key theoretical tool in the study of the UST and USFs is **Wilson's algorithm** [228], which allows us to sample the UST of a finite graph by joining together loop-erasures of random walk paths. In their seminal work [44], Benjamini, Lyons, Peres, and Schramm (henceforth referred to as BLPS) extended Wilson's algorithm to infinite transient graphs and used this extension to establish several fundamental properties of the WUSF. For example, they proved that the WUSF of an infinite, locally finite, connected graph is connected almost surely (a.s.) if and only if the sets of vertices visited by two independent random walks on the graph have infinite intersection a.s. This recovered the earlier, pioneering work of Pemantle [190], who proved that the FUSF and WUSF of $\mathbb{Z}^d$ coincide for all d and are a.s. connected if and only if $d \leq 4$. Wilson's algorithm has also been instrumental in the study of scaling limits of uniform spanning trees and forests [28, 163, 193, 208, 209].

Prior to the introduction of Wilson's algorithm, the best known algorithm for sampling the UST of a finite graph was the **Aldous-Broder algorithm** [8, 59], which generates a uniform spanning tree of a finite connected graph G as the collection of first-entry edges of a random walk on G. We now describe this algorithm in detail. Let ρ be a fixed vertex of G, and let $\langle X_n \rangle_{n\geq 0}$ be a simple random walk on G started at ρ. For each vertex v of G, let $e(v)$ be the edge of G incident to v that is traversed by the random walk X_n as it enters v for the first time, and let $T = \big\{e(v) : v \in V \setminus \{\rho\}\big\}$ be set of first-entry edges. Aldous [8] and Broder [59] proved independently that the resulting random spanning tree T is distributed uniformly on the set of spanning trees of G (see also [173, §4.4]). If we orient the edge in the direction opposite to that in which it was traversed by the random walk, then the spanning tree is oriented towards ρ, meaning that every vertex of G other than ρ has exactly one oriented edge emanating from it in the tree.

While the algorithm extends without modification to generate USTs of recurrent infinite graphs, the collection of first entry edges of a random walk on a transient graph might not span the graph. Thus, naively running the Aldous-Broder on a transient graph will not necessarily produce a spanning forest of the graph. Moreover, unlike in Wilson's algorithm, we cannot simply continue the algorithm by starting another random walk from a new location. As such, it has hitherto been unclear how to extend the Aldous-Broder algorithm to infinite transient graphs and, as a result, the Aldous-Broder algorithm has been of limited theoretical use in the study of USFs of infinite graphs.

In this paper, we extend the Aldous-Broder algorithm to infinite, transient graphs by replacing the random walk with the random interlacement process. The interlacement process was originally introduced by Sznitman [216] to study the disconnection of cylinders and tori by a random walk trajectory, and was generalised to arbitrary transient graphs by Teixeira [218]. The interlacement process $\mathscr{I}$ on a transient graph G is a point process on the space $\mathcal{W}^* \times \mathbb{R}$, where $\mathcal{W}^*$ is the space of doubly-infinite paths in G modulo time-shift (see Section 5.3.3 for precise definitions), and should be thought of as a collection of random walk excursions from infinity. We refer the reader to the monographs [82] and [67] for an introduction to the extensive literature on the random interlacement process.

We state our results in the natural generality of networks. Recall that a **network** (G, c) is a connected, locally finite graph $G = (V, E)$, possibly containing self-loops and multiple edges, together with a function $c : E \to (0, \infty)$ assigning a positive **conductance** $c(e)$ to each edge e of G. The conductance $c(v)$ of a vertex v is defined to be the sum of the conductances of the edges emanating from v. Graphs without specified conductances are considered as networks by setting $c(e) \equiv 1$. We will usually suppress the notation of conductances, and write simply G for a network.

See Section 5.3.1 for detailed definitions of the USFs on general networks.

Oriented edges e are oriented from their tail e^- to their head e^+. The reversal of an oriented edge e is denoted $-e$.

Theorem 5.1.1 (Interlacement Aldous-Broder). *Let G be a transient, connected, locally finite network, let $\mathscr{I}$ be the interlacement process on G, and let $t \in \mathbb{R}$. For each vertex v of G, let $\tau_t(v)$ be the smallest time greater than t such that there exists a trajectory $(W_{\tau_t(v)}, \tau_t(v)) \in \mathscr{I}$ passing through v, and let $e_t(v)$ be the oriented edge of G that is traversed by the trajectory $W_{\tau_t(v)}$ as it enters v for the first time. Then*

$$\mathsf{AB}_t(\mathscr{I}) := \{-e_t(v) : v \in V\}$$

has the law of the oriented wired uniform spanning forest of G.

A useful feature of the interlacement Aldous-Broder algorithm is that it allows us to consider the wired uniform spanning forest of an infinite transient graph as the stationary measure of the ergodic Markov process $\langle \mathsf{AB}_t(\mathscr{I})\rangle_{t\in\mathbb{R}}$. Indeed, it is with this stationarity in mind that we consider the interlacement process to be a point process on $\mathcal{W}^* \times \mathbb{R}$ rather than the more usual $\mathcal{W}^* \times \mathbb{R}_+$. For example, a key step in proving that the number of excessive ends of the WUSF is non-random is to show that the number of *indestructible* excessive ends is a.s. monotone in the time evolution of the process $\langle \mathsf{AB}_t(\mathscr{I})\rangle_{t\in\mathbb{R}}$.

5.2 Applications

5.2.1 Ends

Other than connectivity, the most basic topological property of a forest is the number of ends its components have. Here, an infinite, connected graph $G = (V, E)$ is said to be k**-ended** if, over all finite subsets W of V, the subgraph of G induced by $V \setminus W$ has a supremum of k infinite connected components. In particular, an infinite tree is k-ended if and only if there exist exactly k distinct infinite simple paths starting at each vertex of the tree. Components of the WUSF are known to be one-ended a.s. in several large classes of graphs. The first result of this kind is due to Pemantle [190], who proved that the WUSF of $\mathbb{Z}^d$ is one-ended a.s. for $2 \le d \le 4$, and that every component of the WUSF of $\mathbb{Z}^d$ has at most two ends a.s. for every $d \ge 5$ (the WUSF of $\mathbb{Z}$ is the whole of $\mathbb{Z}$ and is therefore two-ended). BLPS [44] later completed this work, showing in particular that every component of the WUSF is one-ended a.s. in any transient Cayley graph. We note that one-endedness of WUSF components has important consequences for the Abelian sandpile model [135, 137, 138].

Taking a different approach, Lyons, Morris and Schramm [170] gave an isoperimetric criterion for one-endedness of WUSF components, from which they deduced that the every component of the WUSF is one-ended in every transitive graph not rough isometric to $\mathbb{Z}$, and also every non-amenable graph. Unlike the earlier results of BLPS, the results of Lyons, Morris and Schramm are robust in the sense that their assumptions depend only upon the coarse geometry of the graph and do not require any kind of homogeneity. They asked [170, Question 7.9] whether the isoperimetric assumption in their theorem could be replaced by the **anchored** version of the same condition, and in particular whether every WUSF component is one-ended a.s. in any graph with anchored expansion (defined below). Unlike classical isoperimetric conditions, anchored isoperimetric conditions are often preserved under random perturbations such as supercritical Bernoulli percolation [70, 194].

Given a network G and a set K of vertices of G, we write $\partial_E K$ for the set of edges of G with exactly one endpoint in K, and write $|K|$ for the sum of the conductances of the vertices in K. Similarly, if W is a set of edges in G, we write $|W|$ for the sum of the conductances of the edges in W. Given an increasing function $f : (0,\infty) \to (0,\infty)$, we say that G satisfies an **anchored $f(t)$-isoperimetric inequality** if

$$\inf\left\{ \frac{|\partial_E K|}{f\left(|K|\right)} : K \subset V \text{ connected}, v \in K, |K| < \infty \right\} > 0$$

for every vertex v of G. (In contrast, the graph is said to satisfy a (non-anchored) $f(t)$-isoperimetric inequality if the infimum $\inf |\partial_E K|/f(|K|)$ is positive when taken over *all* sets of vertices K with $|K| < \infty$.) In particular, G is said to have **anchored expansion** if and only if it satisfies an anchored t-isoperimetric inequality, and is said to satisfy a **d-dimensional anchored isoperimetric inequality** if it satisfies an anchored $t^{(d-1)/d}$-isoperimetric inequality. Such anchored isoperimetric inequalities are known to hold on, for example, supercritical percolation clusters on $\mathbb{Z}^d$ and related graphs, such as half-spaces and wedges [194].

Theorem 5.2.1. *Let G be a network with $c_0 := \inf_e c(e) > 0$, and suppose that G satisfies an anchored $f(t)$-isoperimetric inequality for some increasing function $f : (0,\infty) \to (0,\infty)$ for which there exists a constant α such that $f(t) \leq t$ and $f(2t) \leq \alpha f(t)$ for all $t \in (0,\infty)$. Suppose that f also satisfies each of the following conditions:*

1.
$$\int_{c_0}^{\infty} \frac{1}{f(t)^2} \mathrm{d}t < \infty$$

and

2.
$$\int_{c_0}^{\infty} \exp\left(-\varepsilon \left(\int_s^{\infty} \frac{1}{f(t)^2} \mathrm{d}t \right)^{-1} \right) \mathrm{d}s < \infty$$

for every $\varepsilon > 0$.

Then every component of the wired uniform spanning forest of G is one-ended almost surely.

In particular, Theorem 5.2.1 applies both to every graph with anchored expansion and to every graph satisfying a d-dimensional anchored isoperimetric inequality with $d > 2$. The graph formed by joining two copies of $\mathbb{Z}^2$ together with a single edge between their origins satisfies a 2-dimensional isoperimetric inequality but has a two-ended WUSF. The theorem can fail if edge conductances are not bounded away from zero, as can be seen by attaching an infinite path with exponentially decaying edge conductances to the root of a 3-regular tree.

Theorem 5.2.1 comes very close to giving a complete answer to [170, Question 7.9]. The isoperimetric condition of [170] is essentially that G satisfies an $f(t)$-isoperimetric inequality for some f satisfying all conditions of Theorem 5.2.1 with the possible exception of (2); the precise condition required is slightly weaker than this but also more technical. Our formulation of Theorem 5.2.1 is adapted from the presentation of the results of [170] given in [173, Theorem 10.43]. The difference in requirements on the function $f(t)$ between Theorem 5.2.1 and [173, Theorem 10.43] can be seen by considering $f(t)$ of the form $t^{1/2}\log^{\alpha}(1+t)$: In particular, we observe that [173, Theorem 10.43] applies to graphs satisfying a $t^{1/2}\log^{\alpha}(1+t)$-isoperimetric inequality for some $\alpha > 1/2$, while our theorem applies to graphs satisfying an *anchored* $t^{1/2}\log^{\alpha}(1+t)$-isoperimetric inequality only if $\alpha > 1$.

In Section 5.6, we give an example of two bounded degree, rough-isometric graphs G and G' such that every component of the WUSF of G is one-ended, while the WUSF of G' a.s. contains a component with uncountably many ends. This answers negatively Question 7.6 of [170], and shows that the behaviour of the WUSF of a graph cannot always be determined from the coarse geometric properties of the graph alone.

5.2.2 Excessive ends

One example of a transient graph in which the WUSF has multiply-ended components is the subgraph of $\mathbb{Z}^6$ spanned by the vertex set

$$\Big(\mathbb{Z}^5 \times \{0\}\Big) \cup \Big(\big\{(0,0,0,0,0),(2,0,0,0,0)\big\} \times \mathbb{N}\Big),$$

which is obtained from $\mathbb{Z}^5$ by attaching an infinite path to each of the vertices $u = (0,0,0,0,0)$ and $v = (2,0,0,0,0)$. The WUSF of this graph, which we denote $\mathfrak{F}$, is equal in distribution to the union of the WUSF of $\mathbb{Z}^5$ with each of the two added paths. If u and v are in the same component of $\mathfrak{F}$, then there is a single component of $\mathfrak{F}$ with three ends and all other components are one-ended. Otherwise, u and v are in different components of $\mathfrak{F}$, so that there are exactly two components of $\mathfrak{F}$ that are two-ended and all other components are one-ended. Each of these events has positive probability, so that the event that there exists a two-ended component of the WUSF has probability strictly between 0 and 1. Nevertheless, the **number of excessive ends** of $\mathfrak{F}$, that is, the sum over all components of $\mathfrak{F}$ of the number of ends minus 1, is equal to two a.s.

In light of this example, Lyons, Morris and Schramm [170, Question 7.8] asked whether the number of excessive ends of the WUSF is non-random (i.e., equal to some constant a.s.) for any graph. Our next application of the interlacement Aldous-Broder algorithm is to answer this question positively.

Theorem 5.2.2. *Let G be a network. Then the number of excessive ends of the wired uniform spanning forest of G is non-random.*

When combined with the spatial Markov property of the wired uniform spanning forest, Theorem 5.2.2 has the following immediate corollary, which states that a natural weakening of [170, Question 7.6] has a positive answer. (As mentioned above, we show the original question to have a negative answer in Section 5.6).

Corollary 5.2.3. *Let G be a network, and suppose that G' is a network obtained from G by adding and deleting finitely many edges. Then the wired uniform spanning forests of G and G' have the same number of excessive ends almost surely. In particular, if every tree of the wired uniform spanning forest of G is one-ended a.s., then the same is true of G'.*

5.2.3 Ends in unimodular random rooted graphs

Another generalisation of the one-endedness theorem of BLPS [44] concerns transient **unimodular random rooted graphs**. A **rooted graph** is a connected, locally finite graph G together with a distinguished vertex ρ, the **root**. An isomorphism of graphs is an isomorphism of rooted graphs if it preserves the root. The **local topology** on the space $\mathcal{G}_\bullet$ of isomorphism classes of rooted graphs is defined so that two rooted graphs are close if they have large graph distance balls around their respective roots that are isomorphic as rooted graphs. Similarly, a doubly rooted graph is a graph together with an ordered pair of distinguished vertices, and the local topology on the space $\mathcal{G}_{\bullet\bullet}$ of isomorphism classes of doubly rooted graphs is defined similarly to the local topology on $\mathcal{G}_\bullet$. A random rooted graph (G,ρ) is said to be **unimodular** if it satisfies the **Mass-Transport Principle**, which states that for every Borel function $f : \mathcal{G}_{\bullet\bullet} \to [0,\infty]$ (which we call a **mass transport**),

$$\mathbb{E}\Big[\sum_{v\in V} f(G,\rho,v)\Big] = \mathbb{E}\Big[\sum_{v\in V} f(G,v,\rho)\Big].$$

In other words, the expected mass sent by the root is equal to the expected mass received by the root. Unimodular random rooted networks are defined similarly by allowing the mass-transport to depend on the edge conductances. We refer the reader to Aldous and Lyons [7] for a systematic development and overview of the theory of unimodular random rooted graphs and networks, as well as several examples.

Aldous and Lyons [7] proved that every component of the WUSF is one-ended a.s. in any bounded degree unimodular random rooted graph, and the author of this article [127] later extended this to all transient unimodular random rooted graphs with finite expected degree at the root, deducing that every component of the WUSF is one-ended a.s. in any supercritical Galton Watson tree. (The assumption of finite expected degree was implicit in [127] since there we considered *reversible* random rooted graphs, which correspond to unimodular random rooted graphs with finite expected degree.) Our final application of the interlacement Aldous-Broder algorithm is to extend the main result of [127] by removing the condition that the expected degree of the root is finite.

Theorem 5.2.4. *Let (G,ρ) be a transient unimodular random rooted network. Then every component of the wired uniform spanning forest of G is one-ended almost surely.*

5.3 Background and definitions

5.3.1 Uniform spanning forests

For each finite graph $G=(V,E)$, let UST_G denote the uniform measure on the set of spanning trees of G (i.e. connected, cycle-free subgraphs of G containing every vertex), which are considered for measure-theoretic purposes to be functions $E \to \{0,1\}$. More generally, for each finite network G, let UST_G denote the probability measure on the set of spanning trees of G such that the probability of a tree is proportional to the product of the conductances of its edges.

Let G be an infinite network, and let $\langle V_n\rangle_{n\geq 0}$ be an exhaustion of V by finite connected subsets, i.e. an increasing sequence of finite connected subsets of V such that $\bigcup V_n = V$. For each n, let G_n denote the subnetwork of G induced by V_n, and let G_n^* denote the finite network obtained by identifying (wiring) every vertex of G in $V \setminus V_n$ into a single vertex ∂_n, and deleting the infinitely many self-loops from ∂_n to itself. The **wired uniform spanning forest** measure is defined as the weak limit of the uniform spanning tree measures on G_n^*. That is, for every finite set $S \subset E$,

$$\mathsf{WUSF}_G(S \subseteq \mathfrak{F}) := \lim_{n\to\infty} \mathsf{UST}_{G_n^*}(S \subseteq T),$$

where $\mathfrak{F}$ is a sample of the WUSF of G and T is a sample of the UST of G_n^*. In contrast, the **free uniform spanning forest** measure is defined as the weak limit of the uniform spanning tree measures on the finite induced subnetworks G_n. It is easily seen that both the free and wired measures are supported on the set of **essential spanning forests** of G, that is, the set of cycle-free subgraphs of G that contain every vertex and do not have any finite connected components.

It will also be useful to consider oriented trees and forests. Given an infinite network G with exhaustion $\langle V_n\rangle_{n\geq 0}$, let $\mathsf{OUST}_{G_n^*}$ denote the law of a uniform spanning tree of G_n^* that has been oriented towards the boundary vertex ∂_n, meaning that every vertex of G_n^* other than ∂_n has exactly one oriented edge emanating from it in the tree. BLPS [44] proved that if G is transient, then the measures $\mathsf{OUST}_{G_n^*}$ converge weakly to a measure OWUSF, the **oriented wired uniform spanning forest** (OWUSF) measure. This measure is supported on the set of essential spanning forests of G that are oriented so that every vertex of G has exactly one oriented edge emanating from it in the forest. The WUSF of a transient graph can be obtained from the OWUSF of the graph by forgetting the orientations of the edges.

5.3.2 The space of trajectories

Let G be a graph. For each $-\infty \leq n \leq m \leq \infty$, let $L(n,m)$ be the graph with vertex set $\{i \in \mathbb{Z} : n \leq i \leq m\}$ and with edge set $\{(i,i+1) : n \leq i \leq m-1\}$. We define $\mathcal{W}(n,m)$ to be the set of multigraph homomorphisms from $L(n,m)$ to G such that the preimage of each vertex in G is finite, and define $\mathcal{W}$ to be the union

$$\mathcal{W} := \bigcup \left\{ \mathcal{W}(n,m) : -\infty \leq n \leq m \leq \infty \right\}.$$

For each set $K \subseteq V$, we let $\mathcal{W}_K(n,m)$ denote the set of $w \in \mathcal{W}(n,m)$ that visit K (that is, for which there exists $n \leq i \leq m$ such that $w(i) \in K$), and let $\mathcal{W}_K$ be the union $\mathcal{W}_K = \bigcup\{\mathcal{W}_K(n,m) : -\infty \leq n \leq m \leq \infty\}$.

Given $w \in \mathcal{W}(n,m)$ and $a \leq b \in \mathbb{Z}$, we define $w|_{[a,b]} \in \mathcal{W}(n \vee a, m \wedge b)$ to be the restriction of w to the subgraph $L(n \vee a, m \wedge b)$ of $L(n,m)$. Given $w \in \mathcal{W}_K(n,m)$, let $H_K^-(w) = \inf\{n \leq i \leq m : w(i) \in K\}$, let $H_K^+(w) = \sup\{n \leq i \leq m : w(i) \in K\}$, and let

$$w_K = w|_{[H_k^-(w), H_K^+(w)]}$$

be the restriction of w to between the first it visits K and last time it visits K. We equip $\mathcal{W}$ with the topology generated by open sets of the form

$$\{w \in \mathcal{W} : w \text{ visits K and } w_K = w'_K\},$$

where $K \subset V$ is finite and $w' \in \mathcal{W}_K$. (Note that this topology is not the weakest topology making the evaluation maps $w \mapsto w(i)$ and $w \mapsto w(i, i+1)$ continuous. First and last hitting times of finite sets are not continuous with respect to that topology, but are continuous with respect to ours. The Borel σ-algebras generated by the two topologies are the same.) We also equip $\mathcal{W}$ with the Borel σ-algebra generated by this topology.

The **time shift** $\theta_k : \mathcal{W} \to \mathcal{W}$ is defined by $\theta_k : \mathcal{W}(n,m) \longrightarrow \mathcal{W}(n-k, m-k)$,

$$\theta_k(w)(i) = w(i+k), \quad \theta_k(w)(i,i+1) = w(i+k, i+k+1).$$

The space $\mathcal{W}^*$ is defined to be the quotient

$$\mathcal{W}^* = \mathcal{W}/\sim \text{ where } w_1 \sim w_2 \text{ if and only if } w_1 = \theta_k(w_2) \text{ for some k.}$$

Let $\pi : \mathcal{W} \to \mathcal{W}^*$ denote the quotient map. $\mathcal{W}^*$ is equipped with the quotient topology and associated quotient σ-algebra. An element of $\mathcal{W}^*$ is called a **trajectory**.

5.3.3 The interlacement process

Given a network $G = (G,c)$, the conductance $c(v)$ of a vertex v is defined to be the sum of the conductances of the edges emanating from v, and, for each pair of vertices (u,v), the conductance $c(u,v)$ is defined to be the sum of the conductances of the edges connecting u to v. Recall that the **random walk** X on the network G is the Markov chain on V with transition probabilities $p(u,v) = c(u,v)/c(u)$. In case G has multiple edges, we will also keep track of the edges crossed by

X, considering X to be a random element of $\mathcal{W}(0,\infty)$. We write either P_u or P_u^G for the law of the random walk started at u on the network G, depending on whether the network under consideration is clear from context. When X is a random walk on a network G and K is a set of vertices in G, we let τ_K denote the first time that X hits K and let τ_K^+ denote the first positive time that X hits K.

Let $G=(V,E)$ be a transient network. Given $w \in \mathcal{W}(n,m)$, let $w^{\leftarrow} \in \mathcal{W}(-n,-m)$ be the reversal of w, defined by setting $w^{\leftarrow}(i) = w(-i)$ for all $-m \le i \le -n$, and $w^{\leftarrow}(i,i+1) = w(-i-1,-i)$ for all $-m \le i \le -n-1$. For each subset $\mathscr{A} \subseteq \mathcal{W}$, let $\mathscr{A}^{\leftarrow}$ denote the set

$$\mathscr{A}^{\leftarrow} := \{w \in \mathcal{W} : w^{\leftarrow} \in \mathscr{A}\}.$$

For each finite set $K \subset V$, define a measure Q_K on $\mathcal{W}_K$ by setting

$$Q_K(\{w \in \mathcal{W} : w(0) \notin K\}) = 0$$

and, for each $u \in K$ and each two Borel subsets $\mathscr{A}, \mathscr{B} \in \mathcal{W}$,

$$\begin{aligned} Q_K\left(\{w \in \mathcal{W} : w|_{(-\infty,0]} \in \mathscr{A},\, w(0) = u \text{ and } w|_{[0,\infty)} \in \mathscr{B}\}\right) \\ = c(u) P_u\big(\langle X_k\rangle_{k\ge 0} \in \mathscr{A}^{\leftarrow} \text{ and } \tau_K^+ = \infty\big) P_u\big(\langle X_k\rangle_{k\ge 0} \in \mathscr{B}\big). \end{aligned}$$

Let $\mathcal{W}_K^* = \pi(W_K)$ be the set of trajectories that visit K.

Theorem 5.3.1 (Sznitman [216] and Teixeira [218]: Existence and uniqueness of the interlacement intensity measure)**.** *Let G be a transient network. There exists a unique σ-finite measure Q^* on $\mathcal{W}^*$ such that for every Borel set $\mathscr{A} \subseteq \mathcal{W}^*$ and every finite $K \subset V$,*

$$Q^*(\mathscr{A} \cap \mathcal{W}_K^*) = Q_K\left(\pi^{-1}(\mathscr{A})\right). \tag{5.3.1}$$

The measure Q^* is referred to as the **interlacement intensity measure**.

Definition 5.3.2. Let Λ denote the Lebesgue measure on $\mathbb{R}$. The **interlacement process** $\mathscr{I}$ on G is defined to be a Poisson point process on $\mathcal{W}^* \times \mathbb{R}$ with intensity measure $Q^* \otimes \Lambda$. For each $t \in \mathbb{R}$, we denote by $\mathscr{I}_t$ the set of $w \in \mathcal{W}^*$ such that $(w,t) \in \mathscr{I}$. We also write $\mathscr{I}_{[a,b]}$ for the intersection of $\mathscr{I}$ with $\mathcal{W}^* \times [a,b]$.

Let $\langle V_n\rangle_{n\ge 0}$ be an exhaustion of an infinite transient network G. The interlacement process on G can be constructed as a limit of Poisson processes on random walk excursions from the boundary vertices ∂_n to itself in the networks G_n^*.

Let N be a Poisson point process on $\mathbb{R}$ with intensity measure $c(\partial_n)\Lambda$. Conditional on N, for every $t \in N$, let W_t be a random walk started at ∂_n and stopped when it first returns to ∂_n, where we consider each W_t to be an element of $\mathcal{W}^*$. We define $\mathscr{I}^n$ to be the point process

$$\mathscr{I}^n := \left\{(W_t, t) : t \in N\right\}.$$

Proposition 5.3.3. *Let G be an infinite transient network and let $\langle V_n\rangle_{n\geq 0}$ be an exhaustion of G. Then the Poisson point processes $\mathscr{I}^n$ converge in distribution to the interlacement process $\mathscr{I}$ as $n\to\infty$.*

A similar construction of the random interlacement process is sketched in [217, §4.5].

Proof. Let $K\subset V$ be finite, and let n be sufficiently large that K is contained in V_n. Define a measure Q_K^n on $\mathcal{W}$ by setting

$$Q_K^n\left(\{w\in\mathcal{W}: w(0)\notin K\}\right)=0$$

and, for each $u\in K$, each $r,m\geq 0$ and each two Borel subsets $\mathscr{A},\mathscr{B}\in\mathcal{W}$,

$$\begin{aligned}Q_K^n&\left(\left\{w\in\mathcal{W}(-r,m): w|_{[-r,0]}\in\mathscr{A}, w(0)=u \text{ and } w|_{[0,m]}\in\mathscr{B}\right\}\right)\\&=c(u)P_u^{G_n^*}\big(\langle X_k\rangle_{k=0}^r\in\mathscr{A}^{\leftarrow}\text{ and }\tau_K^+>\tau_{\partial_n}=r\big)P_u^{G_n^*}\big(\langle X_k\rangle_{k=0}^m\in\mathscr{B}\text{ and }\tau_{\partial_n}=m\big).\end{aligned}\tag{5.3.2}$$

By reversibility of the random walk, the right-hand side of (5.3.2) is equal to

$$\begin{aligned}c(\partial_n)&P_{\partial_n}^{G_n^*}\big(\langle X_k\rangle_{k=0}^n\in\mathscr{A}\text{ and }X_{\tau_K}=u\text{ and }\tau_K=r<\tau_{\partial_n}^+\big)\\&\cdot P_u^{G_n^*}\big(\langle X_k\rangle_{k=0}^m\in\mathscr{B}\text{ and }\tau_{V\setminus V_n}=m\big).\end{aligned}\tag{5.3.3}$$

It follows that $Q_K^n(\mathcal{W})=c(\partial_n)P_{\partial_n}^{G_n^*}(\tau_K<\tau_{\partial_n}^+)$ and that the normalized measure $Q_K^n/Q_K^n(\mathcal{W})$ coincides with the law of a random walk excursion from ∂_n to itself in G_n^* that has been conditioned to hit K and reparameterised so that it first hits K at time 0. Thus, by the splitting property of Poisson processes, $\mathscr{I}^n$ is a Poisson point process on $\mathcal{W}^*\times\mathbb{R}$ with intensity measure $Q^{n*}\otimes\Lambda$, where Q^{n*} satisfies

$$Q^{n*}(\mathscr{A}\cap\mathcal{W}_K^*)=Q_K^n\left(\pi^{-1}(\mathscr{A})\right).$$

We conclude the proof by noting that Q_K^n converges weakly to Q_K as $n\to\infty$. □

5.3.4 Hitting probabilities

Recall that the **capacity** (a.k.a. the **conductance to infinity**) of a finite set of vertices K in a network G is defined to be

$$\mathrm{Cap}(K)=\sum_{v\in K}c(v)P_v(\tau_K^+=\infty),$$

and observe that $Q_K(\mathcal{W})=\mathrm{Cap}(K)$ for every finite set of vertices K. If K is infinite, we define $\mathrm{Cap}(K)=\lim_{n\to\infty}\mathrm{Cap}(K_n)$, where K_n is any increasing sequence of finite sets of vertices with $\bigcup K_n=K$. We say that a set K of vertices is **hit** by $\mathscr{I}_{[a,b]}$ if there exists $(W,t)\in\mathscr{I}_{[a,b]}$ such that W hits K. By the definition of $\mathscr{I}$, we have that

$$\mathbb{P}(K\text{ hit by }\mathscr{I}_{[a,b]})=1-\exp\big(-(b-a)Q_K(\mathcal{W})\big)=1-\exp\big(-(b-a)\,\mathrm{Cap}(K)\big)$$

for each finite set K. This formula extends to infinite sets by taking limits over exhaustions: If $K \subseteq V$ is infinite, let K_n be an exhaustion of K by finite sets. Then

$$\begin{aligned}\mathbb{P}(K \text{ hit by } \mathscr{I}_{[a,b]}) &= \lim_{n\to\infty} \mathbb{P}(K_n \text{ hit by } \mathscr{I}_{[a,b]}) \\ &= 1 - \lim_{n\to\infty} \exp\left(-(b-a)\operatorname{Cap}(K_n)\right) = 1 - \exp\left(-(b-a)\operatorname{Cap}(K)\right).\end{aligned}$$

Similarly, the expected number of trajectories in $\mathscr{I}_{[a,b]}$ that hit K is equal to $(b-a)\operatorname{Cap}(K)$. We apply the above formulas to deduce the following simple 0-1 law.

Lemma 5.3.4. *Let G be a transient network, let $\mathscr{I}$ be the interlacement process on G. Then for all $a < b \in \mathbb{R}$ and every set of vertices $K \subseteq V$, we have*

$$\mathbb{P}\Big(K \textit{ is hit by infinitely many trajectories in } \mathscr{I}_{[a,b]}\Big) = \mathbb{1}\left(\operatorname{Cap}(K) = \infty\right).$$

Proof. If $\operatorname{Cap}(K)$ is finite then the expected number of trajectories in $\mathscr{I}_{[a,b]}$ that hit K is finite, so that the number of trajectories in $\mathscr{I}_{[a,b]}$ that hit K is finite a.s. Conversely, if $\operatorname{Cap}(K)$ is infinite, then there is a trajectory in $\mathscr{I}_{[b-2^{-n},b-2^{-n-1}]}$ that hits K a.s. for every $n \geq 1$ a.s. Since $b-2^{-n} \geq a$ for all but finitely many n, it follows that $\mathscr{I}_{[a,b]}$ hits K infinitely often a.s. □

We next prove that any set that has a positive probability to be hit infinitely often by any single trajectory will in fact be hit by infinitely many trajectories. Recall the **method of random paths** [191, Theorem 10.1]: If G is an infinite network, A is a finite subset of G and Γ is a random infinite simple path in G starting at A, then

$$\operatorname{Cap}(A)^{-1} \leq \sum_{e\in E} \mathbb{P}(e \in \Gamma)^2.$$

In particular, if the sum on the right hand side is finite for some random infinite simple path Γ starting at A, then the capacity of A is positive and G is therefore transient. Moreover, for every finite set A in a transient network, there exists a random infinite simple path Γ starting in A such that

$$\operatorname{Cap}(A)^{-1} = \sum_{e\in E} \mathbb{P}(e \in \Gamma)^2.$$

The following lemma is presumably well-known, but we were unable to find a reference.

Lemma 5.3.5. *Let G be a transient network and let $K \subseteq V$. If the random walk on G hits K infinitely often with positive probability, then* $\operatorname{Cap}(K) = \infty$.

Proof. Let c be the conductance function of G. First suppose that K is hit infinitely often with probability one. Let $X = \langle X_i \rangle_{i\geq 0}$ be a random walk on G started at a vertex of K, and let N_i be the ith time that X visits K. Define

$$c_K(u,v) = c(u)P_u(X_{N_1} = v)$$

for all $u, v \in K$, so that X_{N_i} is the random walk on the (non-locally finite) network $H := ((K, K^2), c_K)$. Note that if A is a finite subset of K, then the capacity of A considered as a set of vertices in H is the same as the capacity of A considered as a set of vertices in G. That is,

$$\mathrm{Cap}(A) = \sum_{v \in A} c(v) P_v(\langle X_i \rangle_{i \geq 0} \text{ returns to } A) = \sum_{v \in A} c(v) P_v(\langle X_{N_i} \rangle_{i \geq 0} \text{ returns to } A).$$

Let $\langle K_n \rangle_{n \geq 1}$ be an increasing sequence of finite sets with $\bigcup_{n \geq 1} K_n = K$, and let Γ be a random infinite simple path in H starting at K_1 such that

$$\frac{1}{2} \sum_{u,v \in V} \mathbb{P}(\{u, v\} \in \Gamma)^2 = \mathrm{Cap}(K_1)^{-1} < \infty.$$

For each $n \geq 2$, let Γ_n be the subpath of Γ beginning at the last time Γ visits K_n. Then, by the monotone convergence theorem,

$$\mathrm{Cap}(K) = \lim_{n \to \infty} \mathrm{Cap}(K_n) \geq \lim_{n \to \infty} \left(\frac{1}{2} \sum_{u,v \in V} \mathbb{P}(\{u, v\} \in \Gamma_n)^2 \right)^{-1} = \infty.$$

This concludes the proof in this case.

Now suppose that K is hit infinitely often by the random walk on G with positive probability, and, for each vertex u of G, let $h(u)$ be the probability that a random walk on G started at u hits K infinitely often. We have that, for each two vertices u and v of G,

$$P_u(X_1 = v \mid X \text{ hits } K \text{ infinitely often}) = \frac{P_u(X_1 = v) P_v(X \text{ hits } K \text{ infinitely often})}{P_u(X \text{ hits } K \text{ infinitely often})} = \frac{h(v) c(u, v)}{h(u) c(u)}.$$

It follows by an elementary calculation that the random walk on G conditioned to hit K infinitely often is reversible, and is equal to the random walk on the network $(G, \hat{c})$, where the conductances $\hat{c}$ are defined by

$$\hat{c}(u, v) = c(u, v) h(u) h(v), \quad \hat{c}(u) = c(u) h(u)^2.$$

This is an example of Doob's h-transform. Since $h \leq 1$, Rayleigh monotonicity implies that the capacity of K with respect to the h-transformed conductances $\hat{c}$ is less than the capacity of K with respect to the original conductances c. Thus, we may conclude the proof by applying the argument of the previous paragraph to the h-transformed network. $\square$

The following lemma is an immediate consequence of Lemma 5.3.4 and Lemma 5.3.5.

Lemma 5.3.6. *Let G be a transient network, let $\mathscr{I}$ be the interlacement process on G. Then for* *all $a < b \in \mathbb{R}$ and every set of vertices $K \subseteq V$, we have*

$$\mathbb{P}\Big(\textit{infinitely many vertices of } K \textit{ are hit by } \mathscr{I}_{[a,b]}\Big) = \\ \mathbb{P}\Big(K \textit{ is hit by infinitely many trajectories in } \mathscr{I}_{[a,b]}\Big) = \mathbb{1}\left(\mathrm{Cap}(K) = \infty\right).$$

5.4 Interlacement Aldous-Broder

In this section we describe the Interlacement Aldous-Broder algorithm and investigate its basic properties. Let G be an infinite transient network. For each set $A \subseteq \mathcal{W}^* \times \mathbb{R}$, and each vertex $v \in V$, define

$$\tau_t(A, v) := \inf\left\{s \geq t : \exists (W, s) \in A \text{ such that } W \text{ hits } v\right\}.$$

Let $\mathscr{I}$ be the interlacement process on G and write $\tau_t(v) = \tau_t(\mathscr{I}, v)$. Let $\mathcal{A}$ be the set of subsets A of $\mathcal{W}^* \times \mathbb{R}$ that satisfy the property that for every for every vertex $v \in V$ and every $t \in \mathbb{R}$ there exists a unique trajectory $W_{\tau_t(A,v)}$ such that $(W_{\tau_t(A,v)}, \tau_t(A, v)) \in A$ and $W_{\tau_t(A,v)}$ hits v. It is clear that $\mathscr{I} \in \mathcal{A}$ a.s. Define $e_t(v) = e_t(\mathscr{I}, v)$ to be the oriented edge pointing into v that is traversed by the trajectory $W_{\tau_t(v)}$ as it enters v for the first time. For each $t \in \mathbb{R}$ and $T \in (t, \infty]$, we define the set $\mathsf{AB}_t^T(\mathscr{I}) \subseteq E$ by

$$\mathsf{AB}_t^T(\mathscr{I}) := \left\{-e_t(v) : v \in V,\, \tau_t(v) \leq T\right\}. \tag{5.4.1}$$

We write $\mathsf{AB}_t(\mathscr{I}) = \mathsf{AB}_t^\infty(\mathscr{I})$. We define $\mathsf{AB}_t^T(A)$ similarly for all $A \in \mathcal{A}$. Let $\langle V_n \rangle_{n \geq 0}$ be an exhaustion of G, and for each $n \geq 0$ let $\mathscr{I}^n$ be defined as in Section 5.3.3. Since the process $\mathscr{I}^n$ is just a decomposition of a single random walk trajectory into excursions, Theorem 5.1.1 follows immediately from the following lemma together with the correctness of the classical Aldous-Broder algorithm.

Lemma 5.4.1. *Let G be a transient network with exhaustion $\langle V_n \rangle_{n \geq 0}$ and let $\mathscr{I}$ be the interlacement process on G. Then $\mathsf{AB}_t^T(\mathscr{I}^n)$ converges weakly to $\mathsf{AB}_t^T(\mathscr{I})$ for each $t \in \mathbb{R}$ and $T \in [t, \infty]$.*

Proof. Let $E^{\rightarrow}$ be the set of oriented edges of G, let S be a finite subset of $E^{\rightarrow}$, and consider the set

$$C_t^T(S) := \Big\{A \subseteq \mathcal{W}^* \times \mathbb{R} : S \subseteq \mathsf{AB}_t^T(A)\Big\}.$$

Let $\partial C_t^T(S)$ denote the topological boundary of $C_t^T(S)$. Observe that

$$\partial C_t^T(S) \subseteq \Big\{A \subseteq \mathcal{A} : \tau_t(e^-) \in \{t, T\} \text{ for some } e \in S\Big\} \cup (\mathcal{W}^* \times \mathbb{R} \setminus \mathcal{A}).$$

It follows that $\mathbb{P}(\mathscr{I} \in \partial C_t^T(S)) = 0$. The Portmanteau Theorem [155, Theorem 13.16] therefore implies that

$$\mathbb{P}\left(S \subseteq \mathsf{AB}_t^T(\mathscr{I})\right) = \mathbb{P}\left(\mathscr{I} \in C_t^T(S)\right) = \lim_{n\to\infty} \mathbb{P}\left(\mathscr{I}^n \in C_t^T(S)\right) = \lim_{n\to\infty} \mathbb{P}\left(S \subseteq \mathsf{AB}_t^T(\mathscr{I}^n)\right)$$

for every finite set $S \subseteq E^{\rightarrow}$. □

We next establish the basic properties of the process $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}} = \langle \mathsf{AB}_t^\infty(\mathscr{I})\rangle_{t\in\mathbb{R}}$. We first recall that a process $\langle X_t\rangle_{t\in\mathbb{R}}$ is said to be **ergodic** if $\mathbb{P}(\langle X_t\rangle_{t\in\mathbb{R}} \in \mathscr{A}) \in \{0,1\}$ whenever $\mathscr{A}$ is an event that is shift invariant in the sense that $\langle X_t\rangle_{t\in\mathbb{R}} \in \mathscr{A}$ implies that $\langle X_{t+s}\rangle_{t\in\mathbb{R}} \in \mathscr{A}$ for every $s \in \mathbb{R}$. The process $\langle X_t\rangle_{t\in\mathbb{R}}$ is said to be **mixing** if for every two events $\mathscr{A}$ and $\mathscr{B}$,

$$\mathbb{P}\left(\langle X_t\rangle_{t\in\mathbb{R}} \in \mathscr{A} \text{ and } \langle X_{t+s}\rangle_{t\in\mathbb{R}} \in \mathscr{B}\right) \xrightarrow[s\to\infty]{} \mathbb{P}\left(\langle X_t\rangle_{t\in\mathbb{R}} \in \mathscr{A}\right)\mathbb{P}\left(\langle X_t\rangle_{t\in\mathbb{R}} \in \mathscr{B}\right).$$

Every mixing process is clearly ergodic, but the converse need not hold in general [155, §20.5].

Proposition 5.4.2. *Let G be a transient network and let $\mathscr{I}$ be the interlacement process on G. Then $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}} = \langle \mathsf{AB}_t^\infty(\mathscr{I})\rangle_{t\in\mathbb{R}}$ is an ergodic, mixing, Markov process.*

Proof. The fact that $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}}$ is a Markov process follows from the following identity, which immediately implies that $\langle \mathfrak{F}_s \rangle_{s\le t}$ and $\langle \mathfrak{F}_s \rangle_{s\ge t}$ are conditionally independent given $\mathfrak{F}_t$ for each $t \in \mathbb{R}$: whenever $t \in \mathbb{R}$ and $s \geq 0$,

$$\mathsf{AB}_{t-s}(\mathscr{I}) = \mathsf{AB}_{t-s}^t(\mathscr{I}) \cup \{e \in \mathsf{AB}_t(\mathscr{I}) : e^- \text{ is not hit by } \mathscr{I}_{[t-s,t)}\}. \tag{5.4.2}$$

We now prove that $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}}$ is mixing. Let $a_1, a_2, \ldots, a_n$ and $b_1, b_2, \ldots, b_m$ be two increasing sequences of real numbers, and let $A_1, A_2, \ldots, A_n$ and $B_1, B_2, \ldots B_m$ be finite subsets of $E^{\rightarrow}$. Let K denote the set of endpoints of edges in the union $\bigcup_{i=1}^n A_i$. Let s be sufficiently large that $b_1+s \geq a_n$. Let $\mathscr{A} = \{A_i \subseteq \mathsf{AB}_{a_i}^\infty(\mathscr{I}) \text{ for all } 1 \leq i \leq n\}$, $\mathscr{B}_s = \{B_i \subseteq \mathsf{AB}_{b_i+s}^\infty(\mathscr{I}) \text{ for all } 1 \leq i \leq m\}$, and $\mathscr{A}_s' = \{A_i \subseteq \mathsf{AB}_{a_i}^{b_1+s}(\mathscr{I}) \text{ for all } 1 \leq i \leq n\} \subseteq \mathscr{A}$. The events $\mathscr{A}_s'$ and $\mathscr{B}_s$ are independent and

$$\mathbb{P}(\mathscr{A} \setminus \mathscr{A}_s') \leq \mathbb{P}(\mathscr{I}_{[a_n, b_1+s]} \text{ does not hit some vertex in } K) \leq \sum_{v\in K} \exp\left(-(b_1+s-a_n)\mathrm{Cap}(v)\right).$$

We deduce that

$$\begin{aligned}
|\mathbb{P}(\mathscr{A} \cap \mathscr{B}_s) - \mathbb{P}(\mathscr{A})\mathbb{P}(\mathscr{B}_s)| &\leq |\mathbb{P}(\mathscr{A} \cap \mathscr{B}_s) - \mathbb{P}(\mathscr{A}_s' \cap \mathscr{B}_s)| + |\mathbb{P}(\mathscr{A}_s' \cap \mathscr{B}_s) - \mathbb{P}(\mathscr{A})\mathbb{P}(\mathscr{B}_s)| \\
&= |\mathbb{P}(\mathscr{A} \cap \mathscr{B}_s) - \mathbb{P}(\mathscr{A}_s' \cap \mathscr{B}_s)| + |\mathbb{P}(\mathscr{A}_s')\mathbb{P}(\mathscr{B}_s) - \mathbb{P}(\mathscr{A})\mathbb{P}(\mathscr{B}_s)| \\
&\leq 2\sum_{v\in K} \exp\left(-(b_1+s-a_n)\mathrm{Cap}(v)\right) \xrightarrow[s\to\infty]{} 0.
\end{aligned}$$

Since events of the form $\{A_i \subseteq \mathsf{AB}_{a_i}^\infty(\mathscr{I}) \text{ for all } 1 \leq i \leq n\}$ generate the Borel σ-algebra on the space of $(E^{\rightarrow})^{\{0,1\}}$-valued processes, it follows that $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}}$ is mixing and therefore ergodic. □

5.5 Proof of Theorems 5.2.1, 5.2.2 and 5.2.4

Recall that an **end** of a tree is an equivalence class of infinite simple paths in the tree, where two infinite simple paths are equivalent if their traces have finite symmetric difference. Similarly, if $\mathfrak{F}$ is a spanning forest of a network G, we define an end of $\mathfrak{F}$ to be an equivalence class of infinite simple paths in G that eventually only use edges of $\mathfrak{F}$, where, again, two infinite simple paths are equivalent if their traces have finite symmetric difference. If an infinite tree T is oriented so that every vertex of the tree has exactly one oriented edge emanating from it, then there is exactly one end ξ of T for which the paths representing ξ eventually follow the orientation of T. We call this end the **primary end** of T. We call ends of T that are not the primary end **excessive ends** of T. Note that if ξ is an excessive end of T and γ is a simple path representing ξ, then all but finitely many of the edges traversed by the path γ are traversed in the opposite direction to their orientation in T.

If $\mathscr{I}$ is the interlacement process on a transient network G, we write $\mathcal{I}_{[a,b]}$ for the set of vertices of G hit by $\mathscr{I}_{[a,b]}$.

The proofs of Theorems 5.2.1 and 5.2.4 both rely on the following criterion. See Figure 5.1 for an illustration of the proof.

Lemma 5.5.1. *Let G be a transient network, let $\mathscr{I}$ be the interlacement process on G, and let $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}} = \langle \mathsf{AB}_t(\mathscr{I}) \rangle_{t\in\mathbb{R}}$. If the connected component containing v of $\operatorname{past}_{\mathfrak{F}_0}(v) \setminus \mathcal{I}_{[-\varepsilon,0]}$ is finite a.s. for every vertex v of G and every $\varepsilon > 0$, then every component of $\mathfrak{F}_0$ is one-ended a.s.*

Proof. If u is in the past of v in $\mathfrak{F}_{-\varepsilon}$, then $\tau_{-\varepsilon}(u) \geq \tau_{-\varepsilon}(v)$. Thus, on the event that v is not hit by $\mathscr{I}_{[-\varepsilon,0]}$, the past of v in $\mathfrak{F}_{-\varepsilon}$ is equal to the component containing v in the subgraph of $\operatorname{past}_{\mathfrak{F}_0}(v)$ induced by the complement of $\mathcal{I}_{[-\varepsilon,0]}$ (see Figure 5.1). By assumption, the connected component containing v in this subgraph is finite a.s., and so, by stationarity,

$$\mathbb{P}(\operatorname{past}_{\mathfrak{F}_0}(v) \text{ is infinite}) = \mathbb{P}(\operatorname{past}_{\mathfrak{F}_{-\varepsilon}}(v) \text{ is infinite}) \leq \mathbb{P}(v \in \mathcal{I}_{[-\varepsilon,0]}) = 1 - e^{-\varepsilon \operatorname{Cap}(v)} \xrightarrow[\varepsilon\to 0]{} 0.$$

Since v was arbitrary, we deduce that every component of $\mathfrak{F}_0$ is one-ended a.s. □

The proof of Theorem 5.2.1 requires both of the following theorems.

Theorem 5.5.2 (Lyons, Morris and Schramm [170]; Lyons and Peres [173, Theorem 6.41])**.** *Let G be a network satisfying an anchored $f(t)$-isoperimetric inequality, where f is an increasing function such that $f(t) \leq t$, $f(2t) \leq \alpha f(t)$ for some constant α, and $\int_1^\infty f(t)^{-2}\mathrm{d}t < \infty$. Then for every vertex v of G there exists a positive constant c_v such that, for every connected set K containing v,*

$$\operatorname{Cap}(K) \geq \frac{c_v^2}{4\alpha^2}\left(\int_{|K|}^\infty \frac{1}{f(t)^2}\right)^{-1}.$$

In particular, G is transient.

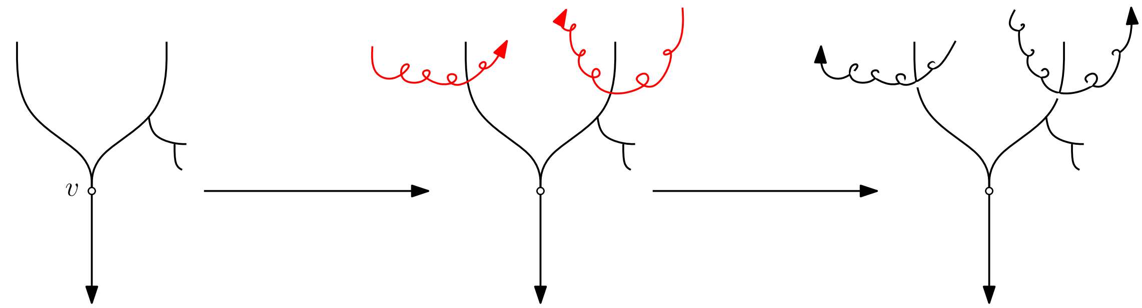

Figure 5.1: A schematic illustration of the proof of Lemma 5.5.1. Left: a tree in $\mathfrak{F}_0$ with two excessive ends in the past of the vertex v. Centre: each of the excessive ends is hit by a trajectory of $\mathscr{I}_{[-\varepsilon,0]}$ (red), but v is not hit by such a trajectory. Right: the resulting forest $\mathfrak{F}_{-\varepsilon}$.

Theorem 5.5.3 (Morris [185, Theorem 9]: WUSF components are recurrent)**.** *Let G be an infinite network with edge conductances bounded above. Then every component of the wired uniform spanning forest of G is recurrent a.s.*

An equivalent statement of Theorem 5.5.3 is the following.

Lemma 5.5.4. *Let G be an infinite network with $\inf_e c(e) > 0$. Then every component of the wired uniform spanning forest of G is recurrent when given unit conductances a.s.*

Proof. Form a network $\hat{G}$ by replacing each edge e of G with $\lceil c(e) \rceil$ parallel edges each with conductance $c(e)/\lceil c(e) \rceil$. It follows immediately from the definition of the UST of a network that the WUSF $\hat{\mathfrak{F}}$ of $\hat{G}$ may be coupled with the WUSF $\mathfrak{F}$ of G so that an edge e of G is contained in $\mathfrak{F}$ if and only if one of the edges corresponding to e in $\hat{G}$ is contained in $\hat{\mathfrak{F}}$. Since $\hat{G}$ has edge conductances bounded above, Theorem 5.5.3 implies that every component of $\hat{\mathfrak{F}}$ is recurrent a.s. Since the edge conductances of $\hat{G}$ are bounded away from zero, it follows by Rayleigh monotonicity that every component of $\hat{\mathfrak{F}}$ is recurrent a.s. when given unit conductances, and consequently that the same is true of $\mathfrak{F}$. $\square$

Proof of Theorem 5.2.1. By Theorem 5.5.2, G is transient. Let $\mathscr{I}$ be the interlacement process on G. Let v be a fixed vertex of G. Then for each vertex u of G contained in the same component of $\mathfrak{F} = \mathsf{AB}_0(\mathscr{I})$ as v, the conditional probability given $\mathfrak{F}$ that u is connected to v in $\mathfrak{F} \setminus \mathcal{I}_{[-\varepsilon,0]}$ is equal to $\exp(-\varepsilon \operatorname{Cap}(\gamma_{u,v}))$, where $\gamma_{u,v}$ is the trace of the path connecting u to v in $\mathfrak{F}$. Write $d_{\mathfrak{F}}(u,v)$ for the graph distance in $\mathfrak{F}$ between two vertices u and v in G. Since the conductances of G are bounded below by some positive constant δ, we have that $|\gamma_{u,v}| \geq \delta d_{\mathfrak{F}}(u,v)$ and so, by Theorem 5.5.2,

$$\mathbb{P}\left(u \text{ connected to } v \text{ in } \mathfrak{F} \setminus \mathcal{I}_{[-\varepsilon,0]} \mid \mathfrak{F}\right) \\ \leq \mathbb{1}(u \text{ connected to } v \text{ in } \mathfrak{F}) \exp\left(-\varepsilon \frac{c_v^2}{4\alpha^2} \left(\int_{\delta d_{\mathfrak{F}}(u,v)}^{\infty} \frac{1}{f(t)^2}\right)^{-1}\right). \quad (5.5.1)$$

Lemma 5.5.5. *Let T be an infinite tree, let v be a vertex of T and let ω be a random subgraph of T. For each vertex u of G, let $\|u\|$ denote the distance between u and v in T, and suppose that there exists a function $p : \mathbb{N} \to [0,1]$ such that*

$$\mathbb{P}(u \text{ is connected to } v \text{ in } \omega) \leq p\left(\|u\|\right)$$

for every vertex u in T and

$$\sum_{n\geq 1} p(n) < \infty.$$

Then

$$\mathbb{P}(\text{The component of } v \text{ in } \omega \text{ is infinite}) \leq \operatorname{Cap}(v) \sum_{n\geq 1} p(n).$$

In particular, if T is recurrent then the component containing v in ω is finite a.s.

Proof. Suppose that the component containing v in ω is infinite with positive probability; the inequality holds trivially otherwise. Denote this event $\mathscr{A}$. Fix a drawing of T in the plane rooted at v. On the event $\mathscr{A}$, let $\Gamma = \Gamma(\omega)$ be the leftmost simple path from v to infinity in ω. Observe that

$$\mathbb{P}(u \in \Gamma \mid \mathscr{A}) \leq \frac{\mathbb{P}(u \text{ is connected to } v \text{ in } \omega)}{\mathbb{P}(\mathscr{A})} \leq \frac{p(n)}{\mathbb{P}(\mathscr{A})} \quad \text{and} \quad \sum_{\|u\|=n} \mathbb{P}(u \in \Gamma \mid \mathscr{A}) = 1,$$

so that

$$\sum_{\|u\|=n} \mathbb{P}(u \in \Gamma \mid \mathscr{A})^2 \leq \frac{p(n)}{\mathbb{P}(\mathscr{A})}.$$

and hence

$$\sum_{e\in E} \mathbb{P}(e \in \Gamma \mid \mathscr{A})^2 = \sum_{u\in V\setminus\{v\}} \mathbb{P}(u \in \Gamma \mid \mathscr{A})^2 = \sum_{n\geq 1}\sum_{\|u\|=n} \mathbb{P}(u \in \Gamma \mid \mathscr{A})^2 \leq \frac{1}{\mathbb{P}(\mathscr{A})}\sum_{n\geq 1} p(n).$$

Applying the method of random paths (taking our measure on random paths to be the conditional distribution of Γ given $\mathscr{A}$), we deduce that

$$\operatorname{Cap}(v) \geq \left(\frac{1}{\mathbb{P}(\mathscr{A})}\sum_{n\geq 1} p(n)\right)^{-1},$$

which rearranges to give the desired inequality. $\square$

Comparing sums with integrals, hypothesis (2) of Theorem 5.2.1 implies that

$$\sum_{n\geq 1} \exp\left(-\varepsilon \frac{c_v^2}{4\alpha^2}\left(\int_{\delta n}^{\infty} \frac{1}{f(t)^2}\right)^{-1}\right) < \infty$$

for every $\varepsilon > 0$, and we deduce from eq. (5.5.1), Lemma 5.5.5 and Lemma 5.5.4 that the component containing v in $\mathfrak{F} \setminus \mathcal{I}_{[-\varepsilon,0]}$ is finite a.s. Since v was arbitrary, every component of $\mathfrak{F} \setminus \mathcal{I}_{[-\varepsilon,0]}$ is finite a.s. We conclude by applying Lemma 5.5.1. $\square$

5.5.1 Unimodular random rooted graphs

Proof of Theorem 5.2.4. Let (G,ρ) be a transient unimodular random rooted network, let $\mathscr{I}$ be the interlacement process on G and let $\mathfrak{F} = \mathsf{AB}_0(\mathscr{I})$. It is known [7, Theorem 6.2, Proposition 7.1] that every component of $\mathfrak{F}$ has at most two ends a.s. Suppose for contradiction that $\mathfrak{F}$ contains a two-ended component with positive probability. The **trunk** of a two-ended component of $\mathfrak{F}$ is defined to be the unique doubly infinite simple path that is contained in the component. Define $\mathrm{trunk}(\mathfrak{F})$ to be the set of vertices of G that are contained in the trunk of some two-ended component of $\mathfrak{F}$. For each vertex v of G, let $e(v)$ be the unique oriented edge of G emanating from v that is contained in $\mathfrak{F}$. For each vertex $v \in \mathrm{trunk}(\mathfrak{F})$, let $s(v)$ be the unique vertex in $\mathrm{trunk}(\mathfrak{F})$ that has $e(s(v))^+ = v$, and let $s^n(v)$ be defined recursively for $n \geq 0$ by $s^0(v) = v$, $s^{n+1}(v) = s(s^n(v))$.

Let $\varepsilon > 0$. We claim that $s^n(\rho) \in \mathcal{I}_{[-\varepsilon,0]}$ for infinitely many n a.s. on the event that $\rho \in \mathrm{trunk}(\mathfrak{F})$. Let $k \geq 1$ and define the mass transport

$$f_k(G,u,v,\mathfrak{F},\mathcal{I}_{[-\varepsilon,0]}) = \mathbb{1}\left(\begin{array}{l} u \text{ is in the trunk of its component in } \mathfrak{F}, \\ v = s^k(u) \text{ and } \mathcal{I}_{[-\varepsilon,0]} \cap \{s^n(v) : n \geq 0\} \neq \emptyset \end{array}\right).$$

Applying the mass-transport principle to f_k, we deduce that

$$\mathbb{P}\left(\mathcal{I}_{[-\varepsilon,0]} \cap \{s^n(\rho) : n \geq 0\} \neq \emptyset \mid \rho \in \mathrm{trunk}\right) = \mathbb{P}\left(\mathcal{I}_{[-\varepsilon,0]} \cap \{s^n(\rho) : n \geq k\} \neq \emptyset \mid \rho \in \mathrm{trunk}\right)$$

for all $k \geq 0$. (Here we are using the fact that $(G,\rho,\mathfrak{F},\mathcal{I}_{[-\varepsilon,0]})$ is a unimodular rooted marked graph, see [7] for appropriate definitions.) By taking the limit as $k \to \infty$, we deduce that

$$\mathbb{P}\left(s^n(\rho) \in \mathcal{I}_{[-\varepsilon,0]} \text{ for infinitely many } n \mid \rho \in \mathrm{trunk}\right) \\ = \mathbb{P}\left(\mathcal{I}_{[-\varepsilon,0]} \cap \{s^n(\rho) : n \geq 0\} \neq \emptyset \mid \rho \in \mathrm{trunk}\right).$$

It follows that $s^n(\rho) \in \mathcal{I}_{[-\varepsilon,0]}$ for infinitely many n almost surely on the event that $\rho \in \mathrm{trunk}(\mathfrak{F}) \cap \mathcal{I}_{[-\varepsilon,0]}$. Since $\rho \in \mathrm{trunk}(\mathfrak{F}) \cap \mathcal{I}_{[-\varepsilon,0]}$ with positive probability conditional on $\mathfrak{F}$ and the event that $\rho \in \mathrm{trunk}(\mathfrak{F})$, it follows from Lemma 5.3.6 that infinitely many vertices of $\{s^n(\rho) : n \geq 1\}$ are hit by $\mathcal{I}_{[-\varepsilon,0]}$ a.s. on the event that $\rho \in \mathrm{trunk}(\mathfrak{F})$.

It follows from [7, Lemma 2.3] that for every vertex $v \in \mathrm{trunk}(\mathfrak{F})$, $s^n(v) \in \mathcal{I}_{[-\varepsilon,0]}$ for infinitely many n a.s., and consequently that the component containing v in $\mathrm{past}_{\mathfrak{F}}(v) \setminus \mathcal{I}_{[-\varepsilon,0]}$ is finite for every vertex v of G a.s. We conclude by applying Lemma 5.5.1. $\square$

5.5.2 Excessive ends

Proof of Theorem 5.2.2. We may assume that G is transient: if not, the WUSF of G is connected, the number of excessive ends of G is tail measurable, and the claim follows by tail-triviality of the WUSF [44]. Let $\mathscr{I}$ be the interlacement process on G and let $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}} = \langle \mathsf{AB}_t(\mathscr{I})\rangle_{t\in\mathbb{R}}$. The event that $\mathfrak{F}_0$ has uncountably many ends is tail measurable, and hence has probability either 0 or 1, again by tail-triviality of the WUSF. If the number of ends of $\mathfrak{F}_0$ is uncountable a.s., then $\mathfrak{F}_0$ must also have uncountably many excessive ends a.s., since the number of components of $\mathfrak{F}_0$ is countable. Thus, it suffices to consider the case that $\mathfrak{F}_0$ has countably many ends a.s.

For each $t \in \mathbb{R}$, we call an excessive end ξ of $\mathfrak{F}_t$ **indestructible** if $\mathrm{Cap}\left(\{\gamma_i : i \geq 0\}\right)$ is finite for some (and hence every) simple path $\langle \gamma_i \rangle_{i\geq 0}$ in G representing ξ, and **destructible** otherwise. Given a simple path $\gamma = \langle \gamma_i \rangle_{i\geq 0}$, write $\gamma_{i,i+1}$ for the oriented edge that is traversed by γ as it moves from γ_i to γ_{i+1}, and let $\gamma_{i,i-1} = -\gamma_{i-1,i}$. Observe, as we did at the beginning of Section 5.5, that a simple path $\langle \gamma_i \rangle_{i\geq 0}$ in G represents an excessive end of $\mathfrak{F}_t$ if and only if $e_t(\gamma_i, \mathscr{I}) = \gamma_{i,i-1}$ for all sufficiently large values of i (equivalently, if and only if the reversed oriented edges $-\gamma_{i,i+1}$ are contained in $\mathfrak{F}_t$ for all sufficiently large values of i). Since $\mathfrak{F}_0$ has countably many ends a.s. by assumption, it follows from Lemma 5.3.4 that for every destructible end ξ of $\mathfrak{F}_0$ and every infinite simple path $\langle \gamma_i \rangle_{i\geq 0}$ in G representing ξ, the trace $\{\gamma_i : i \geq 0\}$ of γ is hit by $\mathscr{I}_{[-\varepsilon,0]}$ infinitely often a.s. for every $\varepsilon > 0$. Recall from the proof of Lemma 5.5.1 that, on the event that $v \notin \mathcal{I}_{[-\varepsilon,0]}$, the past of v in $\mathfrak{F}_{-\varepsilon}$ is contained in subgraph of $\mathrm{past}_{\mathfrak{F}_0}(v)$ induced by the complement of $\mathcal{I}_{[-\varepsilon,0]}$. It follows that v a.s. does not have any destructible ends in its past in $\mathfrak{F}_{-\varepsilon}$ a.s. on the event that $v \notin \mathcal{I}_{[-\varepsilon,0]}$, and so, by stationarity,

$$\mathbb{P}(v \text{ has a destructible end in its past in } \mathfrak{F}_0) \leq \mathbb{P}(v \in \mathcal{I}_{[-\varepsilon,0]}) \xrightarrow[\varepsilon\to 0]{} 0.$$

Since the vertex v was arbitrary, we deduce that $\mathfrak{F}_0$ does not contain any destructible excessive ends a.s.

Since every excessive end of $\mathfrak{F}_0$ is indestructible a.s., it follows from Lemma 5.3.6 that for every excessive end ξ of $\mathfrak{F}_0$ and every path $\langle v_i \rangle$ representing ξ, only finitely many of the vertices v_i are hit by $\mathscr{I}_{[t,0]}$ a.s. for every $t \leq 0$. Since $\mathfrak{F}_0$ has at most countably many excessive ends a.s., we deduce that every path $\langle v_i \rangle_{i\geq 0}$ that represents an excessive end of $\mathfrak{F}_0$ also represents an excessive end of $\mathfrak{F}_t$ for every $t \leq 0$. In particular, the cardinality of the set of excessive ends of $\mathfrak{F}_t$ is at least the cardinality of the set of excessive ends of $\mathfrak{F}_0$ a.s. for every $t \leq 0$. Since, by Proposition 5.4.2, $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}}$ is stationary and ergodic, we deduce that the cardinality of the set of excessive ends of $\mathfrak{F}_0$ is a.s. equal to some constant. □

Proof of Corollary 5.2.3. Let G'' be the network that has all the edges of both G and G'. By symmetry, it suffices to show that the wired uniform spanning forests of G and G'' have the same number of excessive ends a.s. Let $\mathfrak{F}$ and $\mathfrak{F}''$ be samples of the WUSFs of G and G'' respectively, and let A be the set of edges of G'' that are not edges of G. Since A is finite and G is connected, the event $\mathscr{A} = \{A \cap \mathfrak{F}'' = \emptyset\}$ has positive probability (this implication is easily proven in several ways, e.g. using either Wilson's algorithm, the Aldous-Broder algorithm, or the Transfer Current Theorem [62]). The spatial Markov property of the WUSF implies that the conditional distribution of $\mathfrak{F}''$ given $\mathscr{A}$ is equal to the distribution of $\mathfrak{F}$, and in particular the conditional distribution of the number of excessive ends of $\mathfrak{F}''$ given $\mathscr{A}$ has the same distribution as the number of excessive ends of $\mathfrak{F}$. The claim now follows from Theorem 5.2.2. □

5.6 Ends and rough isometries

Recall that a **rough isometry** from a graph $G = (V, E)$ to a graph $G' = (V', E')$ is a function $\phi : V \to V'$ such that, letting d_G and $d_{G'}$ denote the graph distances on V and V', there exist positive constants α and β such that the following conditions are satisfied:

1. **(ϕ roughly preserves distances.)** For every pair of vertices $u, v \in V$,

$$\alpha^{-1} d_G(u, v) - \beta \leq d_{G'}(\phi(u), \phi(v)) \leq \alpha d_G(u, v) + \beta.$$

2. **(ϕ is almost surjective.)** For every vertex $v' \in V'$, there exists a vertex $v \in V$ such that $d_G(\phi(v), v') \leq \beta$.

For background on rough isometries, see [173, §2.6]. The final result of this paper answers negatively Question 7.6 of Lyons, Morris and Schramm [170], which asked whether the property of having one-ended WUSF components is preserved under rough isometry of graphs.

Theorem 5.6.1. *There exist two rough-isometric, bounded degree graphs G and G' such that every component of the wired uniform spanning forest of G has one-end a.s., but the wired uniform spanning forest of G' contains a component with uncountably many ends a.s.*

The proof of Theorem 5.6.1 uses Wilson's algorithm rooted at infinity. We refer the reader to [173, Proposition 10.1] for an exposition of this algorithm. The description as a branching process of the past of the WUSF of a regular trees with height-dependent exponential edge stretching is adapted from [44, §11], and first appeared in the work of Häggström [111].

Proof of Theorem 5.6.1. Let $T = (V, E)$ be a 3-regular tree with root ρ. We write $\|u\|$ for the distance between $u \in V$ and ρ. For each positive integer k, let $T_k = (V_k, E_k)$ denote the tree obtained from T by replacing every edge connecting a vertex u of T to its parent by a path of length $k^{\|u\|}$. We identify the degree 3 vertices of T_k with the vertices of T. For each vertex $u \in V$, let $S(u)$ be a binary tree with root ρ_u and let $S_k(u)$ be the tree obtained from $S(u)$ by replacing every edge with a path of length $k^{\|u\|+1}$. Finally, for each pair of positive integers (k, m), let G_k^m be the graph obtained from T_k by, for each vertex $u \in V$, adding a path of length $k^{\|u\|+1}$ connecting u to ρ_u and then replacing every edge in each of these added paths and every edge in each of the trees $S_k(u)$ by m parallel edges. The vertex degrees of G_k^m are bounded by $3 + m$, and the identity

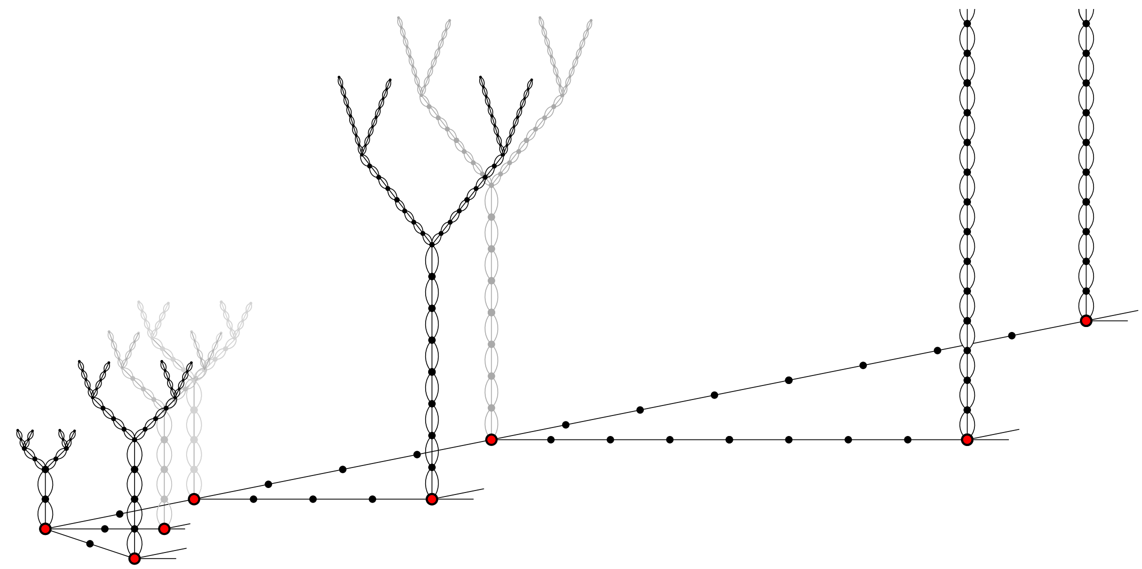

Figure 5.2: An illustration of the graph G_2^3. Red vertices correspond to vertices of the 3-regular tree T. Only three generations of each of the trees $S(v)$ are pictured.

map is an isometry (and hence a rough isometry) between G_k^m and $G_k^{m'}$ whenever k, m and m' are positive integers. See Figure 1 for an illustration.

Let k and m be positive integers. Observe that for every vertex v of T and every child u of v in T, the probability that simple random walk on G_k^m started at u ever hits v does not depend on the choice of v or u. Denote this probability $p(m,k)$. We can bound $p(m,k)$ as follows.

$$\frac{k}{k+2+m} \leq p(m,k) \leq \frac{k+2}{k+2+2m}. \tag{5.6.1}$$

The lower bound of $k/(k+2+m)$ is exactly the probability that the random walk started at u visits v before visiting any other vertex of T or visiting ρ_u. The upper bound of $(k+2)/(k+2+2m)$ is exactly the probability that the random walk started at u ever visits a neighbour of u in T. This can be computed by a straightforward network reduction (see [173] for background): The conductance to infinity from the root of a binary tree is 1, so that, by the series and parallel laws, the effective conductance to infinity from u in the subgraph of G_k^m spanned by the vertices of $S_k(u)$ and the path connecting u to $S_k(u)$ is $2mk^{-\|u\|-1}$. On the other hand, the effective conductance between u and its parent v is $k^{-\|u\|}$, while the effective conductance between u and each of its children is $k^{-\|u\|-1}$. It follows that the probability that a random walk started at u ever visits a neighbour of u in T is exactly

$$\frac{k^{-\|u\|} + 2k^{-\|u\|-1}}{k^{-\|u\|} + 2k^{-\|u\|-1} + 2mk^{-\|u\|-1}} = \frac{k+2}{k+2+2m}$$

as claimed.

Let $\mathfrak{F}_k^m$ be a sample of $\mathsf{WUSF}_{G_k^m}$ generated using Wilson's algorithm on G_k^m, starting with the root ρ of T. Let ξ be the loop-erased random walk in G_k^m beginning at ρ that is used to start our forest. The path ξ includes either one or none of the neighbours of ρ in T and so, in either case, there are at least two neighbours v_1 and v_2 of ρ in T that are not contained in this path. Continuing to run Wilson's algorithm from v_1 and v_2, we see that, conditional on ξ, the events $A_1 = \{v_1$ is in the past of ρ in $\mathfrak{F}_k^m\}$ and $A_2 = \{v_2$ is in the past of ρ in $\mathfrak{F}_k^m\}$ are independent and each have probability $p(m,k)$. Furthermore, on the event A_i, we add only the path connecting v_i and ρ in G_k^m to the forest during the corresponding step of Wilson's algorithm. Recursively, we see that the restriction to T of the past of ρ in $\mathfrak{F}$ contains a Galton-Watson branching process with Binomial offspring distribution $(2, p(m,k))$. If $k \geq m+3$ this branching process is supercritical, so that $\mathfrak{F}_k^m$ contains a component with uncountably many ends with positive probability. By tail triviality of the WUSF [173, Theorem 10.18], $\mathfrak{F}_k^m$ contains a component with uncountably many ends a.s. when $k \geq m+3$.

On the other hand, a similar analysis shows that the restriction to T of past of ρ in $\mathfrak{F}_k^m$ is stochastically dominated by a binomial $(3, p(m,k))$ branching process. (The 3 here is to account for the possibility that every child of ρ in T is in its past). If $m \geq k+2$, this branching process is either critical or subcritical, and we conclude that the restriction to T of the past of ρ in $\mathfrak{F}_k^m$ is finite a.s. Condition on this restriction. Similarly again to the above, the restriction to $S_k(v)$ of the past of v in $\mathfrak{F}_m^k$ is stochastically dominated by a critical binomial $(2, 1/2)$ branching process for each vertex v of T, and is therefore finite a.s. We conclude that the past of ρ in $\mathfrak{F}_k^m$ is finite a.s. whenever $m \geq k+2$. A similar analysis shows that the past in $\mathfrak{F}_k^m$ of every vertex of G_k^m is finite a.s., and consequently that every component of $\mathfrak{F}_k^m$ is one-ended a.s. whenever $m \geq k+2$.

Since $4 \geq 1+3$ and $6 \geq 4+2$, the wired uniform spanning forest $\mathfrak{F}_4^1$ of G_4^1 contains an infinitely-ended component a.s., and every component of the wired uniform spanning forest $\mathfrak{F}_4^6$ of G_4^6 is one-ended a.s. $\square$

5.7 Closing discussion and open problems

5.7.1 The FMSF of the interlacement ordering

One way to think about the Interlacement Aldous-Broder algorithm is as follows. Given the interlacement process $\mathscr{I}$ on a transient network G, we can define a total ordering of the edges of G according to the order in which they are traversed by the trajectories of $\mathscr{I}_{[0,\infty)}$. That is, we define a strict total ordering $\prec$ of E by setting $e_1 \prec e_2$ if and only if either e_1 is first traversed by a trajectory of $\mathscr{I}_{[0,\infty)}$ at a smaller time than e_2 is first traversed by a trajectory of $\mathscr{I}_{[0,\infty)}$, or if e_1 and e_2 are both traversed for the first time by the same trajectory of $\mathscr{I}_{[0,\infty)}$, and this trajectory traverses e_1 before it traverses e_2. We call $\prec$ the **interlacement ordering** of the edge set E.

It is easily verified that $\mathsf{AB}_0(\mathscr{I})$ is the **wired minimal spanning forest** of G with respect to the interlacement ordering. That is, an edge $e \in E$ is included in $\mathsf{AB}_0(\mathscr{I})$ if and only if there does not exist either a finite cycle or a bi-infinite path in G containing e for which e is the $\prec$-maximal element. See [173] for background on minimal spanning forests. In light of this, it is natural to wonder what might be said about the *free* minimal spanning forest of the interlacement ordering, that is, the spanning forest of G that includes an edge $e \in E$ if and only if there does not exist a finite cycle in G containing e for which e is the $\prec$-maximal element. Indeed, if this forest were the FUSF of G, this could be used to solve the monotone coupling problem [173, Question 10.6] (see also [57, 177, 181]) and the almost-connectivity problem [173, Question 10.12].

Unfortunately there is little reason for this to be the case other than wishful thinking. Indeed, let $\mathscr{I}$ be the interlacement process on a transient network G, and define

$$t_c = \inf\left\{t \in (0,\infty) : \mathcal{I}_{[0,t]} \text{ is connected a.s.}\right\}.$$

Teixeira and Tykesson [219] proved that if G is transitive, then t_c is positive if and only if G is nonamenable. (The amenable case of their result generalises the corresponding result for $\mathbb{Z}^d$, due to Sznitman [216].) We can apply this result to prove that the free minimal spanning forest of the interlacement ordering is distinct from the WUSF on any nonamenable transitive graph: This is similar to how the usual FMSF and WMSF (where the edge weights are i.i.d.) are distinct if and only there is a nonempty nonuniqueness phase for Bernoulli bond percolation [175]. Since there are many nonamenable transitive graphs where the WUSF and FUSF coincide (e.g. the product of a 3-regular tree with $\mathbb{Z}$, see [173, Chapter 10]), we deduce that there are transitive graphs (indeed, Cayley graphs) for which the FUSF does not coincide with the free minimal spanning forest of the interlacement ordering.

We now give a quick sketch of this argument. Suppose that G is a transitive nonamenable graph. Observe that for every $t < t_c$, there must exist a connected component of $\mathcal{I}_{[0,t)}$ and a vertex u of G such that a random walk started at u has a positive probability not to hit the component: If not, we would have that $\mathcal{I}_{[0,s]}$ was connected for every $s > t$, contradicting the assumption that $t < t_c$. Moreover, by finding a path from u to the component and considering the last vertex of the path before we reach the component, the vertex u can be taken to be adjacent to the component. Let τ be the first time after $t_c/2$ that u is hit by a trajectory of $\mathscr{I}$, and let $e_{t_c/2}(u)$ be the oriented edge that is traversed by this trajectory as it enters u for the first time. Denote this trajectory by W. No other trajectories of $\mathscr{I}$ appear at time τ a.s. In light of the above discussion, by making local modifications to finitely many trajectories in $\mathscr{I}_{[0,\tau)}$, we see that the following event occurs with positive probability: τ is strictly less than t_c, the vertices u and $e_{t_c/2}(u)^-$ are both in different components of $\mathcal{I}_{[0,\tau)}$ (and, in particular, are both in $\mathcal{I}_{[0,\tau)}$), and W hits the component of u in $\mathcal{I}_{[0,\tau)}$ for the first time at u. On this event we must have that $e_{t_c/2}(u)$ is included in the free minimal spanning forest of the interlacement ordering, but is not in $\mathsf{AB}_0(\mathscr{I})$, and hence the two forests do not coincide.

5.7.2 Exceptional times

A natural question raised by the Interlacement Aldous-Broder algorithm concerns the existence or non-existence of **exceptional times** for the process $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}} = \langle \mathsf{AB}_t(\mathscr{I})\rangle_{t\in\mathbb{R}}$, that is, times at which $\mathfrak{F}_t$ has properties markedly different from the a.s. properties of $\mathfrak{F}_0$. For example, we might ask whether, considering the process $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}}$ on $\mathbb{Z}^d$ $(d \geq 3)$, there are exceptional times when the forest has multiply ended components, is disconnected (if $d = 3, 4$), or is connected (if $d \geq 5$). (Note that the proof of Theorem 5.2.2 implies that there do not exist exceptional times at which $\mathfrak{F}_t$ contains *indestructible* excessive ends.)

The answers to the first of these questions turn out to rather simple. Given a trajectory W in a graph G and a vertex u of G visited by the path, we define $e(W, u)$ to be the oriented edge pointing into u that is traversed by W as it enters u for the first time, and define

$$\mathsf{AB}(W) = \{-e(W, u) : u \text{ is visited by } W\}.$$

Note that if the trace of W is infinite then $\mathsf{AB}(W)$ is an infinite oriented tree. We define the tree of first entry edges $\mathsf{AB}(X)$ similarly when $X = \langle X_n \rangle_{n\geq 0}$ is a path in G.

Proposition 5.7.1 (Exceptional times for excessive ends)**.** *Let G be a transient network, let $\mathscr{I}$ be the interlacement process, and let $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}} = \langle \mathsf{AB}_t(\mathscr{I}) \rangle_{t\in\mathbb{R}}$. Let $\mathscr{E}$ be the set of times $t \in \mathbb{R}$ such that $\mathfrak{F}_t$ has a multiply ended component, and let $\mathscr{E}'$ be the set of times $t \in \mathbb{R}$ for which there exists a trajectory W_t in $\mathscr{I}_t$ such that $\mathsf{AB}(W)$ is multiply ended. If every component of $\mathfrak{F}_0$ is one-ended almost surely, then the following hold almost surely.*

1. *$\mathscr{E} = \mathscr{E}'$, and $\mathscr{E} = \emptyset$ if and only if $\mathfrak{F}_0$ is connected almost surely.*
2. *For every $t \in \mathscr{E}$, there is exactly one two-ended component of $\mathfrak{F}_t$, and all other components are one-ended. The unique two-ended component is the union of the tree $\mathsf{AB}(W_t)$ with some finite bushes.*

Since Proposition 5.7.1 is tangential to the paper, we leave out some details from the proof.

Proof. We first prove that $\mathscr{E} = \mathscr{E}'$ almost surely. The containment $\mathscr{E}' \subseteq \mathscr{E}$ is immediate, and holds deterministically. Let Ω be the almost sure event that every component of $\mathfrak{F}_t$ is one-ended for every rational t, and that no two trajectories of $\mathscr{I}$ have the same arrival time. We claim that $\mathscr{E} = \mathscr{E}'$ pointwise on the event Ω. Suppose that Ω holds and that $t \in \mathscr{E}$, so that there exists a sequence of vertices $\langle v_i \rangle_{i\geq 0}$ such that $v_i = e_t(v_{i+1})^-$ for each $i \geq 0$. In particular, the arrival times $\tau_t(v_i)$ are increasing. We claim that we must have $\tau_t(v_i) = t$ for all $i \geq 0$. Indeed, if $\tau_t(v_i) \geq t + \varepsilon$ for some $\varepsilon > 0$ and all i larger than some i_0, then we would have that $e_{t+\delta}(v_i) = e_t(v_i)$ for all $0 < \delta \leq \varepsilon$ and all $i \geq i_0$. In this situation, we would therefore have that $\mathfrak{F}_{t+\delta}$ contained a multiply ended component for every $0 < \delta \leq \varepsilon$, contradicting the assumption that the event Ω occured. Thus, we must have that there exists a trajectory $W_t \in \mathscr{I}_t$ (which is unique by definition of Ω), and the sequence v_i gives an excessive end in the tree $\mathsf{AB}(W)$. Since the sequence v_i represented an arbitrary excessive end of $\mathfrak{F}_t$, it follows that every excessive end of $\mathfrak{F}_t$ arises from the tree $\mathsf{AB}(W)$ on the event Ω.

Now, if $\mathfrak{F}_0$ is not connected a.s., then there is a vertex v of G such that two independent random walks from v do not intersect with positive probability, and it follows that there a.s. exist trajectories in $\mathscr{I}$ such that $\mathsf{AB}(W)$ has at least two ends. It remains to prove that the trees $\mathsf{AB}(W)$ have at most two-ends for every trajectory W in $\mathscr{I}$, and are all one-ended if $\mathfrak{F}_0$ is a.s. connected. Since there are only countably many trajectories in $\mathscr{I}$, it suffices to analyze a single bi-infinite random walk. To prove this, it is convenient to introduce a variant of the interlacement Aldous-Broder in which we first run a simple random walk started from a fixed vertex (considered to arrive at time zero), and then run the interlacement process $\mathscr{I}_{[0,\infty)}$, and form a forest from the first entry edges. It is not difficult to see, by a slight modification of the proof Theorem 5.1.1, that the forest produced this was is the wired uniform spanning forest: In the finite exhaustion, this corresponds to first running a random walk from v until hitting the distinguished boundary vertex, and then decomposing the rest of the walk into excursions from the boundary vertex. Using this algorithm, it follows that $\mathsf{AB}(X)$ is one-ended a.s. whenever X is a random walk on a transient graph G for which the wired uniform spanning forest is one-ended.

Now suppose that $W = \langle W_n \rangle_{n\in\mathbb{Z}}$ is a bi-infinite random walk. If $\langle v_i \rangle_{i\geq 0}$ is a sequence of vertices in G corresponding to an excessive end of $\mathsf{AB}(W)$, then we must have that $v_i = e(v_{i+1})^-$ for all i sufficiently large, and it follows that this excessive end must be an end of the tree $\mathsf{AB}(\langle W_i \rangle_{i\geq 0})$, completing the proof that $\mathsf{AB}(W)$ has at most two ends. On the other hand, we note that the unique path to infinity from W_0 in $\mathsf{AB}(\langle W_i \rangle_{i\geq 0})$ is exactly the loop-erasure of $\langle W_i \rangle_{i\geq 0}$, and if $\mathfrak{F}_0$ is connected a.s. then this path is hit infinitely often a.s. by $\langle W_i \rangle_{i<0}$. We deduce that in this case this end is not present in the tree $\mathsf{AB}(W)$, completing the proof. $\square$

We do not know if there exist exceptional times for (dis)connectivity. We expect that such times do not exist, but it would be very interesting if they do. kirchhoff1847ueber

Question 5.7.2. *Let $d \geq 3$, let $\mathscr{I}$ be the interlacement process on $\mathbb{Z}^d$, and let $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}} = \langle \mathsf{AB}_t(\mathscr{I}) \rangle_{t\in\mathbb{R}}$. If $d = 3, 4$, do there exist times at which $\mathfrak{F}_t$ is disconnected? If $d \geq 5$, do there exist times at which $\mathfrak{F}_t$ is connected?*

If the answer to Question 5.7.2 is positive, it would be interesting to further understand the structure of the set of exceptional times and the geometry of the forest $\mathfrak{F}_t$ at a typical exceptional time. It is easy to see that, unlike for excessive ends, the arrival times of trajectories are *not* exceptional times for connectivity, so if exceptional times do exist they are likely to have a more interesting structure. We note that there is a rich theory of exceptional times for other models such as dynamical percolation, addressing many analogous questions. See e.g. [99, 115, 125, 213].

A related question concerns the decorrelation of connectivity events under the dynamics.

Question 5.7.3. *Let $d \geq 5$, let $\mathscr{I}$ be the interlacement process on $\mathbb{Z}^d$, and let $\langle \mathfrak{F}_t \rangle_{t\in\mathbb{R}} = \langle \mathsf{AB}_t(\mathscr{I}) \rangle_{t\in\mathbb{R}}$. How does*

$$\mathbb{P}(x \text{ is connected to } y \text{ in both } \mathfrak{F}_0 \text{ and } \mathfrak{F}_t)$$

behave as a function of the vertices $x, y \in \mathbb{Z}^d$ and the number $t > 0$? Does the behaviour as a function of x, y undergo a phase transition as t is increased?

Recall that for the USF of $\mathbb{Z}^d$, $d \geq 5$, the probability that two vertices x and y are in the same component of the USF decays like $\|x-y\|^{-(d-4)}$ as $\|x-y\| \to \infty$ [41]. A successful approach to Questions 5.7.2 and 5.7.3 might need to draw more deeply on the interlacement literature than we have needed to in this paper.

5.7.3 Excessive ends via update tolerance

A key tool in the study of the USFs carried out in [127, 131, 223] is the **update-tolerance** of the USFs (referred to as **weak insertion tolerance** by Timár [223]). Given a sample $\mathfrak{F}$ of either the WUSF and the FUSF of a network G and an oriented edge e of G not in $\mathfrak{F}$, update-tolerance states that there exists a forest $U(\mathfrak{F}, e)$, obtained from $\mathfrak{F}$ by adding e and deleting some other appropriately chosen edge d, such that the law of $U(\mathfrak{F}, e)$ is absolutely continuous with respect to that of $\mathfrak{F}$. The forest $U(\mathfrak{F}, e)$ is called the update of $\mathfrak{F}$ at e. See [127, 131, 223] for further details.

Since the number of excessive ends does not change when we perform an update, a positive solution of the following conjecture would yield an alternative proof of Theorem 5.2.2. We say that a Borel set $\mathscr{A} \subseteq \{0,1\}^E$ is **update-stable** if for every oriented edge e of G, the updated forest $U(\mathfrak{F}, e)$ is in $\mathscr{A}$ if and only if $\mathfrak{F}$ is in $\mathscr{A}$ almost surely. The conjecture would also imply a positive solution to [44, Question 15.7].

Conjecture 5.7.4. *Let G be an infinite network, and let $\mathfrak{F}$ be either the wired or free spanning forest of G. Then for every update-stable Borel set $\mathscr{A} \subseteq \{0,1\}^E$, the probability that $\mathfrak{F}$ is in $\mathscr{A}$ is either zero or one.*

5.7.4 Ends in uniformly transient networks

The following natural question remains open. If true, it would strengthen the results of Lyons, Morris and Schramm [170]. A network is said to be **uniformly transient** if the capacities of the vertices of the network are bounded below by a positive constant.

Question 5.7.5. *Let G be a uniformly transient network with $\inf_e c(e) > 0$. Does it follow that every component of the wired uniform spanning forest of G is one-ended almost surely?*

The argument used in the proof of Theorem 5.2.2 can be adapted to show that, under the hypotheses of Question 5.7.5, every component of the WUSF is either one-ended or has uncountably many ends, with no isolated excessive ends. To answer Question 5.7.5 positively, it remains to rule this second case out.

Chapter 6

Indistinguishability of trees in uniform spanning forests

Summary. We prove that in both the free and the wired uniform spanning forest (FUSF and WUSF) of any unimodular random rooted network (in particular, of any Cayley graph), it is impossible to distinguish the connected components of the forest from each other by invariantly defined graph properties almost surely. This confirms a conjecture of Benjamini, Lyons, Peres and Schramm [44].

We also answer positively two additional questions of [44] under the assumption of unimodularity. We prove that on any unimodular random rooted network, the FUSF is either connected or has infinitely many connected components almost surely, and, if the FUSF and WUSF are distinct, then every component of the FUSF is transient and infinitely-ended almost surely. All of these results are new even for Cayley graphs.

6.1 Introduction

The **Free Uniform Spanning Forest** (FUSF) and the **Wired Uniform Spanning Forest** (WUSF) of an infinite graph G are defined as weak limits of the uniform spanning trees on large finite subgraphs of G, taken with either free or wired boundary conditions respectively (see Section 6.1.2 for details). First studied by Pemantle [190], the USFs are closely related many other areas of probability, including electrical networks [62, 154], Lawler's loop-erased random walk [44, 161, 228], sampling algorithms [197, 228], domino tiling [150], the Abelian sandpile model [137, 138, 178], the rotor-router model [126], and the Fortuin-Kasteleyn random cluster model [105, 109]. The USFs are also of interest in group theory, where the FUSFs of Cayley graphs are related to the ℓ^2-Betti numbers [94, 169] and to the fixed price problem of Gaboriau [95], and have also been used to approach the Dixmier problem [87].

Although both USFs are defined as limits of trees, they need not be connected. Indeed, a principal result of Pemantle [190] is that the FUSF and WUSF coincide on $\mathbb{Z}^d$ for all $d \geq 1$ and that they are connected almost surely (a.s.) if and only if $d \leq 4$. A complete characterisation of the connectivity of the WUSF was given by Benjamini, Lyons, Peres and Schramm (henceforth referred to as BLPS) in their seminal work [44], who showed that the WUSF of a graph G is connected a.s. if and only if the traces of two simple random walks started at arbitrary vertices of G a.s. intersect. This recovers Pemantle's result on $\mathbb{Z}^d$, and shows more generally that the WUSF of a Cayley graph is connected a.s. if and only if the corresponding group has polynomial growth of degree at most 4 [124, 173].

Besides connectivity, several other basic features of the WUSF are also understood rather firmly. This understanding mostly stems from Wilson's algorithm rooted at infinity, which allows the WUSF to be sampled by joining together loop-erased random walks [173, 228]. For example, other than connectivity, the simplest property of a forest is the number of *ends* its components have. Here, an infinite graph G is said to be k-ended if, over all finite sets of vertices W, the subgraph induced by $V \setminus W$ has a maximum of k infinite connected components. In particular, an infinite tree is one-ended if and only if it does not contain a simple bi-infinite path. Following earlier work by Pemantle [190], BLPS [44] proved that the number of components of the WUSF of any graph is non-random, that the WUSF of any unimodular transitive graph (e.g., any Cayley graph) is either connected or has infinitely many components a.s., and that in both cases every component of the WUSF is one-ended a.s. unless the underlying graph is itself two-ended. Morris [185] later proved that every component of the WUSF is recurrent a.s. on any graph, confirming a conjecture of BLPS [44, Conjecture 15.1], and several other classes of graphs have also been shown to have one-ended WUSF components [7, 127, 170].

Much less is known about the FUSF. No characterisation of its connectivity is known, nor is it known whether the number of components of the FUSF is non-random on an any graph. In [44] it is proved that if the FUSF and WUSF differ on a unimodular transitive graph, then a.s. the FUSF has a transient tree with infinitely many ends, in contrast to the WUSF. However, it remained an open problem [44, Question 15.8] to prove that, under the same hypotheses, *every* connected component of the FUSF is transient and infinitely ended a.s. In light of this, it is natural to ask the following more general question:

Question. *Let G be a unimodular transitive graph. Can the components of the free uniform spanning forest of G be very different from each other?*

Questions of this form were first studied by Lyons and Schramm [176] in the context *insertion-tolerant* automorphism-invariant random subgraphs. Their remarkable theorem asserts that in any such random subgraph (e.g. Bernoulli bond percolation or the Fortuin-Kasteleyn random cluster model) on a unimodular transitive graph, one cannot distinguish between the infinite connected components using automorphism-invariant graph properties. For example, all such components must have the same volume growth, spectral dimension, value of p_c and so forth (see Section 6.1.3 for further examples). They also exhibited applications of indistinguishability to statements not of this form, including uniqueness monotonicity and connectivity decay. Here, a random subgraph ω of a graph G is insertion-tolerance if for every edge e of G, the law of the subgraph $\omega \cup \{e\}$ formed by inserting e into ω is absolutely continuous with respect to the law of ω. The uniform spanning forests are clearly not insertion-tolerant, since the addition of an edge may close a cycle.

BLPS conjectured [44, Conjecture 15.9] that the components of both the WUSF and FUSF also exhibit this form of indistinguishability. In this paper we confirm this conjecture.

Theorem 6.1.1 (Indistinguishability of USF components)**.** *Let G be a unimodular transitive graph, and let $\mathfrak{F}$ be a sample of either the free uniform spanning forest or the wired uniform spanning forest* *of G. Then for each automorphism-invariant Borel-measurable set $\mathscr{A}$ of subgraphs of G, either every connected component of $\mathfrak{F}$ is in $\mathscr{A}$ or every connected component of $\mathfrak{F}$ is not in $\mathscr{A}$ almost surely.*

As indicated by the above discussion, Theorem 6.1.1 implies the following positive answer to [44, Question 15.8] under the assumption of unimodularity.

Theorem 6.1.2 (Transient trees in the FUSF). *Let G be a unimodular transitive graph and let $\mathfrak{F}$ be a sample of FUSF_G. If the measures FUSF_G and WUSF_G are distinct, then every component of $\mathfrak{F}$ is transient and has infinitely many ends almost surely.*

However, rather than deducing Theorem 6.1.2 from Theorem 6.1.1, we instead prove Theorem 6.1.2 directly and apply it in the proof of Theorem 6.1.1.

We also apply Theorem 6.1.1 to answer another of the most basic open problems about the FUSF [44, Question 15.6] under the assumption of unimodularity.

Theorem 6.1.3 (Number of trees in the FUSF). *Let G be a unimodular transitive graph and let $\mathfrak{F}$ be a sample of the free uniform spanning forest of G. Then $\mathfrak{F}$ is either connected or has infinitely many components almost surely.*

The derivation of Theorem 6.1.3 from Theorem 6.1.1 is inspired by the proof of [176, Theorem 4.1], and also establishes the following result.

Theorem 6.1.4 (Connectivity decay in the FUSF). *Let G be a unimodular transitive graph and let $\mathfrak{F}$ be a sample of the free uniform spanning forest of G. If $\mathfrak{F}$ is disconnected a.s., then for every vertex v of G,*

$$\inf\left\{\mathsf{FUSF}_G(u \in T_{\mathfrak{F}}(v)) : u \in V(G)\right\} = 0\,,$$

where $u \in T_{\mathfrak{F}}(v)$ is the event that u belongs to the component of v in $\mathfrak{F}$.

We prove all of our results in the much more general setting of *unimodular random rooted networks*, which includes all Cayley graphs as well as a wide range of popular infinite random graphs and networks [7]. For example, our results hold when the underlying graph is an infinite supercritical percolation cluster in a Cayley graph, a hyperbolic unimodular random triangulation [37, 75] (for which the FUSF and WUSF are shown to be distinct in the upcoming work [20]), a supercritical Galton-Watson tree, or even a component of the FUSF of another unimodular random rooted network. See Section 6.1.3 for the strongest and most general statements.

Organization. In Section 6.1.1 we describe our approach and the novel ingredients of our proof. The necessary background, including definitions of USFs and unimodular random rooted networks are presented in Section 6.1.2. In Section 6.1.3 we define the graph properties we will work with, state the most general and strongest versions of our theorems (most importantly, Theorem 6.1.9), and provide several illustrative examples. In Section 6.2 we develop the update-tolerance property of the FUSF, and prove, in the setting of Theorem 6.1.9, that if the FUSF and WUSF are distinct then every component of the FUSF is transient and infinitely-ended (Theorem 6.1.12), and then prove indistinguishability of the components in this case. In Section 6.3, still in the case where the FUSF and WUSF are distinct, we prove that the FUSF is either connected or has infinitely many connected components (Theorem 6.1.10) and in the latter case we show that connectivity decay is exhibited (Theorem 6.1.11). In Section 6.4 we show that the WUSF components are indistinguishable, completing the proof of Theorem 6.1.9.

Remark. After this paper was posted on the arXiv, Adam Timár posted independent work [223] in which he proves Theorem 6.1.1 for the FUSF only in the case that $\mathsf{FUSF} \neq \mathsf{WUSF}$, and also proves Theorems 6.1.2 and 6.1.3; Indistinguishability of components in the WUSF is not treated. In the present paper, we prove [44, Conjecture 15.1] in its entirety for both the FUSF and WUSF.

6.1.1 About the proof

In [176], Lyons and Schramm argue that the coexistence of clusters of different types in an invariant edge percolation implies the existence of infinitely many *pivotal edges*, that is, closed edges that change the type of an infinite cluster if they are inserted. When the percolation is insertion-tolerant, this heuristically contradicts the Borel-measurability of the property, as the existence of pivotal edges far away from the origin should imply that we cannot approximate the event that the cluster has the property by a cylinder event. This argument was made precise in [176]. Unimodularity of the underlying graph was used heavily – indeed, indistinguishability can fail without it [176, Remark 3.16].

A crucial ingredient of our proof is an *update-tolerance* property of USFs. This property was introduced for the WUSF by the first author [127] and is developed for the FUSF in Section 6.2.1. This property allows us to make a local modification to a sample of the FUSF or WUSF in such a way that the law of the resulting modified forest is absolutely continuous with respect to the law of the forest that we started with. In this local modification, we add an edge of our choice to the USF and, in exchange, are required to remove an edge emanating from the same vertex. The edge that we are required to remove is random and depends upon both the edge we wish to insert and on the entire sample of the USF.

Update-tolerance replaces insertion-tolerance and allows us to perform a variant of the key argument in [176]. However, several obstacles arise as we are required to erase an edge at the same time as inserting one. In particular, we cannot simply open a closed edge connecting two clusters of different types in order to form a single cluster. These obstacles are particularly severe for the WUSF (and the FUSF in the case that the two coincide), where it is no longer the case that the coexistance of components of different types implies the existence of pivotal edges. To proceed, we separate the component properties into two types, tail and non-tail, according to whether the property is sensitive to finite modifications of the component. Indistinguishability is then proven by a different argument in each case: non-tail properties are handled by a variant of the Lyons-Schramm method, while tail properties are handled by a completely separate argument utilising Wilson's algorithm [228] and the spatial Markov property. The proof that components of the WUSF cannot be distinguished by tail properties also applies to transitive graphs without the assumption of unimodularity.

6.1.2 Background and definitions

Notation

A **tree** is a connected graph with no cycles. A **spanning tree** of a graph $G = (V, E)$ is a connected subgraph of G that contains every vertex and no cycles. A **forest** is a graph with no cycles, and a **spanning forest** of a graph $G = (V, E)$ is a subgraph of G that contains every vertex and no cycles. Given a forest $\mathfrak{F}$ and a vertex v we write $T_{\mathfrak{F}}(v)$ for the connected component of $\mathfrak{F}$ containing v. An **essential** spanning forest is a spanning forest such that every component is infinite. A **branch** of an infinite tree T is an infinite component of $T \setminus v$ for some vertex v. The **core** of an infinite tree T, denoted $\mathrm{core}(T)$, is the set of vertices of T such that $T \setminus v$ at least two infinite connected components.

Recall that an infinite graph G is said to be **k-ended** if removing a finite set of vertices W from G results in a maximum of k distinct infinite connected components. In particular, an infinite tree is one-ended if and only if it does not contain any simple bi-infinite paths. We say that a forest $\mathfrak{F}$ is one-ended if all of its components are one-ended. The **past** of a vertex v in a one-ended forest, denoted $\mathrm{past}_{\mathfrak{F}}(v)$, is the union of v and the finite components of $\mathfrak{F} \setminus v$. The **future** of the vertex v is the set of u such that $v \in \mathrm{past}_{\mathfrak{F}}(u)$.

We write $B_G(v, r)$ for the graph-distance ball of radius r about a vertex v in a graph G.

Uniform Spanning Trees and Forests

We now briefly provide the necessary definitions, notation and background concerning USFs. We refer the reader to [173, §4 and §10] for a comprehensive review of this theory. Given a graph $G = (V, E)$ we will refer to an edge $e \in E$ both as an oriented and unoriented edge and it will always be clear which one from the context. Most frequently we will deal with oriented edges and in this case we orient them from their tail e^- to their head e^+.

A **network** (G, c) is a locally finite, connected multi-graph $G = (V, E)$ together with a function $c : E \to (0, \infty)$ assigning a positive **conductance** to each edge of G. Graphs are considered to be networks by setting $c \equiv 1$. The distinction between graphs and networks does not play much of a role for us, and we will mostly suppress the notation of conductances, writing G to mean either a graph or a network. Write $c(u)$ for the sum of the conductances of the edges $e^- = u$ emanating from u and $c(u, v)$ for the conductance of the sum of the conductances of the (possibly many) edges with endpoints u and v. The **random walk** $\langle X_n \rangle_{n \geq 0}$ on a network G is the Markov chain on V with transition probabilities $p(u, v) = c(u, v)/c(u)$.

The **uniform spanning tree** measure UST_G of a finite connected graph G is the uniform measure on spanning trees of G (considered for measure-theoretic purposes as functions $E \to \{0, 1\}$). When G is a network, UST_G is the probability measure on spanning trees of G such that the probability of a tree is proportional to the product of the conductances of its edges.

Let G be an infinite network. An **exhaustion** $\langle V_n \rangle_{n\geq 0}$ of G is an increasing sequence of finite sets of vertices $V_n \subset V$ such that $\bigcup_{n\geq 0} V_n = V$. Given such an exhaustion, we define G_n to be the subgraph of G induced by V_n together with the conductances inherited from G, and define G_n^* to be the network obtained from G by identifying (or "wiring") $V \setminus V_n$ into a single vertex and deleting all the self-loops that are created. The weak limits of the measures UST_{G_n} and $\mathsf{UST}_{G_n^*}$ exist for any network and do not depend on the choice of exhaustion [109, 190]. The limit of the UST_{G_n} is called the **free uniform spanning forest** measure FUSF_G while the limit of the $\mathsf{UST}_{G_n^*}$ is called the **wired uniform spanning forest** measure WUSF_G. Both limits are clearly concentrated on the set of essential spanning forests of G.

The measures FUSF_G and WUSF_G coincide if and only if G does not support any non-constant harmonic functions of finite Dirichlet energy [44], and in particular the two measures coincide when $G = \mathbb{Z}^d$. The two measures also coincide on every amenable transitive graph [44, Corollary 10.9], and an analogous statement holds for unimodular random rooted networks once an appropriate notion of amenability is adopted [7, §8]. When G is a Cayley graph, the two measures FUSF_G and WUSF_G coincide if and only if the *first ℓ^2-Betti number* of the corresponding group is zero [169]. By taking various free or direct products of groups and estimating their Betti numbers, this characterization allows to construct an abundance of Cayley graphs in which the two measures either coincide or differ [173, §10.2].

A very useful property of the UST and the USFs is the *spatial Markov property*. Let G be a network and let H and F be finite subsets of G. We write $\hat{G} = (G - H)/F$ for the network formed from G by deleting each edge $h \in H$ and contracting (i.e., identifying the two endpoints of) each edge $f \in F$. If G is finite and T is a sample of UST_G, then the law of T conditioned on the event $\{F \subseteq T, H \cap T = \emptyset\}$ (assuming this event has positive probability) is equal to the law of the union of F with an independent copy of $\mathsf{UST}_{\hat{G}}$, considered as a subgraph of G [173, §4]. Now suppose that G is an infinite network with exhaustion $\langle V_n \rangle_{n\geq 0}$ and let $\mathfrak{F}$ be a sample of either FUSF_G or WUSF_G. Applying the Markov property to the finite networks G_n and G_n^* and taking the limit as $n \to \infty$, we see similarly that the conditional distribution of $\mathfrak{F}$ conditioned on the event $\{F \subseteq \mathfrak{F}, H \cap \mathfrak{F} = \emptyset\}$ is equal to the law of the union of F with an independent copy of $\mathsf{FUSF}_{\hat{G}}$ or $\mathsf{WUSF}_{\hat{G}}$ as appropriate. It is important here that H and F are finite.

Lastly, throughout Section 6.4 we will use a recent result of the first author regarding ends of the WUSF's components. Components of the WUSF are known to be one-ended a.s. in several large classes of graphs and networks. The following is proven by the first author in [127], and follows earlier works [7, 44, 170, 190].

Theorem 6.1.5 ([127]). *Let (G, ρ) be transient unimodular random rooted network with $\mathbb{E}[c(\rho)] < \infty$. Then every component of the wired uniform spanning forest of G is one-ended almost surely.*

Unimodular random networks

We present here the necessary definition of unimodular random networks and refer the reader to the comprehensive monograph of Aldous and Lyons [7] for more details and many examples. A **rooted graph** (G, ρ) is a locally finite, connected graph G together with a distinguished vertex ρ, the root. An isomorphism of graphs is an isomorphism of rooted graphs if it preserves the root. The ball of radius r around a vertex v of G, denoted $B_G(v, r)$ is the graph induced on the set of vertices which are at graph distance at most r from v. The **local topology** on the set of isomorphism classes of rooted graphs is defined so that two (isomorphism classes of) rooted graphs (G, ρ) and (G', ρ') are close to each other if and only if the rooted balls $(B_G(\rho, r), \rho)$ and $(B_{G'}(\rho', r), \rho')$ are isomorphic to each other for large r. We denote the space of isomorphism classes of rooted graphs endowed with the local topology by $\mathcal{G}_\bullet$. We define an **edge-marked graph** to be a locally finite connected graph together with a function $m : E(G) \to \mathbb{X}$ for some separable metric space $\mathbb{X}$, the **mark space** (in this paper, $\mathbb{X}$ will be a product of intervals and some copies of $\{0, 1\}$). For example, if $G = (G, c)$ is a network and $\mathfrak{F}$ is a sample of FUSF_G, then $(G, c, \mathfrak{F})$ is a graph with marks in $(0, \infty) \times \{0, 1\}$. The local topology on rooted marked graphs is defined so that two marked rooted graphs are close if for large r there is an isomorphism (of rooted graphs) $\phi : (B_G(\rho, r), m, \rho) \to (B_{G'}(\rho', r), \rho')$ such that $d_{\mathbb{X}}(m'(\phi(e)), m(e))$ is small for every edge e in $B_G(\rho, r)$. We denote the space of edge-marked graphs with marks in $\mathbb{X}$ by $\mathcal{G}_\bullet^{\mathbb{X}}$.

Similarly, we define a **doubly-rooted graph** (G, u, v) to be a graph together with an ordered pair of distinguished vertices. The space $\mathcal{G}_{\bullet\bullet}$ of doubly-rooted graphs is defined similarly to $\mathcal{G}_\bullet$. A random rooted graph (G, ρ) is **unimodular** if it obeys the **mass-transport principle**. That is, for every non-negative Borel function $f : \mathcal{G}_{\bullet\bullet} \to [0, \infty]$ – which we call a **mass transport** – we have that

$$\mathbb{E} \sum_{v \in V} f(G, \rho, v) = \mathbb{E} \sum_{u \in V} f(G, u, \rho).$$

In other words, (G, ρ) is unimodular if for every mass transport f, the expected mass received by the root equals the expected mass sent by the root. Every Cayley graph (rooted at any vertex) is a unimodular random rooted graph (whose law is concentrated on a singleton), as is every unimodular transitive graph [173, §8]. For many examples of a more genuinely random nature, see [7]. Unimodular random rooted networks and other edge-marked graphs are defined similarly.

When (G, ρ) is a unimodular random rooted network and $\mathfrak{F}$ is a sample of either FUSF_G or WUSF_G, $(G, \rho, \mathfrak{F})$ is also unimodular: Since the definitions of FUSF_G and WUSF_G do not depend on the choice of exhaustion, for each mass transport $f : \mathcal{G}_{\bullet\bullet}^{(0,\infty)\times\{0,1\}} \to [0, \infty]$, the expectations

$$f^F(G, u, v) = \mathsf{FUSF}_G \left[f(G, u, v, \mathfrak{F}) \right] \qquad \text{and} \qquad f^W(G, u, v) = \mathsf{WUSF}_G \left[f(G, u, v, \mathfrak{F}) \right]$$

are also mass transports. This allows us to deduce the mass-transport principle for $(G, \rho, \mathfrak{F})$ from that of (G, ρ).

Reversibility and stationarity

Let (G, ρ) be a random rooted network and let $\langle X_n \rangle_{n \geq 0}$ be a random walk on G started at ρ. The random rooted graph (G, ρ) is said to be **stationary** if

$$(G, \rho) \stackrel{d}{=} (G, X_1)$$

and **reversible** if

$$(G, \rho, X_1) \stackrel{d}{=} (G, X_1, \rho).$$

While every reversible random rooted graph is trivially stationary, the converse need not hold in general. Indeed, every transitive graph (rooted arbitrarily) is stationary, while it is reversible if and only if it is unimodular. For example, the *grandfather graph* [173] is transitive but not reversible.

The following correspondence between unimodular and reversible random rooted networks is implicit in [7, §4] and is proven explicitly in [38].

Proposition 6.1.6. *If (G, ρ) is a unimodular random rooted network with $\mathbb{E}[c(\rho)] < \infty$, then biasing the law of (G, ρ) by $c(\rho)$ (that is, reweighting the law of (G, ρ) by the Radon-Nikodym derivative $c(\rho)/\mathbb{E}[c(\rho)]$) yields the law of a reversible random rooted network. Conversely, if (G, ρ) is a reversible random rooted network with $\mathbb{E}[c(\rho)^{-1}] < \infty$ then biasing the law of (G, ρ) by $c(\rho)^{-1}$ yields the law of a unimodular random rooted network.*

$$\left\{ \begin{array}{c} (G,\rho) \textit{ unimodular} \\ \textit{with } \mathbb{E}[c(\rho)] < \infty \end{array} \right\} \; \underset{\textit{bias by } c(\rho)^{-1}}{\overset{\textit{bias by } c(\rho)}{\rightleftarrows}} \; \left\{ \begin{array}{c} (G,\rho) \textit{ reversible} \\ \textit{with } \mathbb{E}[c(\rho)^{-1}] < \infty \end{array} \right\}.$$

For example, a finite rooted network is unimodular if and only if, conditioned on G, its root is uniformly distributed on the network, and is reversible if and only if, conditioned on G, the root is distributed according to the stationary distribution of the random walk on the network.

Thus, to prove an almost sure statement about unimodular random rooted networks with $\mathbb{E}[c(\rho)] < \infty$ we can bias by the conductance at the root and work in the reversible setting, and vice versa.

A useful equivalent characterisation of reversibility is as follows. Let $\mathcal{G}_{\leftrightarrow}$ denote the space of isomorphism classes of graphs equipped with a bi-infinite path $(G, \langle x_n \rangle_{n \in \mathbb{Z}})$, which is endowed with a natural variant of the local topology. Let (G, ρ) be a random rooted graph and let $\langle X_n \rangle_{n \geq 0}$ and $\langle X_{-n} \rangle_{n \geq 0}$ be two independent simple random walks started from $X_0 = \rho$, so that $(G, \langle X_n \rangle_{n \in \mathbb{Z}})$ is a random variable taking values in $\mathcal{G}_{\leftrightarrow}$. Then (G, ρ) is reversible if and only if

$$(G, \langle X_n \rangle_{n \in \mathbb{Z}}) \stackrel{d}{=} (G, \langle X_{n+k} \rangle_{n \in \mathbb{Z}}) \quad \forall\, k \in \mathbb{Z}. \tag{6.1.1}$$

Indeed, $(\rho, X_{-1}, \ldots)$ is a simple random walk started from ρ independent of X_1 and, conditional on (G, X_1), reversibility implies that ρ is uniformly distributed among the neighbours of X_1, so that $(X_1, \rho, X_{-1}, X_{-2}, \ldots)$ has the law of a simple random walk from X_1 and (6.1.1) follows. Conversely, (6.1.1) implies that (G,ρ) is reversible by taking $k=1$ and restricting to the 0th and 1st coordinates of the walk.

A useful variant of Proposition 6.1.6 is the following. Suppose that (G,ρ) is a unimodular random rooted network with $\mathbb{E}[c(\rho)]<\infty$, $\mathfrak{F}$ is a sample of either FUSF_G or WUSF_G, and let $c_{\mathfrak{F}}(v)$ denote the sum of the conductances of the edges of G emanating from v that are included in $\mathfrak{F}$. Then, if we sample $(G,\rho,\mathfrak{F})$ biased by $c_{\mathfrak{F}}(\rho)$ and let $\langle X_n\rangle_{n\geq 0}$ and $\langle X_{-n}\rangle_{n\geq 0}$ be independent random walks on $\mathfrak{F}$ starting at ρ, then, by [7, Theorem 4.1],

$$(G,\langle X_n\rangle_{n\in\mathbb{Z}},\mathfrak{F}) \stackrel{d}{=} (G,\langle X_{n+k}\rangle_{n\in\mathbb{Z}},\mathfrak{F}) \quad \forall\, k\in\mathbb{Z}. \tag{6.1.2}$$

Ergodicity

We say that a unimodular random rooted network with $\mathbb{E}[c(\rho)]<\infty$ is **ergodic** if any (and hence all) of the below hold.

Theorem 6.1.7 (Characterisation of ergodicity [7, §4]). *Let (G,ρ) be a unimodular random rooted network with $\mathbb{E}[c(\rho)]<\infty$. The following are equivalent.*

1. *When the law of (G,ρ) is biased by $c(\rho)$ to give an equivalent reversible random rooted network, the stationary sequence $\langle (G,X_n)\rangle_{n\geq 0}$ is ergodic.*
2. *Every event $A\subset\mathcal{G}_\bullet^{(0,\infty)}$ invariant to changing the root has probability in $\{0,1\}$.*
3. *The law of (G,ρ) is an extreme point of the weakly closed convex set of laws of unimodular random rooted networks.*

A similar statement holds for edge-marked networks. Tail triviality of the USFs [44, Theorem 8.3] implies that if (G,ρ) is an ergodic unimodular random rooted network and $\mathfrak{F}$ is a sample of either FUSF_G or WUSF_G, then $(G,\rho,\mathfrak{F})$ is also ergodic.

The extremal characterisation (3) implies (by Choquet theory) that every unimodular random rooted network with $\mathbb{E}[c(\rho)]<\infty$ can be written as a mixture of ergodic unimodular random rooted networks. Thus, to prove a.s. statements about general unimodular random rooted networks it suffices for us to consider ergodic unimodular random rooted networks.

6.1.3 Component properties and indistinguishability on unimodular random rooted networks

General unimodular random rooted graphs and networks have few automorphisms, so that it is not appropriate at this level of generality to phrase indistinguishability in terms of automorphism-invariant properties. Instead, we consider properties that are invariant under rerooting within a component as follows. Consider the space $\mathcal{G}_\bullet^{\{0,1\}}$ of rooted graphs with edges marked by $\omega(e)\in\{0,1\}$, which we think of as a rooted graph together with a distinguished subgraph spanned by the edges $\omega=\{e:\omega(e)=1\}$. Given such a (G,v,ω) we define $K_\omega(v)$ to be the connected component of v in ω.

Definition 6.1.8. A Borel-measurable set $\mathscr{A} \subset \mathcal{G}_\bullet$ is called a **component property** if and only if it is invariant to rerooting within the component of the root, i.e.,

$$(G, v, \omega) \in \mathscr{A} \implies (G, u, \omega) \in \mathscr{A} \ \ \forall u \in K_\omega(v)\,.$$

Again, this may be formulated for networks with the obvious modifications. This definition is equivalent to the one given in [7, Definition 6.14]. We say that a connected component K of ω has property $\mathscr{A}$ (and abuse notation by writing $K \in \mathscr{A}$) if $(G, u, \omega) \in \mathscr{A}$ for some (and hence every) vertex $u \in K$. We are now ready to state our main theorem in its full generality and strength.

Theorem 6.1.9 (Indistinguishability of USF components)**.** *Let (G, ρ) be a unimodular random network with $\mathbb{E}[c(\rho)] < \infty$, and let $\mathfrak{F}$ be a sample of either FUSF_G or WUSF_G. Then for every component property $\mathscr{A}$, either every connected component of $\mathfrak{F}$ has property $\mathscr{A}$ or none of the connected components of $\mathfrak{F}$ have property $\mathscr{A}$ almost surely.*

And we may now restate Theorems 6.1.3, 6.1.4 and 6.1.2 in their full generality.

Theorem 6.1.10. *Let (G, ρ) be a unimodular random rooted network with $\mathbb{E}[c(\rho)] < \infty$ and let $\mathfrak{F}$ be a sample of FUSF_G. Then $\mathfrak{F}$ is either connected or has infinitely many components almost surely.*

Theorem 6.1.11. *Let (G, ρ) be a unimodular random rooted network with $\mathbb{E}[c(\rho)] < \infty$ and let $\mathfrak{F}$ be a sample of FUSF_G. If $\mathfrak{F}$ is disconnected a.s., then a.s. for every vertex v of G,*

$$\inf\{\mathsf{FUSF}_G(u \in T_{\mathfrak{F}}(v)) : u \in V(G)\} = 0.$$

Theorem 6.1.12. *Let (G, ρ) be a unimodular random rooted network and let $\mathfrak{F}$ be a sample of the FUSF_G. On the event that the measures FUSF_G and WUSF_G are distinct, every component of $\mathfrak{F}$ is transient and has infinitely many ends almost surely. This holds both when the edges of $\mathfrak{F}$ are given the conductances inherited from G and when they are given unit conductances.*

It follows that, under the assumptions of Theorem 6.1.12, every component of the FUSF of G has positive speed and critical percolation probability $p_c < 1$ [7].

We remark that, by [96, Proposition 5], Theorem 6.1.9 is equivalent to the following ergodicity statement.

Corollary 6.1.13. *Let (G, ρ) be an ergodic unimodular random rooted network with $\mathbb{E}[c(\rho)] < \infty$ and let $\mathfrak{F}$ be a sample of either FUSF_G or WUSF_G. Then $(T_{\mathfrak{F}}(\rho), \rho)$ is an ergodic unimodular random rooted network. Moreover, if we bias the distribution of $(G, \rho, \mathfrak{F})$ by $c_{\mathfrak{F}}(\rho)$ and let $\langle X_n \rangle_{n \geq 0}$ and $\langle X_{-n} \rangle_{n \geq 0}$ be independent random walks on $\mathfrak{F}$ started at ρ, then the stationary sequence $\big\langle (G, \langle X_{n+k} \rangle_{n \in \mathbb{Z}}, \mathfrak{F}) \big\rangle_{k \in \mathbb{Z}}$ is ergodic.*

Examples of component properties

Example 6.1.14 (Automorphism-invariant properties)**.** Let G_0 be a transitive graph, and let $\mathscr{A}$ be an automorphism-invariant set of subgraphs of G_0, that is, $\gamma\mathscr{A} = \mathscr{A}$ for any automorphism γ of G_0. Fix an arbitrary vertex v_0 of G_0 and let

$$\mathscr{A}' = \left\{ (G, v, \omega) : \begin{array}{l} \exists \text{ an isomorphism } \phi : (G, v) \to (G_0, v_0) \\ \text{such that } \phi(K_\omega(v)) \in \mathscr{A} \end{array} \right\}$$

Then $\mathscr{A}'$ is a component property such that $(G_0, v_0, \omega) \in \mathscr{A}'$ if and only if $K_\omega(v_0) \in \mathscr{A}$. Thus, Theorem 6.1.1 follows from Theorem 6.1.1 and similarly Theorems 6.1.3, 6.1.4 and 6.1.2 follow from Theorems 6.1.10, 6.1.11 and 6.1.12, respectively.

Example 6.1.15 (Intrinsic properties)**.** A graph H is said to have **volume-growth dimension** d if

$$|B_H(v, r)| = r^{d+o(1)}$$

for any (and hence all) vertices v of H. Let $p_n(\cdot, \cdot)$ denote the n-step transition probabilities of simple random walk on H. We say that H has **spectral dimension** d if

$$p_n(v, v) = n^{-d/2+o(1)}$$

for any (and hence all) vertices v of H. Lastly, recall that the **critical percolation probability** $p_c(H)$ of an infinite connected graph H is the supremum over $p \in [0, 1]$ such that independent percolation with edge probability p a.s. does not exhibit an infinite cluster. Then, under the hypotheses of Corollary 6.1.13, the volume-growth, spectral dimension of $T_{\mathfrak{F}}(\rho)$ (if they exist) and value of p_c are non-random, and consequently are a.s. the same for every tree in $\mathfrak{F}$.

Component properties can be 'extrinsic' and depend upon how the component sits inside of the base graph G.

Example 6.1.16 (Extrinsic properties)**.** A subgraph H of G is said to have **discrete Hausdorff dimension** α if

$$|B_G(v, n) \cap H| = n^{\alpha+o(1)},$$

for any (and hence all) vertices v of G. The event that a component has a particular discrete Hausdorff dimension is a component property, and consequently Theorem 6.1.9 implies that all components of the USF in a unimodular random rooted network have the same discrete Hausdorff dimension (if this dimension exists). In fact, the discrete Hausdorff dimension of every component of the USF in $\mathbb{Z}^d$ was proven to be 4 for all $d \geq 4$ by Benjamini, Kesten, Peres and Schramm [41].

Even for unimodular transitive graphs, the conclusion of Theorem 6.1.9 is strictly stronger than that of Theorem 6.1.1. This is because a component property can also depend on the whole configuration, as the following example demonstrates.

Example 6.1.17. Define $N(v,\omega,r)$ to be the number of distinct components of ω that are adjacent to the ball $B_\omega(v,r)$ of radius r about v in the intrinsic distance on $K_\omega(\rho)$. Asymptotic statements about the growth of $N(v,\omega,r)$, can be used to define component properties that depend on the entire configuration ω, e.g.

$$\mathscr{A} = \{(G,v,\omega) : N(v,\omega,r) = r^{\beta+o(1)}\}.$$

All of the examples above are what we call *tail* properties. That is, these properties can be verified by looking at all but finitely many edges of both the component and of the configuration (see the next section for the precise definition). Let us give now an interesting example of a non-tail property.

Example 6.1.18 (Non-tail property)**.** The component property

$$\mathscr{A}(n) = \left\{ (G,v,\omega) : \begin{array}{l} \text{for each connected component } K \text{ of } \omega\text{, there exists a path } \langle e_i \rangle \\ \text{in } G \text{ connecting } K_\omega(v) \text{ to } K \text{ such that at most } n \text{ of the } e_i\text{'s are} \\ \text{not contained in } \omega. \end{array} \right\}$$

is not a tail property. Benjamini, Kesten, Peres and Schramm [41] proved the remarkable result that the property $\mathscr{A}(n) \setminus \mathscr{A}(n-1)$ holds a.s. for every component of the USF of $\mathbb{Z}^d$ if and only if $4(n-1) < d \le 4n$.

Tail properties

Definition 6.1.19. We say that a component property $\mathscr{A}$ is a **tail component property** if

$$(G,v,\omega) \in \mathscr{A} \implies (G,v,\omega') \in \mathscr{A} \quad \begin{array}{l} \forall \omega' \subseteq E(G) \text{ such that } \omega \triangle \omega' \text{ and} \\ K_\omega(v) \triangle K_{\omega'}(v) \text{ are both finite.} \end{array}$$

Indistinguishability of USF components by properties that are not tail can fail without the assumption of unimodularity – see [176, Remark 3.16] and [36, Example 3.1]. However, our next theorem shows that unimodularity is not necessary for indistinguishability of WUSF components by tail properties when the WUSF components are a.s. one-ended. In [170] it is shown that the last condition holds in every transient transitive graph. The following theorem, which is used in the proof of Theorem 6.1.9, implies that WUSF components are indistinguishable by tail properties in any transient transitive graph (not necessarily unimodular).

Theorem 6.1.20. *Let (G,ρ) be a stationary random network and let $\mathfrak{F}$ be a sample of WUSF_G. Suppose that every component of $\mathfrak{F}$ is one-ended almost surely. Then for every tail component property $\mathscr{A}$, either every connected component of $\mathfrak{F}$ has property $\mathscr{A}$ or none of the connected components of $\mathfrak{F}$ have property $\mathscr{A}$ almost surely.*

Sharpness

We present a construction showing that the condition $\mathbb{E}[c(\rho)] < \infty$ in Theorem 6.1.9 is indeed necessary. For integers n and $k > 2$, denote by $T_n(k)$ the finite network on a binary tree of height n such that edges at distance h from the leaves have conductance k^h. Choose a uniform random root in $T_n(k)$ and take n to ∞ while keeping k fixed. The limit of this process can be seen to be the transient unimodular random rooted network $T(k)$ in which the underlying graph is the canopy tree [5] and edges of distance h from the leaves have conductance k^h.

Consider the finite network G_n obtained by gluing a copy of $T_n(3)$ and a copy of $T_n(4)$ at their leaves in such a way that the resulting network G_n is planar, and let ρ_n a uniformly chosen root vertex of G_n. Then the randomly rooted graphs (G_n, ρ_n) converge to a unimodular random rooted network which is formed by gluing a copy of $T(3)$ and a copy of $T(4)$ at their leaves. It can easily be seen via Wilson's algorithm (see Section 6.4) that the WUSF will contain precisely two one-ended components corresponding to the two infinite rays of $T(3)$ and $T(4)$. These two trees are clearly distinguishable from each other by measuring the frequency of edges with conductances 3 or 4 on their infinite ray. This example can be made into a graph rather than a network simply by replacing an edge with conductance k^h by k^h parallel edges.

It is also possible to construct a unimodular random rooted network on which there are infinitely many WUSF components almost surely and every cluster is distinguishable from every other cluster. Let G be a 3-regular tree, and let $\mathfrak{F}_1$ be a sample of WUSF_T. For every component T of $\mathfrak{F}_1$, let $U(T)$ be i.i.d. uniform $[0, 1]$. For each edge e of G that is contained in $\mathfrak{F}_1$, let $T(e)$ be the component of $\mathfrak{F}_1$ containing e. Define conductances on G by, for each edge e of G, setting $c(e) = 1$ if $e \notin \mathfrak{F}_1$ and otherwise setting

$$c(e) = \exp\big((1 + U(T)) \times |\text{The finite component of } T(e) \setminus e|\big) \ .$$

These strong drifts ensure that a random walk on the network (G, c) will eventually remain in a single component of $\mathfrak{F}_1$. Running Wilson's algorithm on the network (G, c) to sample a copy $\mathfrak{F}_2$ of $\mathsf{WUSF}_{(G,c)}$, we see that the components of $\mathfrak{F}_2$ correspond to the components of $\mathfrak{F}_1$. We can distinguish these components from each other by observing the rate of growth of the conductances along a ray in each component.

6.2 Indistinguishability of FUSF components

Our goal in this section is to prove Theorem 6.1.9 for the FUSF on a unimodular random rooted network (G, ρ) when the measures FUSF_G and WUSF_G are distinct.

6.2.1 Cycle breaking in the FUSF

Let G be a finite network. For each spanning tree t of G and oriented edge e of G that is not a self-loop, we define the **direction** $D(e) = D(t, e)$ to be the first edge in the unique simple path from e^- to e^+ in t.

Lemma 6.2.1. *Let G be an infinite network with exhaustion $\langle V_n \rangle_{n \geq 1}$ and let T_n be a sample of UST_{G_n} for each n. Then for every oriented edge e of G, the random variables $(T_n, D(T_n, e))$ converge in distribution to some limit $(\mathfrak{F}, D(e))$, where $D(e)$ is an edge adjacent to e^- and the marginal distribution of $\mathfrak{F}$ is given by FUSF_G.*

Proof. Since the distribution of T_n converges to FUSF_G, it suffices to show that the conditional probabilities

$$\mathbb{P}(D(T_n, e) = d \mid f_1, \ldots, f_k \in T_n,\ h_1, \ldots h_l \notin T_n) \tag{6.2.1}$$

converge, where $d = (d^-, d^+)$ is any oriented edge with $d^- = e^-$ and $F = \{f_1, \ldots, f_k\}$ and $H = \{h_1, \ldots, h_l\}$ are any two finite collections of edges in G for which the event $\mathscr{A} = \{f_1, \ldots, f_k \in \mathfrak{F},\ g_1, \ldots g_l \notin \mathfrak{F}\}$ has non-zero probability. Fix such d, F and H. If F includes a path from e^- to e^+ then $D(T_n, e)$ is determined and there is convergence in (6.2.1). Also, if $d \in H$ then (6.2.1) is zero. So let us assume now that neither is the case.

Let us first explain why convergence holds in (6.2.1) when $F = H = \emptyset$. In that case, by Kirchhoff's effective resistance formula (see [44, Theorem 4.1]), the probability that the unique path between e^- to e^+ in T_n goes through d equals the amount of current on the edge d when a unit current flows from e^- to e^+ in the network G_n. It is well known that this quantity converges as $n \to \infty$ to the current passing through d in the *free unit current flow* from e^- to e^+ in G [173, Proposition 9.1].

When F and H are non-empty, we take n to be sufficiently large such that the edges e, d and all the edges of F and H are contained in G_n, and that the event $\mathscr{A}_n = \{F \subset T_n,\ H \cap T_n = \emptyset\}$ has non-zero probability (i.e, that $G_n \setminus H$ is connected and there are no cycles in F). We write $(G_n - H)/F$ for the network formed from G_n by deleting each edge $h \in H$ and contracting each edge $f \in F$. By the Markov property (see Section 6.1.2), the laws of T_n conditioned on $\mathscr{A}_n$ and of $\mathfrak{F}$ conditioned on $\mathscr{A}$ can be sampled from by taking the union of F with a sample of the UST of $(G_n - H)/F$ or the FUSF of $(G - H)/F$ respectively. Thus, the same argument as above works when $d \notin F$ and shows that the limit of (6.2.1) is equal to the current passing through d in the free unit current flow from e^- to e^+ in $(G - H)/F$.

Finally suppose that $d \in F$. Let V_d be the set of vertices connected to e^- by a simple path in F passing through d, and let E_d be the set of oriented edges with tail in V_d. By our previous discussion, (6.2.1) equals the sum of the currents flowing through the edges of E_d in the unit current flow from e^- to e^+ in $(G_n - H)/F$. As before, [173, Proposition 9.1] shows that this quantity converges to the corresponding sum of currents in the free unit current flow from e^- to e^+ in $(G - H)/F$. □

For each oriented edge e of G and FUSF_G-a.e. spanning forest f of G, we define the **update** $U(f, e)$ as follows. If e is either a self-loop or already contained in f, set $U(f, e) = f$. Otherwise, sample $D(e)$ from its conditional distribution given $\mathfrak{F} = f$, and set $U(f, e) = \mathfrak{F} \cup \{e\} \setminus D(e)$. It seems likely that this conditional distribution is concentrated on a point.

Question 6.2.2. *Let G be a network and let e be an edge of G. Does $U(f, e)$ coincide FUSF_G-a.e. with some measurable function of f?*

If any additional randomness is required to perform an update, it will always be taken to be independent of any other random variables considered.

Lemma 6.2.3. *Let G be a network and $\mathfrak{F}$ be a sample of FUSF_G. Let v be a vertex of G and let E be an element of the set $\{e : e^- = v\}$ chosen independently of $\mathfrak{F}$ and with probability proportional to its conductance. Then $U(\mathfrak{F}, E)$ and $\mathfrak{F}$ have the same distribution.*

Proof. Let $\langle V_n \rangle_{n \geq 0}$ be an exhaustion of G and let T_n be a sample of the UST on G_n for each n. We may assume that G_n contains v and every edge adjacent to v for all $n \geq 1$. We define the update $U(t, e)$ of a spanning tree t of G_n at the oriented edge e to be

$$U(t, e) = t \cup \{e\} \setminus D(t, e).$$

Since $U(T_n, E)$ converges to $U(\mathfrak{F}, E)$ in distribution, and so it suffices to verify that $U(T_n, E) \stackrel{d}{=} T_n$ for each $n \geq 0$: this may be done by checking that UST_{G_n} satisfies the detailed balance equations for the Markov chain on the set of spanning trees of G with transition probabilities

$$p(t_1, t_2) = \frac{1}{c(v)} c(\{e : e^- = v \text{ and } U(t_1, e) = t_2\}).$$

This simple calculation is carried out in [127, Lemma 6]. □

This has the following immediate consequence.

Corollary 6.2.4 (Update Tolerance for the FUSF)**.** *Let G be a network. Fix an edge e, and let FUSF_G^e denote the joint distribution of a sample $\mathfrak{F}$ of FUSF_G and of the update $U(\mathfrak{F}, e)$. Then*

$$\mathsf{FUSF}_G(\mathfrak{F} \in \mathscr{A}) \geq \frac{c(e)}{c(e^-)} \mathsf{FUSF}_G^e(U(\mathfrak{F}, e) \in \mathscr{A})$$

Proof. Lemma 6.2.3 implies that

$$\begin{aligned} \mathsf{FUSF}_G(\mathfrak{F} \in \mathscr{A}) &= \frac{1}{c(e^-)} \sum_{\hat{e}^- = e^-} c(\hat{e}) \mathsf{FUSF}_G^{\hat{e}}(U(\mathfrak{F}, \hat{e}) \in \mathscr{A}) \\ &\geq \frac{c(e)}{c(e^-)} \mathsf{FUSF}_G^e(U(\mathfrak{F}, e) \in \mathscr{A}). \end{aligned}$$

□

6.2.2 All FUSF components are transient and infinitely-ended

A **weighted tree** is a network whose underlying graph is a tree. Recall that a branch of an infinite tree T is an infinite connected component of $T \setminus v$ for some vertex v.

Lemma 6.2.5. *Let (T, ρ) be a unimodular random rooted weighted tree that is transient with positive probability. On the event that T is transient, T a.s. does not have any recurrent branches.*

Proof. For each vertex u of T, let $V(u)$ be the set of vertices $v \neq u$ such that the component containing u in $T \setminus v$ is recurrent. We first claim that if T is a transient weighted tree and u is a vertex of T, then $V(u)$ is either empty, or a finite simple path starting from a neighbour of u in T, or an infinite transient ray (i.e. a ray such that the sum of the edge resistances along the ray is finite) starting from a neighbour of u in T. First, Rayleigh's monotonicity principle implies that if $v \in V(u)$ then every other vertex on the unique path from u to v in T is also in $V(u)$. Second, if there exist $v_1, v_2 \in V(u)$ which do not lie on a simple path from u in T, then the component of u in $T \setminus v_1$ and the component of u in $T \setminus v_2$ are both recurrent and have all of T as their union implying that T is recurrent.

Thus, if $V(u)$ is not empty, it must be a finite path or ray in T starting at u. Let us rule out the case that $V(u)$ is a recurrent infinite ray. Assume that $V(u)$ is infinite and denote this ray by $(v_0, v_1, v_2, \ldots)$ with $v_0 = u$. For each integer $n \geq 0$ the component of $T \setminus v_{n+1}$ containing v_n is recurrent and so, by Rayleigh's monotonicity principle, the components of $T \setminus v_n$ that do not contain v_{n+1} are all recurrent. Thus T is decomposed to the union of the ray $(v_0, v_1, \ldots)$ and a collection of recurrent branches hanging on this ray. If the ray $V(u)$ is a recurrent, we conclude that T is recurrent as well.

Suppose for contradiction that with positive probability T is transient but there exists an edge e such that the component of ρ in $T \setminus e$ is infinite and recurrent. Take $\varepsilon > 0$ sufficiently small so that this edge e may be taken to have $c(e) \geq \varepsilon$ with positive probability. Denote the event that such an edge exists by $\mathscr{B}_\varepsilon$. We will show that this contradicts the Mass-Transport Principle by exhibiting a mass transport such that every vertex sends a mass of at most one but some vertices receive infinite mass on the event $\mathscr{B}_\varepsilon$.

From each vertex u such that $V(u)$ is finite and non-empty, send mass one to the vertex v in $V(u)$ that is farthest from u in T. From each vertex u such that $V(u)$ is a transient ray, send mass one to the end-point $v = e^-$ of the last edge e in the path from u spanned by $V(u)$ such that $c(e) \geq \varepsilon$. If $V(u)$ is empty, u sends no mass. Clearly every vertex sends a total mass of at most one. However, on the event $\mathscr{B}_\varepsilon$, the vertex v that ρ sends mass to receives infinite mass. Indeed, every vertex in the infinite recurrent component of $T \setminus v$ containing ρ sends mass one to v, contradicting the Mass-Transport Principle. □

Lemma 6.2.6. *Let G be an infinite network and let $\mathfrak{F}$ be a sample of FUSF_G. If with positive probability $\mathfrak{F}$ has a recurrent component and a transient component with a non-empty core, then with positive probability $\mathfrak{F}$ has a component that is a transient tree with a recurrent branch. This holds both when edges of the trees are given the conductances inherited from G and when they are given unit conductances.*

Proof. We consider the case that the edges of the trees are given the conductances inherited from G, the other case is similar. If two such components exist, then one can find a finite path starting at a vertex of a recurrent component T and ending in a vertex of the core of a transient component T'. Moreover, by taking the shortest such path, the starting vertex is the only vertex in T and the end vertex is the only vertex in $\mathrm{core}(T')$.

Thus, there exists a non-random finite simple path $\gamma = \langle \gamma_i \rangle_{i=0}^n$ in G such that the following event, denoted $\mathscr{B}(\gamma)$, holds with positive probability:

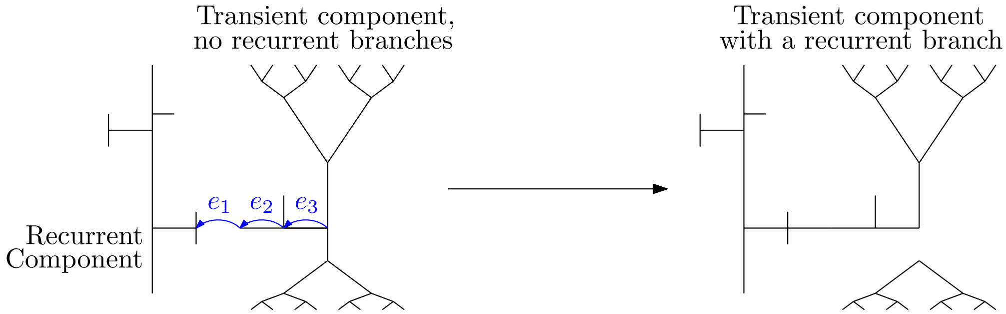

Figure 6.1: When recurrent components and transient components with non-empty cores and no recurrent branches coexist, a finite sequence of updates can create a transient component with a recurrent branch.

- $T_{\mathfrak{F}}(\gamma_0)$ is recurrent,
- $T_{\mathfrak{F}}(\gamma_i) \neq T_{\mathfrak{F}}(\gamma_0)$ for $0 < i \leq n$,
- $T_{\mathfrak{F}}(\gamma_n)$ is transient and
- $\gamma_n \in \mathrm{core}(T_{\mathfrak{F}}(\gamma_n))$ and it is the only such vertex in γ.

For each $1 \leq i \leq n$, let e_i be an oriented edge of G with $e_i^- = \gamma_i$ and $e_i^+ = \gamma_{i-1}$. Define the forests $\langle \mathfrak{F}_i \rangle_{i=0}^n$ by setting $\mathfrak{F}_0 = \mathfrak{F}$ and recursively,

$$\mathfrak{F}_i = U(\mathfrak{F}_{i-1}, e_i)\,, \qquad i = 1, \ldots, n\,.$$

We claim that on the event $\mathscr{B}(\gamma)$, at least one of the two forests $\mathfrak{F}_0$ or $\mathfrak{F}_n$ contains a transient tree with a recurrent branch. If $\mathfrak{F}_0$ contains such a tree we are done, so suppose not. We claim that in this case $T_{\mathfrak{F}_n}(\gamma_n)$ is a transient tree with a recurrent branch. Indeed, at each step of the process we are add the edge e_i and remove some other edge adjacent to γ_i in $T_{\mathfrak{F}_{i-1}}(\gamma_i)$, so that $T_{\mathfrak{F}_n}(\gamma_n)$ contains the tree $T_{\mathfrak{F}_0}(\gamma_0)$ (since $\gamma_i \notin T_{\mathfrak{F}_0}(\gamma_0)$ for $i \geq 1$) and the path $e_1, \ldots, e_n$. Moreover, since $\gamma_i \notin \mathrm{core}(T_{\mathfrak{F}_0}(\gamma_n))$ for all $0 \leq i < n$, the tree $T_{\mathfrak{F}_n}(\gamma_n)$ contains a branch of $T_{\mathfrak{F}_0}(\gamma_n)$ and is therefore transient by our assumption. Thus, removing γ_1 from the transient tree $T_{\mathfrak{F}_n}(\gamma_n)$ yields the recurrent branch $T_{\mathfrak{F}_0}(\gamma_0)$ as required.

Denote by $\mathscr{E}$ the set of subgraphs of G that are transient trees with a recurrent branch. We have shown that $\mathbb{P}(\mathfrak{F}_0 \in \mathscr{E}) + \mathbb{P}(\mathfrak{F}_n \in \mathscr{E}) > 0$, while by update-tolerance (Corollary 6.2.4)

$$\mathbb{P}(\mathfrak{F}_0 \in \mathscr{E}) \geq \Big(\prod_{i=1}^n \frac{c(e_i)}{c(e_i^-)}\Big)\mathbb{P}(\mathfrak{F}_n \in \mathscr{E})$$

so that $\mathbb{P}(\mathfrak{F}_0 \in \mathscr{E}) > 0$ as claimed. □

Proof of Theorem 6.1.12. We may assume that (G, ρ) is ergodic, see Section 6.1.2. We apply [7, Proposition 4.9 and Theorem 6.2] to deduce that whenever (T, ρ) is an infinite unimodular random rooted (unweighed) tree with $\mathbb{E}[c(\rho)] < \infty$, the event that T is infinitely-ended and the event that

T is transient coincide up to a null set and, moreover, T has positive probability to be transient and infinitely-ended if and only if $\mathbb{E}[\deg_T(\rho)] > 2$. The expected degree of the WUSF is 2 in any unimodular random rooted network, and since the FUSF_G stochastically dominates WUSF_G, the assumption that $\mathsf{FUSF}_G \neq \mathsf{WUSF}_G$ implies that $\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] > 2$. Let $M > 0$ and let $\mathfrak{F}'$ be the forest obtained by deleting from $\mathfrak{F}$ every edge e such that $\max(\deg_{\mathfrak{F}}(e^-), \deg_{\mathfrak{F}}(e^+)) \geq M$. If M is sufficiently large then $\mathbb{E}[\deg_{\mathfrak{F}'}(\rho)] > 2$ by the monotone convergence theorem. It follows by the above that $T_{\mathfrak{F}'}(\rho)$ is infinitely-ended and and transient (when given unit conductances) with positive probability, and consequently that the same holds for $T_{\mathfrak{F}}(\rho)$ by Rayleigh monotonicity. Ergodicity of $(G, \rho, \mathfrak{F})$ then implies that the forest $\mathfrak{F}$ contains a component that is infinitely-ended and transient (when given unit conductances) a.s.

Assume for contradiction that with positive probability $\mathfrak{F}$ has a component that is finitely-ended, or equivalently a component that is recurrent when given unit conductances. Lemma 6.2.6 then implies that with positive probability $\mathfrak{F}$ has a transient component with a recurrent branch (when all components are given unit conductances), contradicting Lemma 6.2.5.

Thus, we have that all components of $\mathfrak{F}$ are a.s. infinitely-ended and are transient when given unit conductances. It follows from [7, Proposition 4.10] that every component is also a.s. transient when given the conductances inherited from G. $\square$

6.2.3 Pivotal edges for the FUSF

Let G be a network, let $\mathfrak{F}$ be a sample of FUSF_G and let $\mathscr{A}$ be a component property. We say that an oriented edge e of G is a δ-**additive pivotal** for a vertex v if

1. $e^+ \in T_{\mathfrak{F}}(v)$ and $e^- \notin T_{\mathfrak{F}}(v)$ and,

2. given $\mathfrak{F}$, the components $T_{U(\mathfrak{F},e)}(v)$ and $T_{\mathfrak{F}}(v)$ have different types with probability at least δ.

We say that an oriented edge e is a δ-**subtractive pivotal** for v if

1. $e^- \in T_{\mathfrak{F}}(v)$ and $e^+ \notin T_{\mathfrak{F}}(v)$ and,

2. given $\mathfrak{F}$, the components $T_{U(\mathfrak{F},e)}(v)$ and $T_{\mathfrak{F}}(v)$ have different types with probability at least δ.

We emphasize that when we say "with probability at least δ" above, this is over the randomness of $U(\mathfrak{F}, e)$, rather than of $\mathfrak{F}$.

Lemma 6.2.7. *Let G be a network and let $\mathfrak{F}$ be a sample of FUSF_G. Assume that a.s. all the components of $\mathfrak{F}$ are transient trees with non-empty cores and that with positive probability $\mathfrak{F}$ has components of both types $\mathscr{A}$ and $\neg\mathscr{A}$. Then for some small $\delta > 0$, with positive probability there exists a vertex v and an edge e such that $v \in \mathrm{core}(F)$ and e is a δ-pivotal for v.*

Proof. We argue similarly to Lemma 6.2.6. Due to the assumptions of this lemma, there must be a component T of type $\mathscr{A}$ and a component T' of type $\neg\mathscr{A}$ and an edge $e \notin \mathfrak{F}$ connecting them.

So we may form a path starting with e that ends in a core vertex of T' such that all edges of the path except for e are in T', and the last vertex of the path is the only vertex in $\mathrm{core}(T')$.

Hence, there exists a non-random simple path $\gamma = \langle \gamma_i \rangle_{i=0}^n$ in G such that the following event, denoted $\mathscr{B}(\gamma)$, holds with positive probability:

- $T_{\mathfrak{F}}(\gamma_0)$ has type $\mathscr{A}$,
- $T_{\mathfrak{F}}(\gamma_i) = T_{\mathfrak{F}}(\gamma_j) \neq T_{\mathfrak{F}}(\gamma_0)$ for all $0 < i \le j \le n$,
- $T_{\mathfrak{F}}(\gamma_n)$ has type $\neg\mathscr{A}$ and,
- $\gamma_n \in \mathrm{core}(T_{\mathfrak{F}}(\gamma_n))$ and it is the only such vertex in γ.

For each $1 \le i \le n$, let e_i be an oriented edge of G with $e_i^- = \gamma_i$ and $e_i^+ = \gamma_{i-1}$. Define the forests $\langle \mathfrak{F}_i \rangle_{i=0}^n$ by setting $\mathfrak{F}_0 = \mathfrak{F}$ and, recursively,

$$\mathfrak{F}_i = U(\mathfrak{F}_{i-1}, e_i)\,, \qquad i = 1, \ldots, n.$$

We claim that given $\mathscr{B}(\gamma)$ there exists some small $\delta > 0$ such that either one of the edges e_i is a δ-additive pivotal for γ_0 in the forest $\mathfrak{F}_{i-1}$, or one of the edges e_i is a δ-subtractive pivotal for γ_n in the forest $\mathfrak{F}_{i-1}$. Indeed, if there is $1 \le i \le n$ such that

$$\mathbb{P}\big(T_{\mathfrak{F}_i}(\gamma_0) \in \neg\mathscr{A} \mid \mathfrak{F}_{i-1}\big) > 0\,,$$

(i.e., the component of γ_0 changes type with positive probability in the transition from $\mathfrak{F}_{i-1}$ to $\mathfrak{F}_i$), then for the first such i, the edge e_i is a δ-additive pivotal for γ_0 in $\mathfrak{F}_{i-1}$ (since the cluster of γ_0 only grows) where $\delta > 0$ is the conditional probability above.

If this does not occur, then a.s. $T_{\mathfrak{F}_n}(\gamma_0)$ is of type $\mathscr{A}$. However, $T_{\mathfrak{F}_n}(\gamma_0) = T_{\mathfrak{F}_n}(\gamma_n)$ and $T_{\mathfrak{F}_0}(\gamma_n)$ is of type $\neg\mathscr{A}$, so there is the first $1 \le i \le n$ in which $T_{\mathfrak{F}_i}(\gamma_n) \in \mathscr{A}$ (i.e., the component of γ_n changes type in the transition from $\mathfrak{F}_{i-1}$ to $\mathfrak{F}_i$). For this i we have e_i is a δ'-subtractive pivotal for γ_n in $\mathfrak{F}_{i-1}$ with

$$\delta' = \mathbb{P}\big(T_{\mathfrak{F}_i}(\gamma_n) \in \mathscr{A} \mid \mathfrak{F}_{i-1}\big) > 0\,.$$

Let $\mathscr{E}_\delta$ be the event that there exists a vertex v such that $v \in \mathrm{core}(\mathfrak{F})$ and there exists an edge e that is a δ-additive pivotal or δ-subtractive pivotal for v in $\mathfrak{F}$. We proved that for some small $\delta > 0$ we get $\sum_{i=0}^n \mathbb{P}(\mathfrak{F}_i \in \mathscr{E}_\delta) > 0$. However, by update tolerance (Corollary 6.2.4) it follows that

$$\mathbb{P}(\mathfrak{F}_0 \in \mathscr{E}) \ge \Big(\prod_{j=1}^{i} \frac{c(e_j)}{c(e_j^-)}\Big)\mathbb{P}(\mathfrak{F}_i \in \mathscr{E})\,.$$

Hence $\mathbb{P}(\mathfrak{F}_0 \in \mathscr{E}_\delta) > 0$ for some small $\delta > 0$ as claimed. □

6.2.4 Proof of Theorem 6.1.9 for the FUSF

Our goal in this section is to prove the following theorem.

Theorem 6.2.8. *Let (G, ρ) be a unimodular random network with $\mathbb{E}[c(\rho)] < \infty$, and let $\mathfrak{F}$ be a sample of FUSF_G. On the event that $\mathsf{FUSF}_G \neq \mathsf{WUSF}_G$, we have that for every component property $\mathscr{A}$, either every connected component of $\mathfrak{F}$ has property $\mathscr{A}$ or none of the connected components of $\mathfrak{F}$ have property $\mathscr{A}$ almost surely.*

We follow the strategy of Lyons and Schramm [176] while making the changes necessary to use update-tolerance.

Proof of Theorem 6.2.8. Let (G, ρ) be a unimodular random rooted network with $\mathbb{E}[c(\rho)] < \infty$, let $\mathfrak{F}$ be a sample of FUSF_G and let $\mathscr{A}$ be a component property. Let $\langle X_n \rangle_{n \in \mathbb{Z}}$ be a bi-infinite random walk on $\mathfrak{F}$ started at ρ (that is, the concatenation of two independent random walks starting at ρ, as in Proposition 6.1.6). Conditioned on the random walk $\langle X_n \rangle_{n \in \mathbb{Z}}$, let e_n be oriented edges chosen uniformly and independently from the set of edges at distance at most r from X_n in G. Finally, let $\{U(\mathfrak{F}, e) : e \in E\}$ be updates of $\mathfrak{F}$ at each edge e of G, sampled independently of each other and of $\langle X_n \rangle_{n \in \mathbb{Z}}$ and $\langle e_n \rangle_{n \in \mathbb{Z}}$ conditional on $(G, \rho, \mathfrak{F})$. We bias by $c_{\mathfrak{F}}(\rho)$ so that, by [7, Theorem 4.1],

$$(G, \langle X_n \rangle_{n \in \mathbb{Z}}, \mathfrak{F}) \stackrel{d}{=} (G, \langle X_{n+k} \rangle_{n \in \mathbb{Z}}, \mathfrak{F}) \quad \forall\, k \in \mathbb{Z}. \tag{6.2.2}$$

Let $\widehat{\mathbb{P}}$ denote joint distribution of the random variables (G, ρ), $\mathfrak{F}$, $\langle X_n \rangle_{n \in \mathbb{Z}}$, $\langle e_n \rangle_{n \in \mathbb{Z}}$, and $\{U(\mathfrak{F}, e) : e \in E\}$ under this biasing, and let $\widehat{\mathbb{P}}_{(G,\rho)}$ denote the conditional distribution given (G, ρ) of $\mathfrak{F}$, $\langle X_n \rangle_{n \in \mathbb{Z}}$, $\langle e_n \rangle_{n \in \mathbb{Z}}$ and $\{U(\mathfrak{F}, e) : e \in E\}$ under the same biasing.

By Theorem 6.1.12, every component of $\mathfrak{F}$ is a.s. transient and infinitely-ended. By Lemma 6.2.7, there exists $\delta > 0$ such that with positive probability $\rho \in \operatorname{core}(\mathfrak{F})$ and there exists a either a δ-additive or δ-subtractive pivotal edge e for ρ in $\mathfrak{F}$. By decreasing δ if necessary, it follows that there exists an integer r such that with positive probability $\rho \in \operatorname{core}(\mathfrak{F})$ and there exists a δ-pivotal edge for ρ in $\mathfrak{F}$ such that ρ and e are at graph distance at most r in G and $c(e)/c(e-) \geq \delta$.

Conditional on (G, ρ), for each edge e in G and $n \in \mathbb{Z}$, denote by $\mathscr{E}_e^n$ the event that $e_n = e$ and that the trace $\{X_{n+k}\}_{k \in \mathbb{Z}}$ is disjoint from the components of $\mathfrak{F} \setminus X_n$ containing e^- and e^+. For every essential spanning forest f of G and $n \in \mathbb{Z}$, we have

$$\widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n \mid \mathfrak{F} = f) = \widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n \mid U(\mathfrak{F}, e) = f)\,.$$

Thus, for every event $\mathscr{B} \subseteq \{0,1\}^{E(G)}$ such that $\widehat{\mathbb{P}}_{(G,\rho)}(\mathfrak{F} \in \mathscr{B}) > 0$, we have that

$$\begin{aligned}
\widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n \cap \{\mathfrak{F} \in \mathscr{B}\}) &= \widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n \mid \mathfrak{F} \in \mathscr{B}) \widehat{\mathbb{P}}_{(G,\rho)}(\mathfrak{F} \in \mathscr{B}) \\
&= \widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n \mid \{U(\mathfrak{F}, e) \in \mathscr{B}\}) \widehat{\mathbb{P}}_{(G,\rho)}(\mathfrak{F} \in \mathscr{B}) \\
&= \frac{\widehat{\mathbb{P}}_{(G,\rho)}(\mathfrak{F} \in \mathscr{B})}{\widehat{\mathbb{P}}_{(G,\rho)}(U(\mathfrak{F}, e) \in \mathscr{B})} \widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n \cap \{U(\mathfrak{F}, e) \in \mathscr{B}\}) \\
&= \frac{\mathsf{FUSF}_G[c_{\mathfrak{F}}(\rho) \mathbb{1}(\mathfrak{F} \in \mathscr{B})]}{\mathsf{FUSF}_G^e[c_{\mathfrak{F}}(\rho) \mathbb{1}(U(\mathfrak{F}, e) \in \mathscr{B})]} \widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n \cap \{U(\mathfrak{F}, e) \in \mathscr{B}\}).
\end{aligned}$$

Observe that if e does not have ρ as an endpoint then, by Lemma 6.2.3,

$$\begin{aligned}\mathsf{FUSF}_G[c_{\mathfrak{F}}(\rho)\mathbb{1}(\mathfrak{F}\in\mathscr{B})] &= \frac{1}{c(e^-)}\sum_{e'^-=e^-} c(e')\mathsf{FUSF}_G^{e'}[c_{U(\mathfrak{F},e')}(\rho)\mathbb{1}(U(\mathfrak{F},e')\in\mathscr{B})]\\ &\geq \frac{c(e)}{c(e^-)}\mathsf{FUSF}_G^{e}[c_{U(\mathfrak{F},e)}(\rho)\mathbb{1}(U(\mathfrak{F},e)\in\mathscr{B})]\\ &= \frac{c(e)}{c(e^-)}\mathsf{FUSF}_G^{e}[c_{\mathfrak{F}}(\rho)\mathbb{1}(U(\mathfrak{F},e)\in\mathscr{B})],\end{aligned}$$

and so, for every edge e of G not having ρ as an endpoint,

$$\widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n\cap\{\mathfrak{F}\in\mathscr{B}\}) \geq \frac{c(e)}{c(e^-)}\widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n\cap\{U(\mathfrak{F},e)\in\mathscr{B}\})\,. \tag{6.2.3}$$

Update-tolerance also implies that (6.2.3) holds trivially when $\widehat{\mathbb{P}}_{(G,\rho)}(\mathfrak{F}\in\mathscr{B})=0$.

Fix $\varepsilon>0$, and let R be sufficiently large such that there exists an event $\mathscr{A}'$ that is measurable with respect to (G,ρ) and $\mathfrak{F}\cap B_G(\rho,R)$ and satisfies $\widehat{\mathbb{P}}(\mathscr{A}\,\triangle\,\mathscr{A}')\leq\varepsilon$. Such an $\mathscr{A}'$ exists by the assumption that $\mathscr{A}$ is measurable. Define the disjoint unions

$$\mathscr{E}^n := \bigcup_{c(e)/(e^-)\geq\delta}\mathscr{E}_e^n \quad\text{ and }\quad \mathscr{E}_R^n := \bigcup_{e^-\notin B_G(\rho,R)\,,c(e)/c(e^-)\geq\delta}\mathscr{E}_e^n.$$

Condition on (G,ρ), and let

$$\mathscr{B}=\{\omega\in\{0,1\}^E : (G,\rho,\omega)\in\mathscr{A}'\setminus\mathscr{A}\}.$$

Summing over (6.2.3) with this $\mathscr{B}$ yields that, for every $R\geq 1$,

$$\begin{aligned}\widehat{\mathbb{P}}_{(G,\rho)}(\mathfrak{F}\in\mathscr{B}) &\geq \widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_R^n\cap\{\mathfrak{F}\in\mathscr{B}\})\\ &\geq \delta\widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_R^n\cap\{U(\mathfrak{F},e_n)\in\mathscr{B}\})\end{aligned}$$

and hence, taking expectations,

$$\widehat{\mathbb{P}}((G,\rho,\mathfrak{F})\in\mathscr{A}'\setminus\mathscr{A}) \geq \delta\widehat{\mathbb{P}}\big(\mathscr{E}_R^n\cap\{(G,\rho,\mathfrak{F})\in\mathscr{A}'\setminus\mathscr{A})\}\big).$$

By the definition of $\mathscr{A}'$ we have that

$$\mathscr{E}_R^n\cap\{(G,\rho,U(\mathfrak{F},e_n))\in\mathscr{A}'\} = \mathscr{E}_R^n\cap\{(G,\rho,\mathfrak{F})\in\mathscr{A}'\},$$

and so

$$\widehat{\mathbb{P}}((G,\rho,\mathfrak{F})\in\mathscr{A}'\setminus\mathscr{A}) \geq \delta\widehat{\mathbb{P}}\Big(\mathscr{E}_R^n\cap\big\{(G,\rho,\mathfrak{F})\in\mathscr{A}'\big\}\cap\big\{(G,\rho,U(\mathfrak{F},e_n))\in\neg\mathscr{A}\big\}\Big).$$

Let $\mathscr{P}_n$ denote the event that e_n is either a δ-additive or δ-subtractive pivotal edge for X_n. On the event $\mathscr{E}_R^n$, the vertices ρ and X_n are in the same component of $U(\mathfrak{F}, e_n)$, so that, on the event $\mathscr{P}_n \cap \mathscr{E}_R^n \cap \{(G, \rho, \mathfrak{F}) \in \mathscr{A}\}$, we have that $(G, \rho, U(\mathfrak{F}, e_n)) \in \neg\mathscr{A}$ with probability at least δ. Thus,

$$\begin{aligned}\widehat{\mathbb{P}}((G, \rho, \mathfrak{F}) \in \mathscr{A}' \setminus \mathscr{A}) &\geq \delta^2 \widehat{\mathbb{P}}(\mathscr{E}_R^n \cap \{(G, \rho, \mathfrak{F}) \in \mathscr{A}'\} \cap \mathscr{P}_n) \\ &\geq \delta^2 \widehat{\mathbb{P}}(\mathscr{E}_R^n \cap \{(G, \rho, \mathfrak{F}) \in \mathscr{A}\} \cap \mathscr{P}_n) - \delta^2 \varepsilon\,,\end{aligned}$$

by definition of $\mathscr{A}'$. Since $\langle X_n \rangle_{n\geq 0}$ is transient, we can take n to be sufficiently large that $\widehat{\mathbb{P}}(\{(G, \rho, \mathfrak{F}) \in \mathscr{A}\} \cap \mathscr{E}^n \setminus \mathscr{E}_R^n) \leq \varepsilon$. Thus, for such n,

$$\widehat{\mathbb{P}}((G, \rho, \mathfrak{F}) \in \mathscr{A}' \setminus \mathscr{A}) \geq \delta^2 \widehat{\mathbb{P}}(\mathscr{E}^n \cap \{(G, \rho, \mathfrak{F}) \in \mathscr{A}\} \cap \ \mathscr{P}_n) - 2\delta^2 \varepsilon$$

since $\mathscr{A}$ is a component property. Stationarity of $\langle (G, \mathfrak{F}, \langle X_{n+k} \rangle_{k\in\mathbb{Z}}) \rangle_{n\in\mathbb{Z}}$ implies that $\widehat{\mathbb{P}}(\mathscr{E}^n \cap \{(G, \rho, \mathfrak{F}) \in \mathscr{A}\} \cap \mathscr{P}_n)$ does not depend on n, so that it suffices to show it is positive to obtain a contradiction by choosing $\varepsilon > 0$ sufficiently small.

As mentioned earlier, with positive probability $\rho \in \mathrm{core}(\mathfrak{F})$ and there exists a δ-pivotal edge e in $\mathfrak{F}$ at distance at most r from ρ such that $c(e)/(e^-) \geq \delta$. Hence, either

$$\widehat{\mathbb{P}}(\{\rho \in \mathrm{core}(\mathfrak{F})\} \cap \{(G, \rho, \mathfrak{F}) \in \mathscr{A}\} \cap \mathscr{P}_0) > 0$$

or

$$\widehat{\mathbb{P}}(\{\rho \in \mathrm{core}(\mathfrak{F})\} \cap \{(G, \rho, \mathfrak{F}) \in \neg\mathscr{A}\} \cap \mathscr{P}_0) > 0.$$

Since $\neg\mathscr{A}$ is also a component property, we may assume without loss of generality that the former occurs. Conditioned on the events $\{\rho \in \mathrm{core}(\mathfrak{F})\}, \{(G, \rho, \mathfrak{F}) \in \mathscr{A}\}$ and $\mathscr{P}_0$, the event $\mathscr{E}^n$ is the event that two independent random walks from ρ stay within the components of $\mathfrak{F} \setminus \rho$ that do not contain e_0^- or e_0^+. This occurs with positive probability since every infinite component of $T_{\mathfrak{F}}(\rho) \setminus \rho$ is transient by Theorem 6.1.12 and Lemma 6.2.5, concluding the proof. □

6.3 The FUSF is either connected or has infinitely many components

Our goal in this section is to prove Theorems 6.1.10 and 6.1.11. Let (G, ρ) be a unimodular random rooted network with $\mathbb{E}[c(\rho)] < \infty$ and let $\mathfrak{F}$ be a sample of FUSF_G. We may assume that (G, ρ) is ergodic, otherwise we take an ergodic decomposition. We may also assume that $\mathsf{FUSF}_G \neq \mathsf{WUSF}_G$ a.s., since otherwise the result follows from [44].

The following is an adaptation of [176, Lemma 4.2] from the unimodular transitive graph setting to our setting. We omit the proof, which is very similar to that of [176].

Lemma 6.3.1 (Component Frequencies)**.** *Let (G, ρ) be an ergodic unimodular random rooted network with $\mathbb{E}[c(\rho)] < \infty$ and let $\mathfrak{F}$ be a sample of FUSF_G. Conditional on (G, ρ), let P_v denote the law of a random walk $\langle X_n \rangle_{n\geq 0}$ on G started at v for each vertex v of G, independent of $\mathfrak{F}$. Then* *there exists a measurable function* $\mathrm{Freq} : \mathcal{G}_\bullet^{\{0,1\}} \to [0,1]$ *such that for every vertex* v *of* G *and every component* T *of* $\mathfrak{F}$,

$$\lim_N \frac{1}{N}\sum_{n=1}^N \mathbb{1}(X_n \in T) = \mathrm{Freq}(G,\rho,T) \quad P_v\text{-a.s.}$$

For each subset W of V, we refer to $\mathrm{Freq}(W) = \mathrm{Freq}(G,\rho,W)$ as the **frequency** of W.

Lemma 6.3.2. *Let* (G,ρ) *be an ergodic unimodular random rooted network with* $\mathbb{E}[c(\rho)] < \infty$ *such that* $\mathsf{FUSF}_G \neq \mathsf{WUSF}_G$ *a.s. and let* $\mathfrak{F}$ *be a sample of* FUSF_G. *Conditioned on* $\mathfrak{F}$ *let* P_x *denote the law of a random walk* $\langle X_n \rangle_{n\geq 0}$ *on* G *started at* x*, independent of* $\mathfrak{F}$. *Assume that with positive probability there exist a component of* $\mathfrak{F}$ *with positive frequency. Then on this event, we a.s. have that for every vertex* u *of* G *such that* $\mathrm{Freq}(T_\mathfrak{F}(u)) > 0$ *and every edge* $e \in T_\mathfrak{F}(u)$ *such that* $\mathfrak{F} \setminus e$ *consists of two infinite components* K_1 *and* K_2*,*

$$P_x\left(\limsup_{N\to\infty} \frac{1}{N}\sum_{n=1}^N \mathbb{1}(X_n \in K_i) > 0 \,\middle|\, (G,\rho,\mathfrak{F}) \right) > 0.$$

for both $i = 1,2$ *and every vertex* $x \in G$.

Proof. We argue similarly to Lemma 6.2.5. For each vertex u of G and each edge e in $T_\mathfrak{F}(u)$, let $K_u(e)$ denote the component of $T_\mathfrak{F}(u)\setminus e$ containing u. For every vertex u such that $\mathrm{Freq}(T_\mathfrak{F}(u)) > 0$, let $E(u)$ be the set of edges e in $T_\mathfrak{F}(u)$ such that

$$P_x\left(\limsup_{N\to\infty} \frac{1}{N}\sum_{n=1}^N \mathbb{1}(X_n \in K_u(e)) > 0 \,\middle|\, (G,\rho,\mathfrak{F}) \right) = 0.$$

for some vertex x (and hence every vertex).

Similarly to the proof of Lemma 6.2.5, we have that for every vertex u, if $E(u)$ is non-empty then it is either a finite simple path or a ray in $T_\mathfrak{F}(u)$ starting at u. Indeed, if $e \in E(u)$ then every edge on the path between u and e is also in $E(u)$, while if e_1 and e_2 do not lie on a simple path from u in $T_\mathfrak{F}(u)$ then the union of $K_u(e_1)$ and $K_u(e_2)$ is all of $T_\mathfrak{F}(u)$ hence

$$\limsup_{N\to\infty} \frac{1}{N}\sum_{n=1}^N \mathbb{1}(X_n \in K_u(e_1)) + \limsup_{N\to\infty} \frac{1}{N}\sum_{n=1}^N \mathbb{1}(X_n \in K_u(e_2)) \geq \mathrm{Freq}(T_\mathfrak{F}(u)) > 0$$

a.s. for every starting point x of the random walk X_n.

First suppose that with positive probability there exists a vertex $u \in \mathrm{core}(F)$ such that $E(u)$ is a finite path. Define a mass transport by sending mass one from each vertex u such that $E(u)$ is finite, to the endpoint of the last edge in $E(u)$; from all other vertices send no mass. Every vertex sends a mass of at most one while, if $E(u)$ is a finite path for some $u \in \mathrm{core}(F)$ and $e = (e^-, e^+)$ is the last edge of this path, then e^+ receives mass one from every vertex in the infinite set $K_u(e)$, contradicting the Mass-Transport Principle.

Next suppose that with positive probability there exists a vertex u such that $E(u)$ is an infinite ray emanating from u. In this case, for any other vertex $u' \in T_{\mathfrak{F}}(u)$ all but finitely many edges e of $E(u)$ satisfy that $u' \in K_u(e)$ and it follows that $E(u')$ is also an infinite ray and $E(u') \triangle E(u)$ is finite. By Theorem 6.2.8 and ergodicity of $(G, \rho, \mathfrak{F})$, the set $E(u)$ is therefore an infinite ray for every vertex u in G a.s. Transport unit mass from each vertex u to the first vertex following u in the ray $E(u)$. Then every vertex u sends unit mass, and receives $\deg_{\mathfrak{F}}(u) - 1$ mass. By the Mass-Transport Principle $\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] = 2$, contradicting [7, Proposition 7.1] and the assumption that $\mathsf{FUSF}_G \neq \mathsf{WUSF}_G$ a.s.

By Theorem 6.1.12 each component of $\mathfrak{F}$ has a non-empty core a.s. and by the above argument $E(u) = \emptyset$ for every core vertex u for which $\mathrm{Freq}(T_{\mathfrak{F}}(u)) > 0$, concluding our proof. $\square$

Proof of Theorem 6.1.10. Suppose that $\mathfrak{F}$ has some finite number $k \geq 2$ of components a.s., which we denote $T_1, \dots T_k$. Then for every N and v

$$\sum_{i=1}^{k} \frac{1}{N} \sum_{n=1}^{N} \mathbb{1}(X_n \in T_i) = 1$$

and so $\sum_{i=1}^{k} \mathrm{Freq}(T_i) = 1$. The frequency of a component is a component property and so, by Theorem 6.2.8 we must have that $\mathrm{Freq}(T_i) = 1/k$ for all $i = 1, \dots, k$.

As in Lemma 6.2.5, conditional on $(G, \rho, \mathfrak{F})$ there exists a simple path $\langle \gamma_i \rangle_{j \geq 0}^{m}$ in G that does not depend on $\mathfrak{F}$ such that, with positive probability

- $T_{\mathfrak{F}}(\gamma_j) \neq T_{\mathfrak{F}}(\gamma_0)$ for $j = 1, \dots, m$ and
- All the edges (γ_j, γ_{j+1}) for $j = 1, \dots, m-1$ belong to the same component, and
- $\gamma_m \in \mathrm{core}(T_{\mathfrak{F}}(\gamma_m))$ and it is the only such vertex in γ.

Denote this event by $\mathscr{B}(\gamma)$. For each $1 \leq j \leq m$, let e_j be an oriented edge of G with $e_j^- = \gamma_j$ and $e_j^+ = \gamma_{j-1}$. Define the forests $\langle \mathfrak{F}_j \rangle_{j=0}^{m}$ recursively by setting $\mathfrak{F}_0 = \mathfrak{F}$ and

$$\mathfrak{F}_j = U(\mathfrak{F}_{j-1}, e_j).$$

Condition on this event $\mathscr{B}(\gamma)$. The choice of the edges $\langle e_j \rangle_{j \geq 0}$ is such that

$$\mathfrak{F}_m = \mathfrak{F} \cup \{e_1\} \setminus \{d\}$$

for some edge d, which we orient so that $d^- = \gamma_m$, that disconnects $T_{\mathfrak{F}}(\gamma_m)$ and whose removal disconnects $T_{\mathfrak{F}}(\gamma_m)$ into two infinite connected components. We denote the component in $\mathfrak{F}$ containing γ_m by K_1 and the component containing d^+ by K_2. As sets of vertices, we have that

$$T_{\mathfrak{F}_m}(\gamma_0) = T_{\mathfrak{F}}(\gamma_0) \cup K_1 \quad \text{and} \quad T_{\mathfrak{F}_m}(d^+) = K_2.$$

Thus, by update-tolerance (Corollary 6.2.4), we have a.s. that

$$\mathrm{Freq}(T_{\mathfrak{F}_m}(\gamma_0)) = \mathrm{Freq}(T_{\mathfrak{F}}(\gamma_0)) + \mathrm{Freq}(K_1) = \frac{1}{k}\,, \tag{6.3.1}$$

and so $\mathrm{Freq}(K_1) = 0$ a.s., contradicting Lemma 6.3.2. □

Proof of Theorem 6.1.11. By Theorem 6.1.10 and the assumption that $\mathfrak{F}$ is disconnected a.s., $\mathfrak{F}$ has infinitely many components. For every N and v, letting $\langle X_n \rangle_{n\geq 0}$ be a simple random walk independent of $\mathfrak{F}$ started at v,

$$\sum_{T_i \text{ is a component of } \mathfrak{F}} \frac{1}{N}\sum_{n=1}^{N} \mathbb{1}(X_n \in T_i) = 1$$

and so, taking the limit

$$\sum_{T_i \text{ a component of } \mathfrak{F}} \mathrm{Freq}(T_i) \leq 1.$$

By Theorem 6.2.8 all the component frequencies are equal and so must all equal zero. Thus, for every component T of $\mathfrak{F}$,

$$\lim \frac{1}{N}\sum_{n=1}^{N} \mathbb{1}(X_n \in T) \to 0 \quad \text{ a.s.}\,,$$

which in particular implies that for any $\varepsilon > 0$ there exists some n for which $\mathbb{P}(X_n \in T_{\mathfrak{F}}(\rho)) \leq \varepsilon$, hence there exists a vertex v with $\mathbb{P}(v \in T_{\mathfrak{F}}(\rho)) \leq \varepsilon$, as required. □

Remark 6.3.3. It is possible to remove the application of indistinguishability in the above proofs. If there is a unique component with non-zero frequency a.s., Lemma 6.3.2 allows us to perform updates to create two components of non-zero frequency, contradicting update-tolerance and ergodicity. Otherwise, consider the component of maximal frequency. The maximal component frequency is non-random in an ergodic unimodular random rooted network. However, updating allows us to attach an infinite part of another positive-frequency component to the maximal frequency component, increasing its frequency by Lemma 6.3.2, contradicting update-tolerance.

6.4 Indistinguishability of WUSF components

6.4.1 Indistinguishability of WUSF components by tail properties.

Let G be a transient network. Recall that **Wilson's algorithm rooted at infinity** [44, 197, 228] allows us to generate a sample of WUSF_G by joining together loop-erased random walks on G. Let γ be a path in G that is either finite or transient, i.e. visits each vertex of G at most finitely many times. The **loop-erasure** $\mathsf{LE}(\gamma)$ of the path γ (introduced by Lawler [161]) is formed by erasing cycles from γ chronologically as they are created. More formally, we define $\mathsf{LE}(\gamma)_i = \gamma_{t_i}$ where the times t_i are defined recursively by $t_0 = 0$ and $t_i = 1 + \max\{t \geq t_{i-1} : \gamma_t = \gamma_{t_{i-1}}\}$.

Wilson's algorithm rooted at infinity generates a sample $\mathfrak{F}$ of WUSF_G as follows. Let $\{v_j : j \in \mathbb{N}\}$ be an enumeration of the vertices of G and define a sequence of forests $\langle \mathfrak{F}_i \rangle_{i \geq 0}$ in G as follows:

1. Let $\mathfrak{F}_0 = \emptyset$.
2. Given $\mathfrak{F}_j$, start an independent random walk from v_{j+1} stopped if and when it hits the set of vertices already included in $\mathfrak{F}_j$.
3. Form the loop-erasure of this random walk path and let $\mathfrak{F}_{j+1}$ be the union of $\mathfrak{F}_j$ with this loop-erased path.
4. Let $\mathfrak{F} = \bigcup_{j \geq 0} \mathfrak{F}_j$.

It is a fact that the choice of enumeration does not affect the law of $\mathfrak{F}$. We will require the following slight variation on Wilson's algorithm rooted at infinity.

Lemma 6.4.1. *Let W be a finite set of vertices in G, and let $\{v_j : j \in \mathbb{N}\}$ and $\{w_j : 1 \leq j \leq |W|\}$ be enumerations of $V(G) \setminus W$ and W respectively. Let $\mathfrak{F}$ be the spanning forest of G generated as follows.*

1. *Let $\mathfrak{F}'_0 = \emptyset$.*
2. *Given $\mathfrak{F}'_j$, start an independent random walk from v_{j+1} stopped if and when it hits the set of vertices already included in $\mathfrak{F}'_j$.*
3. *Form the loop-erasure of this random walk path and let $\mathfrak{F}'_{j+1}$ be the union of $\mathfrak{F}'_j$ with this loop-erased path.*
4. *Let $\mathfrak{F}_0 = \bigcup_{j \geq 0} \mathfrak{F}'_j$.*
5. *Given $\mathfrak{F}_j$, start an independent random walk from w_{j+1} stopped if and when it hits the set of vertices already included in $\mathfrak{F}_j$.*
6. *Let $\mathfrak{F} = \bigcup_{j=0}^{|W|} \mathfrak{F}_j$.*

Then $\mathfrak{F}$ is distributed according to WUSF_G.

Proof. Let $j_0 = \max\{j : v_j \text{ is adjacent to } W\}$ and consider the enumeration

$$v_1, \ldots, v_{j_0}, w_1, \ldots, w_{|W|}, v_{j_0+1}, \ldots$$

of $V(G)$. Let $\{X^v : v \in V(G)\}$ be a collection of independent random walks, one from each vertex v of G. Let $\mathfrak{F}$ be generated using the random walks $\{X^v : v \in V(G)\}$ as above and let $\mathfrak{F}'$ be a sample of WUSF_G generated using Wilson's algorithm rooted at infinity, using the enumeration $v_1, \ldots, v_{j_0}, w_1, \ldots, w_{|W|}, v_{j_0+1}, \ldots$ and the same collection of random walks $\{X^v\}$. Then $\mathfrak{F}' = \mathfrak{F}$, and so $\mathfrak{F}'$ has distribution WUSF_G as claimed. $\square$

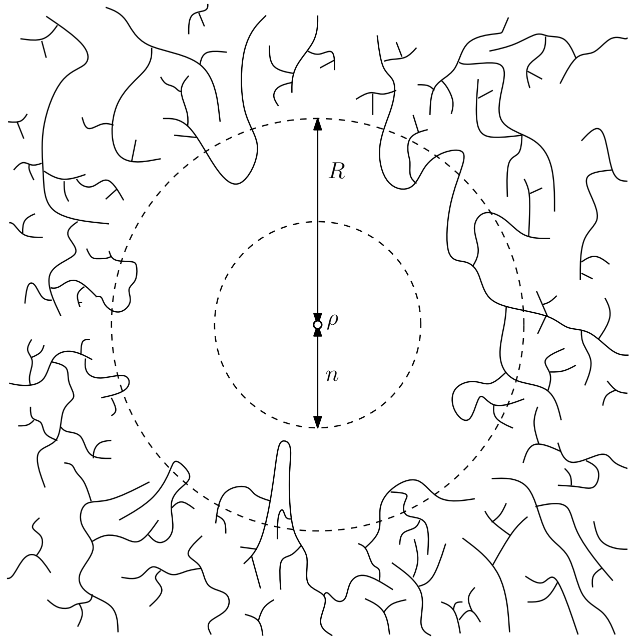

Figure 6.2: Illustration of the forest $\mathfrak{F}_R$ and the event $\mathscr{C}_{R,n}$. The forest $\mathfrak{F}_R$ is defined to be the union of the futures in $\mathfrak{F}$ of each of the vertices of G outside the ball of radius R about ρ. $\mathscr{C}_{R,n}$ is the event that none of these futures intersect the ball of radius n about ρ.

Proof of Theorem 6.1.20. If G is recurrent, then its WUSF is a.s. connected and the statement holds trivially, so let us assume that G is a.s. transient.

For each $r > 0$ let $\mathscr{G}_r$ be the σ-algebra generated by the random variables (G, ρ) and $\mathfrak{F} \cap B_G(\rho, r)$. Let $\widehat{G}$ be the network obtained from G by contracting every edge of $\mathfrak{F} \cap B_G(\rho, r)$ and deleting every edge of $B_G(\rho, r) \setminus \mathfrak{F}$. Conditional on $\mathscr{G}_r$, the forest $\mathfrak{F}$ is distributed as the union of $\mathfrak{F} \cap B_G(\rho, r)$ and a sample of $\mathsf{WUSF}_{\widehat{G}}$ by the WUSF's spatial Markov property (see Section 6.1.2).

For each integer $R \geq 0$, let $\mathfrak{F}_R$ be the subgraph of $\mathfrak{F}$ defined to be the union of the futures in $\mathfrak{F}$ of every vertex in $G \setminus B_G(\rho, R)$. For each R, we can sample $\mathfrak{F}_R$ conditioned on $\mathscr{G}_r$ by running Wilson's algorithm on $\widehat{G}$ starting from every vertex $\widehat{G} \setminus B_G(\rho, R)$ (in arbitrary order). A vertex v is contained in $\mathfrak{F}_R$ if and only if its past $\mathrm{past}_{\mathfrak{F}}(v)$ intersects $G \setminus B_G(\rho, R)$, and so $\bigcap_{R \geq 0} \mathfrak{F}_R = \emptyset$ a.s. by the assumption that $\mathfrak{F}$ has one-ended components a.s. Conditional on $\mathscr{G}_r$ and $\mathfrak{F}_R$, the rest of $\mathfrak{F}$ may be sampled by finishing the run of Wilson's algorithm on $\widehat{G}$ as described in Lemma 6.4.1.

For each R and n such that $R \geq n$, let $\mathscr{C}_{R,n}$ be the event that $\mathfrak{F}_R \cap B_G(\rho, n) = \emptyset$, so that $\lim_{R \to \infty} \mathbb{P}(\mathscr{C}_{R,n}) = 1$ for each fixed n (see Figure 6.2). Let $\langle X_i \rangle_{i \geq 0}$ be a random walk on G started at ρ and let $\langle \widehat{X}_i \rangle_{i \geq 0}$ be a random walk on $\widehat{G}$ started at (the equivalence class in $\widehat{G}$ of) ρ. Fix $r \leq n \leq R$ and condition on $\mathscr{G}_r, \mathfrak{F}_R$ and on the event $\mathscr{C}_{R,n}$. The definition of $\mathfrak{F}_R$ ensures that $\mathfrak{F} \setminus \mathfrak{F}_R$ is finite and that $T_{\mathfrak{F}}(v) \setminus T_{\mathfrak{F}_R}(v)$ is finite for every $v \in \mathfrak{F}_R$. Thus, since $\mathscr{A}$ is a tail property, the event that $T_{\mathfrak{F}}(v)$ has property $\mathscr{A}$ is already determined by (G, ρ) and $\mathfrak{F}_R$ for every for vertex $v \in \mathfrak{F}_R$. Thus, by Lemma 6.4.1, the conditional probability that ρ is in an $\mathscr{A}$ cluster equals the conditional probability that the random walk $\langle \widehat{X}_i \rangle_{i \geq 0}$ first hits $\mathfrak{F}_R$ at a vertex that belongs to an

$\mathscr{A}$ cluster, so that

$$\mathbb{P}(T_{\mathfrak{F}}(\rho) \in \mathscr{A} \mid \mathscr{G}_r, \mathfrak{F}_R, \mathscr{C}_{R,n}) = \mathbb{P}(\langle \widehat{X}_i \rangle_{i \geq 0} \text{ first hits } \mathfrak{F}_R \text{ at an } \mathscr{A} \text{ cluster} \mid \mathscr{G}_r, \mathfrak{F}_R, \mathscr{C}_{R,n}).$$

On the event $\mathscr{C}_{R,n}$ the walk $\langle \widehat{X}_i \rangle_{i \geq 0}$ must hit $\mathfrak{F}_R$ at time $n - r$ or greater, and so for all $j < n - r$ we have

$$\begin{aligned}\mathbb{P}(T_{\mathfrak{F}}(\rho) \in \mathscr{A} \mid \mathscr{G}_r, \mathfrak{F}_R, \mathscr{C}_{R,n}) &= \mathbb{P}(\langle \widehat{X}_i \rangle_{i \geq j} \text{ first hits } \mathfrak{F}_R \text{ at an } \mathscr{A} \text{ cluster} \mid \mathscr{G}_r, \mathfrak{F}_R, \mathscr{C}_{R,n}) \\ &= \mathbb{P}(T_{\mathfrak{F}}(\widehat{X}_j) \in \mathscr{A} \mid \mathscr{G}_r, \mathfrak{F}_R, \mathscr{C}_{R,n})\end{aligned}$$

where the last equality is due to Lemma 6.4.1. That is,

$$\mathbb{P}(\{T_{\mathfrak{F}}(\rho) \in \mathscr{A}\} \cap \mathscr{C}_{R,n} \mid \mathscr{G}_r, \mathfrak{F}_R) = \mathbb{P}(\{T_{\mathfrak{F}}(\widehat{X}_j) \in \mathscr{A}\} \cap \mathscr{C}_{R,n} \mid \mathscr{G}_r, \mathfrak{F}_R)$$

a.s. for all $j < n - r$. Taking conditional expectations with respect to $\mathscr{G}_r$ gives

$$\mathbb{P}(\{T_{\mathfrak{F}}(\rho) \in \mathscr{A}\} \cap \mathscr{C}_{R,n} \mid \mathscr{G}_r) = \mathbb{P}(\{T_{\mathfrak{F}}(\widehat{X}_j) \in \mathscr{A}\} \cap \mathscr{C}_{R,n} \mid \mathscr{G}_r)$$

a.s. for all $j < n - r$. Since the events $\mathscr{C}_{R,n}$ are increasing in R and $\mathbb{P}(\mathscr{C}_{R,n}) \to 1$ as $R \to \infty$, taking the limit $R \to \infty$ in the above equality gives that

$$\mathbb{P}(T_{\mathfrak{F}}(\rho) \in \mathscr{A} \mid \mathscr{G}_r) = \mathbb{P}(T_{\mathfrak{F}}(\widehat{X}_j) \in \mathscr{A} \mid \mathscr{G}_r) \tag{6.4.1}$$

for all $j < n - r$. This equality holds for all j by taking n to infinity.

Let τ and $\widehat{\tau}$ be the last times that $\langle X_n \rangle_{n \geq 0}$ and $\langle \widehat{X}_n \rangle_{n \geq 0}$ visit $B_G(\rho, r+1)$ respectively. Then for each vertex $v \in B_G(\rho, r+1)$, the conditional distributions

$$\begin{aligned}&\langle X_{\tau+n} \rangle_{n \geq 0} \text{ conditioned on } X_\tau = v \text{ and } (G, \rho) \text{ and} \\ &\langle \widehat{X}_{\widehat{\tau}+n} \rangle_{n \geq 0} \text{ conditioned on } \widehat{X}_{\widehat{\tau}} = v \text{ and } \mathscr{G}_r\end{aligned} \tag{6.4.2}$$

are equal.

Let $\mathcal{I}$ denote the invariant σ-algebra of the stationary sequence $\langle (G, X_n, \mathfrak{F}) \rangle_{n \geq 0}$. The Ergodic Theorem implies that

$$\frac{1}{N} \sum_{i=1}^{N} \mathbb{1}(T_{\mathfrak{F}}(X_i) \in \mathscr{A}) \xrightarrow[N \to \infty]{\text{a.s.}} Y := \mathbb{P}(T_{\mathfrak{F}}(\rho) \in \mathscr{A} \mid \mathcal{I}).$$

Moreover, the random variable Y is measurable with respect to the completion of the σ-algebra generated by $(G,\rho,\mathfrak{F})$: to see this, note that for every $a<b\in[0,1]$, Levy's 0-1 law implies that

$$\mathbb{P}\left(\lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\mathbb{1}(T_{\mathfrak{F}}(X_i)\in\mathscr{A})\in[a,b]\,\middle|\,(G,\rho),\mathfrak{F},\langle X_n\rangle_{n=0}^{k}\right) \xrightarrow[k\to\infty]{a.s.} \mathbb{1}\left(\lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\mathbb{1}(T_{\mathfrak{F}}(X_i)\in\mathscr{A})\in[a,b]\right).$$

But

$$\mathbb{P}\left(\lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\mathbb{1}(T_{\mathfrak{F}}(X_i)\in\mathscr{A})\in[a,b]\,\middle|\,(G,\rho),\mathfrak{F},\langle X_n\rangle_{n=0}^{k}\right) = \mathbb{P}\left(\lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\mathbb{1}(T_{\mathfrak{F}}(X_{k+i})\in\mathscr{A})\in[a,b]\,\middle|\,(G,X_k),\mathfrak{F}\right)$$

and so, by stationarity, $\mathbb{P}\big(Y\in[a,b]\,\big|\,(G,\rho,\mathfrak{F})\big)\in\{0,1\}$ a.s. In particular, Y is independent of X_τ given (G,ρ).

Since G is transient, τ is finite a.s. and so

$$\frac{1}{N}\sum_{i=1}^{N}\mathbb{1}(T_{\mathfrak{F}}(X_{\tau+i})\in\mathscr{A})\xrightarrow[N\to\infty]{\text{a.s.}} Y.$$

In particular, this a.s. convergence holds conditioned on $X_\tau=v$ for each v such that $\mathbb{P}(X_\tau=v)>0$. Since the support of $\widehat{X}_{\widehat{\tau}}$ is contained in the support of X_τ and the conditioned measures in (6.4.2) are equal, we have by the above that there exists a random variable $\widehat{Y}$ such that

$$\frac{1}{N}\sum_{i=1}^{N}\mathbb{1}(T_{\mathfrak{F}}(\widehat{X}_{\widehat{\tau}+i})\in\mathscr{A})\xrightarrow[N\to\infty]{\text{a.s.}} \widehat{Y}\,,$$

and the distribution of $\widehat{Y}$ given $\mathscr{G}_r$ and $\widehat{X}_{\widehat{\tau}}=v$ is equal to the distribution of Y given (G,ρ) and $X_\tau=v$, so that $\widehat{Y}$ is in fact independent of $\mathscr{G}_r$ and $\widehat{X}_{\widehat{\tau}}$ given (G,ρ). That is, for every $a<b\in[0,1]$,

$$\mathbb{P}(\widehat{Y}\in[a,b]\mid\mathscr{G}_r,\widehat{X}_{\widehat{\tau}}=v)=\mathbb{P}(Y\in[a,b]\mid(G,\rho),X_\tau=v)=\mathbb{P}(Y\in[a,b]\mid(G,\rho))$$

so that, taking conditional expectations with respect to (G,ρ),

$$\mathbb{P}(\widehat{Y}\in[a,b]\mid(G,\rho))=\mathbb{P}(Y\in[a,b]\mid(G,\rho))=\mathbb{P}(\widehat{Y}\in[a,b]\mid\mathscr{G}_r,\widehat{X}_{\widehat{\tau}}=v) \tag{6.4.3}$$

establishing the independence of $\widehat{Y}$ from $\mathscr{G}_r$ and $\widehat{X}_{\widehat{\tau}}$ conditional on (G,ρ).

Since $\widehat{\tau}$ is finite a.s. we also have that

$$\frac{1}{N}\sum_{i=1}^{N}\mathbb{1}(T_{\mathfrak{F}}(\widehat{X}_i)\in\mathscr{A})\xrightarrow[N\to\infty]{\text{a.s.}}\widehat{Y}.$$

Hence, by (6.4.1) and the conditional Dominated Convergence Theorem,

$$\begin{aligned}\mathbb{P}(T_{\mathfrak{F}}(\rho)\in\mathscr{A}\mid\mathscr{G}_r)&=\mathbb{E}\big[\mathbb{1}(T_{\mathfrak{F}}(\widehat{X}_j)\in\mathscr{A})\mid\mathscr{G}_r\big]\\&=\mathbb{E}\left[\frac{1}{N}\sum_{j=1}^{N}\mathbb{1}(T_{\mathfrak{F}}(\widehat{X}_j)\in\mathscr{A})\,\middle|\,\mathscr{G}_r\right]\xrightarrow[N\to\infty]{a.s.}\mathbb{E}\left[\widehat{Y}\,\middle|\,\mathscr{G}_r\right]=\mathbb{E}\left[\widehat{Y}\,\middle|\,(G,\rho)\right].\end{aligned}$$

It follows by similar reasoning to (6.4.3) that the event $\{T_{\mathfrak{F}}(\rho)\in\mathscr{A}\}$ is independent of $\mathfrak{F}\cap B_G(\rho,r)$ conditional on (G,ρ) for every r. It follows that $\mathsf{WUSF}_G(T_{\mathfrak{F}}(\rho)\in\mathscr{A})\in\{0,1\}$ a.s., and hence $\mathsf{WUSF}_G(T_{\mathfrak{F}}(v)\in\mathscr{A})\in\{0,1\}$ for every vertex v of G a.s. by stationarity. But, given (G,ρ), every vertex v of G has positive probability of being in the same component of $\mathfrak{F}$ as ρ, and so we must have that the probabilities

$$\mathsf{WUSF}_G(T_{\mathfrak{F}}(v)\in\mathscr{A})=\mathsf{WUSF}_G(T_{\mathfrak{F}}(\rho)\in\mathscr{A})\in\{0,1\}$$

agree for every vertex v of G a.s., so that either every component of $\mathfrak{F}$ has type $\mathscr{A}$ a.s. conditional on (G,ρ) or every component of $\mathfrak{F}$ does not have type $\mathscr{A}$ a.s. conditional on (G,ρ), completing the proof. □

6.4.2 Indistinguishability of WUSF components by non-tail properties.

Wired cycle-breaking for the WUSF

Let G be an infinite network and let f be a spanning forest of G such that every component of f is infinite and one-ended. For every oriented edge e of G, we define the **update** $U(f,e)$ of f at e as follows:

If e is a self-loop, or is already contained in f, let $U(f,e)=f$. Otherwise:

1. If e^- and e^+ are in the same component of f, so that $f\cup\{e\}$ contains a cycle, let d be the unique edge of f that is both contained in this cycle and adjacent to e^- and let $U(f,e)=f\cup\{e\}\setminus\{d\}$.
2. Otherwise, let d be the unique edge of f such that d is adjacent to e^- and the component containing e^- in $f\setminus\{d\}$ is finite, and let $U(f,e)=f\cup\{e\}\setminus\{d\}$.

The following are proved in [127] (in which an appropriate update rule is also developed for the case that the WUSF has multiply-ended components) and can also be proved similarly to the proof in Section 6.2.1.

Proposition 6.4.2. *Let G be a network, let $\mathfrak{F}$ be a sample of WUSF_G and suppose that every component of $\mathfrak{F}$ is one-ended almost surely. Let v be a fixed vertex of G and let E be an edge chosen from the set $\{e : e^- = v\}$ independently of $\mathfrak{F}$ and with probability proportional to its conductance. Then $U(\mathfrak{F}, E)$ and $\mathfrak{F}$ have the same distribution.*

Corollary 6.4.3 (Update-tolerance for the WUSF)**.** *Let G be a network and let $\mathfrak{F}$ be a sample of WUSF_G. If every component of $\mathfrak{F}$ is one-ended almost surely, then for every event $\mathscr{A} \subset \{0,1\}^{E(G)}$ and every oriented edge e in G,*

$$\mathsf{WUSF}_G(\mathfrak{F} \in \mathscr{A}) \geq \frac{c(e)}{c(e^-)}\mathsf{WUSF}_G(U(\mathfrak{F}, e) \in \mathscr{A}).$$

Proof. By Proposition 6.4.2,

$$\begin{aligned}\mathsf{WUSF}_G(\mathfrak{F} \in \mathscr{A}) &= \frac{1}{c(e^-)} \sum_{\hat{e}^- = e^-} c(\hat{e})\mathsf{WUSF}_G(U(\mathfrak{F}, \hat{e}) \in \mathscr{A}) \\ &\geq \frac{c(e)}{c(e^-)}\mathsf{WUSF}_G(U(\mathfrak{F}, e) \in \mathscr{A}).\end{aligned}$$ □

Pivotal edges for the WUSF

Let G be a network, and let f be a spanning forest of G such that every component of f is infinite and one-ended and let $\mathscr{A}$ be a component property. We call an oriented edge e of G a **good pivotal edge** for a vertex v of G if either

1. $e^+ \in T_f(v)$, $e^- \in T_f(v)$, and the type of $T_{U(f,e)}(v)$ is different from the type of $T_f(v)$ (in which case we say e is a **good internal pivotal** edge for v),

2. $e^+ \notin T_f(v)$, $e^- \notin T_f(v)$, and the type of $T_{U(f,e)}(v)$ is different from the type of $T_f(v)$ (in which case we say e is a **good external pivotal** edge for v),

3. $e^+ \in T_f(v)$, $e^- \notin T_f(v)$, and the type of $T_{U(f,e)}(v)$ is different from the type of $T_f(v)$ (in which case we say e is a **good additive pivotal** edge for v), or

4. $e^- \in T_f(v)$, $e^+ \notin T_f(v)$, the component of v in $T_f(v) \setminus \{e^-\}$ is infinite and the type of $T_{U(f,e)}(v)$ is different from the type of $T_f(v)$ (in which case we say e is a **good subtractive pivotal** edge for v).

In particular, e is a good pivotal for some vertex v only if infinitely many vertices change the type of their component when we update from f to $U(f, e)$.

Lemma 6.4.4. *Let G be a network, let $\mathfrak{F}$ be a sample of WUSF_G and let $\mathscr{A}$ be a component property. If every component of $\mathfrak{F}$ is one-ended a.s., then either*

> *there exists a good pivotal edge for some vertex v with positive probability*

or

$\mathscr{A}$ is WUSF_G-equivalent to a tail component property. That is, there exists a tail component property $\mathscr{A}'$ such that

$$\mathsf{WUSF}_G((G,v,\mathfrak{F}) \in \mathscr{A} \triangle \mathscr{A}') = 0.$$

for every vertex v of G.

Proof. Suppose that no good pivotal edges exist a.s. and let $\mathscr{A}' \subseteq \mathcal{G}_\bullet^{(0,\infty)\times\{0,1\}}$ be the component property

$$\mathscr{A}' = \left\{ (G,v,\omega) : \begin{array}{l} \text{There exists a vertex } u \in K_\omega(v) \text{ and a one-ended essential} \\ \text{spanning forest } f \text{ of } G \text{ such that } (G,u,f) \in \mathscr{A} \text{ and the} \\ \text{symmetric differences } \omega \triangle f \text{ and } K_\omega(u) \triangle K_f(u) \text{ are finite.} \end{array} \right\}.$$

Note that by definition $\mathscr{A}'$ is a tail component property. Our goal is to show that $\mathscr{A}$ and $\mathscr{A}'$ have $\mathsf{WUSF}_G((G,v,\mathfrak{F}) \in \mathscr{A} \triangle \mathscr{A}') = 0$ for every vertex v of G. One part of this assertion is easy, indeed, let $\Omega_0 \subset \mathcal{G}_\bullet^{(0,\infty)\times\{0,1\}}$ be the event $\Omega_0 = \{(G,v,\omega) : \omega \text{ is a one-ended essential spanning forest}\}$ so that $\mathsf{WUSF}_G(\Omega_0) = 1$ by assumption. Then $\mathscr{A} \cap \Omega_0 \subset \mathscr{A}'$ since one can take $f = \omega$ and $u = v$ in the definition of $\mathscr{A}'$.

The second part of this assertion is slightly more difficult and requires the use of update-tolerance. Given a one-ended essential spanning forest $\mathfrak{F}$ and a finite sequence of oriented edges $\langle e_i \rangle_{i=1}^n$ of G we define $U(\mathfrak{F}; e_1, \ldots, e_n)$ recursively by $U(\mathfrak{F}; e_1) = U(\mathfrak{F}, e_1)$ and

$$U(\mathfrak{F}; e_1, \ldots, e_n) = U(U(\mathfrak{F}; e_1, \ldots, e_{n-1}), e_n).$$

Let $\mathfrak{F}$ be a sample of WUSF_G and let Ω_1 be the event that for any finite sequence of edges $\langle e_i \rangle_{i=1}^n$ the forest $U(\mathfrak{F}; e_1, \ldots, e_n)$ has no good pivotal edges. By update-tolerance and the assumption that $\mathfrak{F}$ has no good pivotal edges a.s., $\mathsf{WUSF}_G(\Omega_1) = 1$. Thus, it suffices to show that $\mathscr{A}' \cap \Omega_0 \cap \Omega_1 \subset \mathscr{A}$.

Let $(G,v,\mathfrak{F}) \in \mathscr{A}' \cap \Omega_0 \cap \Omega_1$ and let f be a one-ended essential spanning forest such that $\mathfrak{F} \triangle f$ and $T_\mathfrak{F}(u) \triangle T_f(u)$ are finite and $(G,u,f) \in \mathscr{A}$ for some vertex $u \in T_\mathfrak{F}(v)$. We will prove by induction on $|f \setminus \mathfrak{F}|$ that there exists a vertex $u' \in T_\mathfrak{F}(v)$ with $(G,u',f) \in \mathscr{A}$ and a finite sequence of oriented edges $\langle e_i \rangle_{i=1}^n$ of G such that $U(\mathfrak{F}; e_1, \ldots, e_n) = f$ and $T_\mathfrak{F}(u')$ and $T_{U(\mathfrak{F}; e_1, \ldots, e_i)}(u')$ have the same type for every $1 \le i \le n$. Since $(G,u',f) \in \mathscr{A}$ by assumption, this will imply that $(G,v,\mathfrak{F}) \in \mathscr{A}$ as desired.

To initialize the induction assume that $|f \setminus \mathfrak{F}| = 0$. Then $f \subset \mathfrak{F}$ and, since both $\mathfrak{F}$ and f are one-ended essential spanning forests, we must have that $\mathfrak{F} \subset f$ since any addition of an edge to f creates either a cycle or a two-ended component, so that $\mathfrak{F} = f$ and the claim is trivial.

Next, assume that $|f \setminus \mathfrak{F}| > 0$ and let $h \in f \setminus \mathfrak{F}$. Since $\mathfrak{F}$ is a one-ended essential spanning forest, $\mathfrak{F} \cup \{h\}$ contains either a cycle or a two-ended component and we can therefore find an edge $g \in \mathfrak{F} \setminus f$ such that $\mathfrak{F}' = \mathfrak{F} \cup \{h\} \setminus \{g\}$ is a one-ended essential spanning forest. The choice of g is not unique, and will be important in the final case below.

First suppose $\mathfrak{F} \cup \{h\}$ contains a cycle. In this case the choice of an edge g as above is not important. The edge g must be contained in this cycle since otherwise the cycle would be contained in $\mathfrak{F}'$. Let $e_1, \ldots, e_k$ be an oriented simple path on this cycle so that $e_1 = g$ and $e_k = h$. We have that $\mathfrak{F}' = U(\mathfrak{F}; e_k, e_{k-1}, \ldots, e_2)$ by definition of the update operation. Since none of the forests $U(\mathfrak{F}; e_k, \ldots, e_i)$ have any good internal or external pivotal edges, we have that $T_{\mathfrak{F}'}(u)$ has the same type as $T_{\mathfrak{F}}(u)$. Lastly, $(G, u, \mathfrak{F}) \in \Omega_0 \cap \Omega_1$ and $|\mathfrak{F}' \cap f| < |\mathfrak{F} \cap f|$, so that our induction hypothesis provides us with a vertex $u' \in T_{\mathfrak{F}'}(u)$ and a sequence of edges $e'_1, \ldots, e'_m$ such that $U(\mathfrak{F}; e'_1, \ldots, e'_m) = f$ and $T_{\mathfrak{F}}(u')$ and $T_{U(\mathfrak{F}; e_1, \ldots, e_i)}(u')$ have the same type for every $1 \leq i \leq n$. Since this also holds when u' is replaced by any vertex u'' in the future of u' in $\mathfrak{F}'$, we may take u'' such that the above hold and $u'' \in T_{\mathfrak{F}}(v)$. We conclude the induction step by concatenating the two sequences $e'_1, \ldots, e'_m, e_k, \ldots, e_2$.

Now suppose that $\mathfrak{F} \cup \{h\}$ contains a two-ended component. Let us first consider the easier case in which u is not contained in this two-ended component, which is the case if and only if neither of the endpoints of h are in $T_{\mathfrak{F}}(u)$. In this case the choice of an edge g as above is not important. The edge g must be such that the removal of g disconnects the component of $\mathfrak{F} \cup \{h\}$ containing g into two infinite connected components. We orient h so that its tail is in the component of g in $\mathfrak{F}$ and orient g so that its head is in the component of $\mathfrak{F} \setminus \{g\}$ containing h^-. We then take an oriented simple path in $\mathfrak{F}$ from g^+ to h^- and append to it the edge h. As above, performing the updates from the last edge of the path (that is, h) to the first (the edge in the path touching g^+) yields $\mathfrak{F}'$. Since none of the forests $U(\mathfrak{F}; e_k, \ldots, e_i)$ have any good external pivotal edges, we have that $T_{\mathfrak{F}'}(u)$ has the same type as $T_{\mathfrak{F}}(u)$. We may now apply our induction hypothesis to $(G, u, \mathfrak{F}')$ as before to complete the induction step in this case.

Finally, if $\mathfrak{F} \cup \{h\}$ contains a two-ended component and one of the endpoints of h is in $T_{\mathfrak{F}}(u)$. The choice of g is important in this case. Orient h so that $h^+ \in T_{\mathfrak{F}}(u)$ and consider the unique infinite rays from h^+ and h^- in $\mathfrak{F}$, denoted $e_1, e_2, \ldots$ and $e_{-1}, e_{-2}, \ldots$ respectively. Orient the ray $\langle e_i \rangle_{i \geq 1}$ towards infinity and the ray $\langle e_{-i} \rangle_{i \geq 0}$ towards h^- so that, writing $e_0 = h$, $\langle e_i \rangle_{i \in \mathbb{Z}}$ is an oriented bi-infinite path in $\mathfrak{F} \cup \{h\}$.

Next consider the unique infinite ray from u in $\mathfrak{F}$. Since the symmetric difference $T_{\mathfrak{F}}(u) \triangle T_f(u)$ is finite, all but finitely many of the edges in the infinite ray from u in $\mathfrak{F}$ must also be contained in the component of u in f. Let u' be the first vertex in the infinite ray from u in $\mathfrak{F}$ such that u' is contained in the ray from h^+ in $\mathfrak{F}$ and all of the ray from u' in $\mathfrak{F}$ is contained in f, so that $u' = e_k^+ = e_{k+1}^-$ for some $k \geq 0$.

Since f is a one-ended essential spanning forest and contains the ray $\langle e_i \rangle_{i \geq k+1}$, there exists an edge e_l with $l < k$ such that $e_l \notin f$. By the definition of the update operation, we have that $\mathfrak{F}' = U(\mathfrak{F}; -e_0, -e_1, \ldots, -e_{l-1})$ if $l > 0$ and $\mathfrak{F}' = U(\mathfrak{F}, e_0, e_{-1}, \ldots, e_{l+1})$ if $l < 0$. Let $\mathfrak{F}_j$ denote either $U(\mathfrak{F}; -e_0, -e_1, \ldots, -e_j)$ or $U(\mathfrak{F}, e_0, e_{-1}, \ldots, e_{-j})$ for each $j \leq l - 1$ as appropriate. In either case, u' is in an infinite connected component of $\mathfrak{F}_j \setminus e_{j+1}$ for each j and so, since good pivotal edges do not exist for any of the $\mathfrak{F}_j$, the type of $T_{\mathfrak{F}'}(u')$ is the same as the type of $T_{\mathfrak{F}}(u')$. Lastly, we also have that $(G, u', \mathfrak{F}) \in \Omega_0 \cap \Omega_1$ and $|\mathfrak{F}' \cap f| < |\mathfrak{F} \cap f|$, and so we may use apply our induction hypothesis to $(G, u', \mathfrak{F}')$ as before, completing the proof. $\square$

Indistinguishability of WUSF components by non-tail properties

Our goal in this section is to prove the following, theorem, which in conjunction with Theorem 6.2.8 completes the proof of Theorem 6.1.9.

Theorem 6.4.5. *Let (G,ρ) be a unimodular random network with $\mathbb{E}[c(\rho)] < \infty$, and let $\mathfrak{F}$ be a sample of WUSF_G. Then for every component property $\mathscr{A}$, either every connected component of $\mathfrak{F}$ has property $\mathscr{A}$ or none of the connected components of $\mathfrak{F}$ have property $\mathscr{A}$ almost surely.*

Proof of Theorem 6.4.5. We may assume that G is transient, since otherwise $\mathfrak{F}$ is connected a.s. and the claim is trivial. Since (G,ρ) becomes reversible when biased by $c(\rho)$, Theorem 6.1.20 implies that the components of $\mathfrak{F}$ are indistinguishable by tail properties (and therefore also by properties equivalent to tail properties), so that we may assume from now on that $\mathscr{A}$ is not equivalent to a tail property. In this case, Lemma 6.4.4 implies that good pivotal edges exist for ρ with positive probability. Without loss of generality, we may assume further that, with positive probability, $T_{\mathfrak{F}}(\rho)$ has property $\mathscr{A}$ and there exists a good pivotal edge for ρ: if not, replace $\mathscr{A}$ with $\neg\mathscr{A}$. In this case, there exist a natural numbers r such that, with positive probability $T_{\mathfrak{F}}(\rho)$ has property $\mathscr{A}$ and there exists a good pivotal edge e for ρ at distance at most r from ρ in G.

Let $\{\theta(e) : e \in E\}$ be i.i.d. uniform $[0,1]$ random variables indexed by the edges of G, and let $\langle \omega_n \rangle_{n\geq 1}$ be Bernoulli $(1-1/(n+1))$-bond percolations on G defined by setting $\omega_n(e) = 1$ if and only if $\theta(e) \geq 1 - 1/(n+1)$. By Theorem 6.1.5, every connected component of $\mathfrak{F}$ is one-ended a.s. and so every component of $\mathfrak{F} \cap \omega_n$ is finite for every n a.s. Given $(G,\rho,\mathfrak{F},\theta)$, for each vertex u of G let $v_n(u)$ be a vertex chosen uniformly at random from the cluster of u in $\mathfrak{F}\cap\omega_n$ and let $e_n(u)$ be an oriented edge chosen uniformly from the ball of radius r about $v_n(u)$ in G, where $(v_n(u), e_n(u))$ and $(v_{n'}(u'), e_{m'}(u'))$ are taken to be independent conditional on $(G,\rho,\mathfrak{F},\theta)$ if $n' \neq n$ or $u' \neq u$. We write $v_n = v_n(\rho), e_n = e_n(\rho)$ and let $\widehat{\mathbb{P}}$ denote the joint law of $(G,\rho,\mathfrak{F},\theta,\langle (v_n(u),e_n(u)) : u \in V\rangle_{n\geq 1})$. The following is a special case of a standard fact about unimodular random rooted networks.

Lemma 6.4.6. *$(G,\rho,v_n,\mathfrak{F},\theta)$ and $(G,v_n,\rho,\mathfrak{F},\theta)$ have the same distribution.*

Proof. Let $\mathscr{B} \subseteq \mathcal{G}_{\bullet\bullet}^{(0,\infty)\times\{0,1\}\times[0,1]}$ be an event, and for each vertex u of G let $K_n(u)$ by the connected component of $\omega_n \cap \mathfrak{F}$ containing u. Define a mass transport by sending mass $1/|K_n(u)|$ from each vertex u to every vertex $v \in K_n(u)$ such that $(G,u,v,\mathfrak{F},\theta) \in \mathscr{B}$ (it may be that no such vertices exist, in which case u sends no mass). Then the expected mass sent by the root is

$$\widehat{\mathbb{E}}\left[\frac{1}{|K_n(\rho)|}\sum_{v\in K_n(\rho)} \mathbb{1}((G,\rho,v,\mathfrak{F},\theta)\in\mathscr{B})\right] = \widehat{\mathbb{P}}((G,\rho,v_n,\mathfrak{F},\theta)\in\mathscr{B})$$

while the expected mass received by the root is

$$\widehat{\mathbb{E}}\left[\frac{1}{|K_n(\rho)|}\sum_{v\in K_n(\rho)} \mathbb{1}((G,v,\rho,\mathfrak{F},\theta)\in\mathscr{B})\right] = \widehat{\mathbb{P}}((G,v_n,\rho,\mathfrak{F},\theta)\in\mathscr{B}).$$

We conclude by applying the Mass-Transport Principle. □

We will also require the following simple lemma.

Lemma 6.4.7. *Let f be an essential spanning forest of G such that every component of f is one-ended.*

1. *For every edge e such that $e \notin f$ but e^+ and e^- are in the same component of f, let $C(f,e)$ denote the unique cycle contained in $f \cup \{e\}$. Then for every vertex u in G,*

$$\widehat{\mathbb{P}}(e_n(u) = e \text{ and } C(f,e) \subseteq \omega_1 \mid (G,\rho),\, \mathfrak{F} = f)$$
$$= \widehat{\mathbb{P}}(e_n(u) = e \text{ and } C(f,e) \subseteq \omega_1 \mid (G,\rho),\, \mathfrak{F} = U(f,e))$$

   for all $n \geq 0$.

2. *For every edge e of G, there exists $\kappa(f,e) > 0$ such that for every vertex u of G for which at least one endpoint of e is not contained in $T_f(u)$ and the component of u in $f \setminus \{e^-\}$ is infinite,*

$$\widehat{\mathbb{P}}(e_n(u) = e \mid (G,\rho),\, \mathfrak{F} = f) \geq \kappa(f,e)\widehat{\mathbb{P}}(e_n(u) = e \mid (G,\rho),\, \mathfrak{F} = U(f,e))$$

   for all $n \geq 0$.

Proof. Item (1) follows immediately from the observation that, under these assumptions, the set of vertices connected to u in $\omega_n \cap f$ and $\omega_n \cap U(f,e)$ are equal on the event that $C(f,e) \subseteq \omega_1$. We now prove item (2). If e^+ and e^- are in the same component of f or if $e^+, e^- \notin T_f(u)$ then the claim holds trivially by setting $\kappa(f,e) = 1$, so suppose not. Recall that $K_{\omega_n \cap f}(u)$ is defined to be the connected component of u in $\omega_n \cap f$. Define

$$\kappa_1(u,f;\omega_n) = \frac{1}{|K_{\omega_n \cap f}(u)|}$$

and

$$\kappa_2(u,f,e;\omega_n) = \sum_{\{v \in K_{\omega_n \cap f}(u)\,:\, d(v,e) \leq r\}} \frac{1}{|\{e' \in E : d(v,e') \leq r\}|}.$$

Then conditional on (G,ρ), $\mathfrak{F} = f$, and ω_n, the probability that $e_n(u) = e$ for each oriented edge e of G equals

$$\kappa_1(u,f;\omega_n)\kappa_2(u,f,e;\omega_n).$$

Let W denote the union of the finite components of $f \setminus \{e^+, e^-\}$. Our assumptions on e, u and f imply that $T_{U(f,e)}(u) \triangle T_f(u)$ is contained in W, so that

$$\kappa_1(u,f;\omega_n)^{-1} \;=\; |K_{\omega_n \cap f}(u)| \;\leq\; |K_{\omega_n \cap U(f,e)}(u)| + |W| \;=\; \kappa_1(u,U(f,e);\omega_n)^{-1} + |W|,$$

and so

$$\kappa_1(u, f; \omega_n) \geq \frac{1}{1+|W|}\kappa_1(u, U(f,e); \omega_n),$$

since $\kappa_1(u, U(f,e); \omega_n) \leq 1$. Let

$$\kappa_2^-(e) = \min\left\{|\{e' \in E : d(v,e') \leq r\}|^{-1} : v \in V(G), d(v,e) \leq r\right\} > 0.$$

Suppose that $\kappa_2(u, U(f,e), e; \omega_n) > 0$. Then there is a vertex x in the tree $K_{\omega_n \cap U(f,e)}(u)$ such that $d(x,e) \leq r$ and x is still connected to u in $K_{\omega_n \cap U(f,e)}(u) \setminus e$. This x is therefore also be connected to u in $\omega_n \cap f$, and so

$$\kappa_2(u, f, e; \omega_n) \geq |\{e' \in E : d(x,e') \leq r\}|^{-1} \geq \kappa_2^-(e)\,,$$

and thus,

$$\kappa_2(u, f, e; \omega_n) \geq \kappa_2^-(e)\mathbb{1}\big(\kappa_2(u, U(f,e), e; \omega_n) > 0\big).$$

But $\kappa_2(u, U(f,e), e; \omega_n)$ is bounded above by

$$\kappa_2(u, U(f,e), e; \omega_n) \leq \kappa_2^+(e) := \sum_{\{v:\, d(v,e) \leq r\}} \frac{1}{|\{e' \in E : d(v,e') \leq r\}|}$$

and so

$$\kappa_2(u, f, e; \omega_n) \geq \frac{\kappa_2^-(e)}{\kappa_2^+(e)}\kappa_2(u, U(f,e), e; \omega_n).$$

We obtain that

$$\kappa_1(u, f, e; \omega_n)\kappa_2(u, f, e; \omega_n) \geq \frac{\kappa_2^-(e)}{(1+|W|)\kappa_2^+(e)}\kappa_1(u, U(f,e); \omega_n)\kappa_2(u, U(f,e), e; \omega_n). \tag{6.4.4}$$

The claim follows by setting

$$\kappa(f,e) = \frac{\kappa_2^-(e)}{(1+|W|)\kappa_2^+(e)}$$

and taking expectations over ω_n in (6.4.4). □

Given $(G, \rho, \mathfrak{F}, \theta)$ and a positive $\delta > 0$, we say that an oriented edge e of G is δ **-update-friendly** if

1. $c(e)/c(e^-) \geq \delta$, and
2. $\kappa(\mathfrak{F}, e) \geq \delta$, and
3. if $e \notin \mathfrak{F}$ but e^+ and e^- are in the same component of $\mathfrak{F}$, then $C(\mathfrak{F}, e) \subseteq \omega_1$.

Note that if e is δ-update-friendly for $(G, \rho, \mathfrak{F}, \theta)$ then it is also δ-update-friendly for $(G, \rho, U(\mathfrak{F}, e), \theta)$. By assumption, there exists $\delta > 0$ such that with positive probability $T_{\mathfrak{F}}(\rho)$ has property $\mathscr{A}$ and there exists a good pivotal edge e for ρ at distance at most r from ρ in G such that e is δ-update-friendly.

Conditional on (G,ρ), for each edge e of G and $n \in \mathbb{Z}$, let $\mathscr{E}_e^n$ denote the event that e is δ-update-friendly and $e_n = e$. Write $\widehat{\mathbb{P}}_{(G,\rho)}$ for $\widehat{\mathbb{P}}$ conditioned on (G,ρ). Applying part (2) of Lemma 6.4.7 if e^+, e^- are both in $T_{\mathfrak{F}}(\rho)$ and part (1) otherwise, we deduce from the definition of δ-update-friendliness that for every event $\mathscr{B} \in \{0,1\}^{E(G)}$ such that $\mathsf{WUSF}_G(\mathfrak{F} \in \mathscr{B}) > 0$,

$$\begin{aligned}
\widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n \cap \{\mathfrak{F} \in \mathscr{B}\}) &= \widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n \mid \mathfrak{F} \in \mathscr{B})\mathsf{WUSF}_G(\mathfrak{F} \in \mathscr{B}) \\
&\geq \delta\widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n \mid \{U(\mathfrak{F},e) \in \mathscr{B}\})\mathsf{WUSF}_G(\mathfrak{F} \in \mathscr{B}) \\
&= \delta\mathbb{1}\left(\frac{c(e)}{c(e^-)} \geq \delta\right)\frac{\mathsf{WUSF}_G(\mathfrak{F} \in \mathscr{B})}{\mathsf{WUSF}_G(U(\mathfrak{F},e) \in \mathscr{B})} \\
&\qquad\qquad\qquad\qquad \cdot \widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n \cap \{U(\mathfrak{F},e) \in \mathscr{B}\}) \\
&\geq \delta^2\widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_e^n \cap \{U(\mathfrak{F},e) \in \mathscr{B}\}),
\end{aligned} \tag{6.4.5}$$

where the last inequality is by update-tolerance (Corollary 6.4.3). Update-tolerance also implies that this inequality holds trivially when $\mathsf{WUSF}_G(\mathfrak{F} \in \mathscr{B}) = 0$.

Fix $\varepsilon > 0$, and let R be sufficiently large that there exists an event $\mathscr{A}'$ that is measurable with respect to the σ-algebra generated by (G,ρ) and $\mathfrak{F} \cap B_G(\rho,R)$ and has $\widehat{\mathbb{P}}((G,\rho,\mathfrak{F}) \in \mathscr{A} \,\triangle\, \mathscr{A}') \leq \varepsilon$. Define the disjoint unions

$$\mathscr{E}^n := \bigcup_{c(e)/(e^-)\geq\delta} \mathscr{E}_e^n \quad \text{ and } \quad \mathscr{E}_R^n := \bigcup_{e^- \notin B_G(\rho,R)\,, c(e)/c(e^-)\geq\delta} \mathscr{E}_e^n.$$

Condition on (G,ρ), and let

$$\mathscr{B} = \{\omega \in \{0,1\}^E : (G,\rho,\omega) \in \mathscr{A}' \setminus \mathscr{A}\}.$$

Summing over (6.4.5) with this $\mathscr{B}$ yields that

$$\begin{aligned}
\widehat{\mathbb{P}}_{(G,\rho)}(\mathfrak{F} \in \mathscr{B}) &\geq \widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_R^n \cap \{\mathfrak{F} \in \mathscr{B}\}) \\
&\geq \delta^2\widehat{\mathbb{P}}_{(G,\rho)}(\mathscr{E}_R^n \cap \{U(\mathfrak{F},e_n) \in \mathscr{B}\})
\end{aligned}$$

and hence, taking expectations,

$$\widehat{\mathbb{P}}((G,\rho,\mathfrak{F}) \in \mathscr{A}' \setminus \mathscr{A}) \geq \delta^2\widehat{\mathbb{P}}\big(\mathscr{E}_R^n \cap \{(G,\rho,\mathfrak{F}) \in \mathscr{A}' \setminus \mathscr{A}\}\big).$$

By the definition of $\mathscr{A}'$ we have that

$$\mathscr{E}_R^n \cap \{(G,\rho,U(\mathfrak{F},e_n)) \in \mathscr{A}'\} = \mathscr{E}_R^n \cap \{(G,\rho,\mathfrak{F}) \in \mathscr{A}'\},$$

and so

$$\widehat{\mathbb{P}}((G,\rho,\mathfrak{F}) \in \mathscr{A}' \setminus \mathscr{A}) \\ \geq \delta^2 \widehat{\mathbb{P}}\Big(\mathscr{E}_R^n \cap \big\{(G,\rho,\mathfrak{F}) \in \mathscr{A}'\big\} \cap \big\{(G,\rho,U(\mathfrak{F},e_n)) \in \neg\mathscr{A}\big\}\Big). \tag{6.4.6}$$

Let $\mathscr{P}_n$ denote the event that e_n is a good pivotal edge for v_n. We claim that if $\mathscr{P}_n$ occurs and ρ is not in the past of v_n, then $T_{U(\mathfrak{F},e_n)}(\rho) = T_{U(\mathfrak{F},e_n)}(v_n)$ and $(G,\rho,U(\mathfrak{F},e_n)) \in \neg\mathscr{A}$. If e_n is a good internal, external or additive pivotal for v_n, then clearly ρ and v_n are in the same component of $U(\mathfrak{F},e_n)$, and, since e_n is pivotal for v_n we deduce that $(G,\rho,U(\mathfrak{F},e_n)) \in \neg\mathscr{A}$. If e_n is a good subtractive pivotal edge for v_n then the component of v_n in $\mathfrak{F} \setminus \{e_n^-, e_n^+\}$ is infinite and, since ρ is not in the past of v_n, ρ and v_n must be in the same component of $\mathfrak{F} \setminus \{e_n^-, e_n^+\}$. It follows that ρ and v_n are in the same component of $U(\mathfrak{F},e_n)$, and so $U(\mathfrak{F},e_n) \in \neg\mathscr{A}_\rho$ as before. Combining this with (6.4.6), we have

$$\widehat{\mathbb{P}}\big((G,\rho,\mathfrak{F}) \in \mathscr{A}' \setminus \mathscr{A}\big) \geq \delta^2 \widehat{\mathbb{P}}\big(\mathscr{E}_R^n \cap \big\{(G,\rho,\mathfrak{F}) \in \mathscr{A}'\big\} \cap \mathscr{P}_n \cap \big\{\rho \notin \mathrm{past}_{\mathfrak{F}}(v_n)\big\}\big)\,.$$

Lemma 6.4.8. $\widehat{\mathbb{P}}(\rho \in \mathrm{past}(v_n)) \to 0$ *as* $n \to \infty$.

Proof. By Lemma 6.4.6,

$$\widehat{\mathbb{P}}\big(\rho \in \mathrm{past}_{\mathfrak{F}}(v_n)\big) = \widehat{\mathbb{P}}\big(v_n \in \mathrm{past}_{\mathfrak{F}}(\rho)\big).$$

Observe that $\mathrm{past}(\rho)$ is finite, while the size of the component of ρ in $T_{\mathfrak{F}}(\rho) \cap \omega_n$ tends to infinity as $n \to \infty$. Since v_n is defined to be a uniform vertex of the this component, it follows that

$$\widehat{\mathbb{P}}(v_n \in \mathrm{past}_{\mathfrak{F}}(\rho) \mid (G,\rho,\mathfrak{F},\theta)) \xrightarrow[n\to\infty]{a.s.} 0$$

and the claim follows by taking expectations. ☐

Thus, taking n sufficiently large that $\widehat{\mathbb{P}}(\rho \in \mathrm{past}_{\mathfrak{F}}(v_n)) < \varepsilon$, we have that

$$\mathbb{P}((G,\rho,\mathfrak{F}) \in \mathscr{A}' \setminus \mathscr{A}) \geq \delta^2 \widehat{\mathbb{P}}(\mathscr{E}_R^n \cap \{(G,\rho,\mathfrak{F}) \in \mathscr{A}'\} \cap \mathscr{P}_n) - \delta^2\varepsilon.$$

By definition of $\mathscr{A}'$, we then have that

$$\mathbb{P}((G,\rho,\mathfrak{F}) \in \mathscr{A}' \setminus \mathscr{A}) \geq \delta^2 \widehat{\mathbb{P}}(\mathscr{E}_R^n \cap \{(G,\rho,\mathfrak{F}) \in \mathscr{A}\} \cap \mathscr{P}_n) - 2\delta^2\varepsilon\,.$$

We can further choose n to be sufficiently large that $\widehat{\mathbb{P}}(\mathscr{E}^n \setminus \mathscr{E}_R^n) \leq \varepsilon$, so that

$$\begin{aligned}\mathbb{P}((G,\rho,\mathfrak{F}) \in \mathscr{A}' \setminus \mathscr{A}) &\geq \delta^2 \widehat{\mathbb{P}}(\mathscr{E}^n \cap \{(G,\rho,\mathfrak{F}) \in \mathscr{A}\} \cap \mathscr{P}_n) - 3\delta^2\varepsilon \\ &= \delta^2 \widehat{\mathbb{P}}(\mathscr{E}^n \cap \{(G,v_n,\mathfrak{F}) \in \mathscr{A}\} \cap \mathscr{P}_n) - 3\delta^2\varepsilon \end{aligned} \tag{6.4.7}$$

where in the second equality we have used the fact that $\mathscr{A}$ is a component property. Observe that, by Lemma 6.4.6, the probability $\widehat{\mathbb{P}}(\mathscr{E}^n \cap \{(G, v_n, \mathfrak{F}) \in \mathscr{A}\} \cap \mathscr{P}_n) > 0$ does not depend on n. It does not depend on ε either, and so (6.4.7) contradicts the definition of $\mathscr{A}'$ when ε is taken to be sufficiently small. $\square$

Part III

Circle Packing and Planar Graphs

Chapter 7

Unimodular hyperbolic triangulations: circle packing and random walk

Summary. We show that the circle packing type of a unimodular random plane triangulation is parabolic if and only if the expected degree of the root is six, if and only if the triangulation is amenable in the sense of Aldous and Lyons [7]. As a part of this, we obtain an alternative proof of the Benjamini-Schramm Recurrence Theorem [50].

Secondly, in the hyperbolic case, we prove that the random walk almost surely converges to a point in the unit circle, that the law of this limiting point has full support and no atoms, and that the unit circle is a realisation of the Poisson boundary. Finally, we show that the simple random walk has positive speed in the hyperbolic metric.

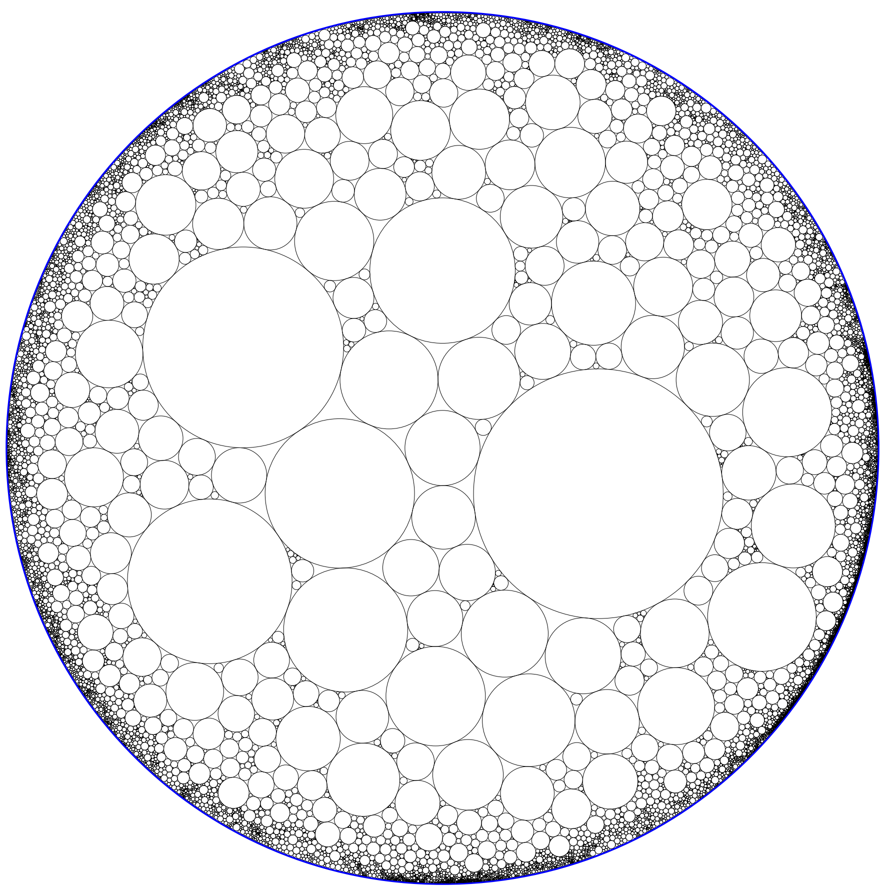

Figure 7.1: A circle packing of a random hyperbolic triangulation.

7.1 Introduction

A **circle packing** of a planar graph G is a set of circles with disjoint interiors in the plane, one for each vertex of G, such that two circles are tangent if and only if their corresponding vertices are adjacent in G. The Koebe-Andreev-Thurston Circle Packing Theorem [156, 221] states that every finite simple planar graph has a circle packing; if the graph is a triangulation (i.e. every face has three sides), the packing is unique up to Möbius transformations and reflections. He and Schramm [121, 122] extended this theorem to infinite, one-ended, simple triangulations, showing that each such triangulation admits a locally finite circle packing either in the Euclidean plane or in the hyperbolic plane (identified with the interior of the unit disc), but not both. See Section 7.3.4 for precise details. This result is a discrete analogue of the Uniformization Theorem, which states that every simply connected, non-compact Riemann surface is conformally equivalent to either the plane or the disc (indeed, there are deep connections between circle packing and conformal maps, see [202, 215] and references therein). Accordingly, a triangulation is called **CP parabolic** if it can be circle packed in the plane and **CP hyperbolic** otherwise.

Circle packing has proven instrumental in the study of random walks on planar graphs [47, 50, 107, 121]. For graphs with bounded degrees, a rich theory has been established connecting the geometry of the circle packing and the behaviour of the random walk. Most notably, a one-ended, bounded degree triangulation is CP hyperbolic if and only if random walk on it is transient [121] and in this case it is also non-Liouville, i.e. admits non-constant bounded harmonic functions [47].

The goal of this work is to develop a similar, parallel theory for *random* triangulations. Particular motivations come from the Markovian hyperbolic triangulations constructed recently in [23] and [75]. These are hyperbolic variants of the UIPT [24] and are conjectured to be the local limits of uniform triangulations in high genus. Another example is the Poisson-Delaunay triangulation in the hyperbolic plane, studied in [49] and [37]. All these triangulations have unbounded degrees, rendering existing methods ineffective (for example methods used in [17, 47, 121]).

Indeed, in the absence of bounded degree the existing theory fails in many ways. For example, in a circle packing of a triangulation with bounded degrees, radii of adjacent circles have uniformly bounded ratios (a fact known as the Ring Lemma [200]). The absence of such a uniform bound invalidates important resistance estimates. This is not a mere technicality: one can add extra circles in the interstices of the circle packing of the triangular lattice to give the random walk drift in arbitrary directions. This does not change the circle packing type, but allows construction of a graph that is CP parabolic but transient or even non-Liouville. Indeed, the main effort in [107] was to overcome this sole obstacle in order to prove that the UIPT is recurrent.

The hyperbolic random triangulations of [75] and [37] make up for having unbounded degrees by a different useful property: *unimodularity* (essentially equivalent to *reversibility*, see Sections 7.3.1 and 7.3.2). This allows us to apply probabilistic and ergodic arguments in place of the analytic arguments appropriate to the bounded degree case. Our first main theorem establishes a probabilistic characterisation of the CP type for unimodular random rooted triangulations, and connects it to the geometric property of *invariant (non-)amenability*, which we define in Section 7.3.3.

Theorem 7.1.1. *Let (G, ρ) be an infinite, simple, one-ended, ergodic unimodular random rooted planar triangulation. Then either*

> $\mathbb{E}[\deg(\rho)] = 6$*, in which case (G, ρ) is invariantly amenable and almost surely CP parabolic,*

or else

> $\mathbb{E}[\deg(\rho)] > 6$*, in which case (G, ρ) is invariantly non-amenable and almost surely CP hyperbolic.*

This theorem can be viewed as a local-to-global principle for unimodular triangulations. That is, it allows us to identify the circle packing type and invariant amenability, both global properties, by calculating the expected degree, a very local quantity. For example, if (G, ρ) is a simple, one-ended triangulation that is obtained as a local limit of planar graphs, then by Euler's formula and Fatou's lemma its average degree is at most 6, so that Theorem 7.1.1 implies it is almost surely CP parabolic. If in addition (G, ρ) has bounded degrees, then it is recurrent by He-Schramm [121]. In particular, this gives an alternative proof of the Benjamini-Schramm Recurrence Theorem [50] in the primary case of a one-ended limit. We handle the remaining cases in Section 7.4.1. Unlike the proof of [50], whose main ingredient is a quantitative estimate for finite circle packings [50, Lemma 2.3], our method works with infinite triangulations directly and implies the following generalisation:

Proposition 7.1.2. *Any unimodular, simple, one-ended random rooted planar triangulation (G, ρ) with bounded degrees and $\mathbb{E}[\deg(\rho)] = 6$ is almost surely recurrent.*

This trivially extends the Benjamini-Schramm result, since any local limit of finite planar graphs is unimodular. An important open question is whether every unimodular random graph is a Benjamini-Schramm limit of finite graphs. In a forthcoming paper [20], we show that any unimodular planar graph G is a limit of some sequence of finite graphs G_n, and that if G is a triangulation with $\mathbb{E}[\deg(\rho)] = 6$ then G_n can also be taken to be planar. In particular, any graph to which Proposition 7.1.2 applies is also a local limit of finite planar graphs with bounded degrees. Consequently there are no graphs to which this result applies and the Benjamini-Schramm Theorem does not. Note however, that for a given unimodular planar triangulation, it may not be obvious how to find this sequence of graphs. We remark that the dichotomy of Theorem 7.1.1 has many extensions, applying to more general maps and holding further properties equivalent. We address these in [20]. See [34] and [121, Theorem 10.2] for earlier connections between the CP type and degree distributions in the deterministic setting.

Our method of proof relies on the deep theorem of Schramm [206] that the circle packing of a triangulation in the disc or the plane is unique up to Möbius transformations fixing the disc or the plane as appropriate. We use this fact throughout the paper in an essential way: it implies that any quantity derived from the circle packing in the disc or the plane that is invariant to Möbius transformations is determined by the graph G and not by our choice of circle packing. Key examples of such quantities are angles between adjacent edges in the associated drawings with hyperbolic or

Euclidean geodesics (see Section 7.4), hyperbolic radii of circles in the hyperbolic case, and ratios of Euclidean radii in the parabolic case.

Boundary Theory. Throughout, we realize the hyperbolic plane as the Poincaré disc $\{|z| < 1\}$ with metric d_{hyp}. The unit circle $\{|z| = 1\}$ is the boundary of the hyperbolic plane in several geometric and probabilistic senses. For a general graph embedded in the hyperbolic plane, the unit circle may or may not coincide with probabilistic notions of the graph's boundary.

When a bounded degree triangulation is circle packed in the disc, Benjamini and Schramm [47] showed that the random walk converges to a point in the circle almost surely and that the law of the limit point has full support and no atoms. More recently, it was shown by the first and third authors together with Barlow and Gurel-Gurevich [17] that the unit circle is a realisation of both the Poisson and Martin boundaries of the triangulation. Similar results regarding square tiling were obtained in [48] and [100].

Again, these theorems fail for some triangulations with unbounded degrees. Starting with any CP hyperbolic triangulation, one can add circles in the interstices of the packing so as to create drifts along arbitrary paths. In this way, one can force the random walk to spiral in the unit disc and not converge to any point in the boundary. One can also create a graph for which the walk can converge to a single boundary point from two or more different angles each with positive probability, so that the exit measure is atomic and the unit circle is no longer a realisation of the Poisson boundary. Our next result recovers the boundary theory in the unimodular setting.

When $\mathcal{C}$ is a circle packing of a graph G in the disc $\mathbb{D}$, we write $\mathcal{C} = (z, r)$ where $z(v)$ is the (Euclidean) centre of the circle corresponding to v, and $r(v)$ is its Euclidean radius. Recall that the hyperbolic metric on the unit disc is defined by

$$|d_{hyp}(z)| = \frac{2|dz|}{1 - |z|^2},$$

and that circles in the Euclidean metric are also hyperbolic circles (with different centres and radii). We write $z_h(v)$ and $r_h(v)$ for the hyperbolic centre and radius of the circle corresponding to v. We use P_v^G and E_v^G to denote the probability and expectation (conditioned on G) with respect to random walk $(X_n)_{n\geq 0}$ on G started from a vertex v.

Theorem 7.1.3. *Let (G, ρ) be a simple, one-ended, CP hyperbolic unimodular random planar triangulation with $\mathbb{E}[\deg^2(\rho)] < \infty$. Let $\mathcal{C}$ be a circle packing of G in the unit disc, and let (X_n) be a simple random walk on G. The following hold conditional on (G, ρ) almost surely:*

1. *$z(X_n)$ and $z_h(X_n)$ both converge to a (random) point denoted $\Xi \in \partial\mathbb{D}$,*
2. *The law of Ξ has full support $\partial\mathbb{D}$ and no atoms.*
3. *$\partial\mathbb{D}$ is a realisation of the Poisson boundary of G. That is, for every bounded harmonic*

function h on G there exists a bounded measurable function $g : \partial\mathbb{D} \to \mathbb{R}$ such that

$$h(v) = E_v^G[g(\Xi)].$$

We refer to the law of Ξ conditional on (G, ρ) as the **exit measure** from v. In Section 7.7 we extend this result to weighted and non-simple triangulations, with the obvious changes. One ingredient in the proof of the absence of atoms is a more general observation, Lemma 7.5.2, which states roughly that exit measures on boundaries of stationary graphs are either non-atomic or trivial almost surely.

Our final result relates exponential decay of the Euclidean radii along the random walk to speed in the hyperbolic metric.

Theorem 7.1.4. *Let (G, ρ) be a simple, one-ended, CP hyperbolic unimodular random rooted planar triangulation with $\mathbb{E}[\deg^2(\rho)] < \infty$ and let $\mathcal{C}$ be a circle packing of G in the unit disc. Then almost surely*

$$\lim_{n\to\infty} \frac{d_{hyp}(z_h(\rho), z_h(X_n))}{n} = \lim_{n\to\infty} \frac{-\log r(X_n)}{n} > 0.$$

In particular, both limits exist. Moreover, the limits do not depend on the choice of packing, and if (G, ρ) is ergodic then this limit is an almost sure constant.

Thus the random walk (X_n) has positive asymptotic speed in the hyperbolic metric, the Euclidean radii along the walk decay exponentially, and the two rates agree.

Organization of the paper. In Section 7.2 we review the motivating examples of unimodular hyperbolic random triangulations to which our results apply. In Section 7.3.1 and Section 7.3.2 we give background on unimodularity, reversibility and related topics. In Section 7.3.3 we recall Aldous and Lyons's notion of invariant amenability [7] and prove one of its important consequences. In Section 7.3.4 we recall the required results on circle packing and discuss measurability. Section 7.4 contains the proof of Theorem 7.1.1 as well as a discussion of how to handle the remaining (easier) cases of the Benjamini-Schramm Theorem. Theorem 7.1.3 is proved in Section 7.5 and Theorem 7.1.4 is proved in Section 7.6. Background on the Poisson boundary is provided before the proof of Theorem 7.1.3(3) in Section 7.5.3. In Section 7.7 we discuss extensions of our results to non-simple and weighted triangulations. We end with some open problems in Section 7.8.

7.2 Examples

Benjamini-Schramm limits of random maps have been objects of great interest in recent years, serving as discrete models of 2-dimensional quantum gravity. Roughly, the idea is to consider a uniformly random map from some class of rooted maps (e.g. all triangulations or quadrangulations of the sphere of size n) and take a local limit as the size of the maps tends to infinity. The first such construction was the UIPT [24]; see also [15, 18, 19, 39, 76].

Curien's PSHT. Recently, hyperbolic versions of the UIPT and related maps have been constructed: half-plane versions in [23] and full-plane versions in [75]. These are constructed directly, and are believed but not yet known to be the limits of finite maps (see below). The full plane triangulations form a one (continuous) parameter family $\{T_\kappa\}_{\kappa\in(0,2/27)}$ (known as the PSHT, for *Planar Stochastic Hyperbolic Triangulation*). They are reversible and ergodic, have anchored expansion and are therefore invariantly non-amenable. The degree of the root in T_κ is known to have an exponential tail, so that all of its moments are finite. These triangulations are not simple, so our main results do not apply to them directly, but by considering their *simple cores* we are still able to obtain a geometric representation of their Poisson boundary (see Section 7.7).

Benjamini-Schramm limits of maps in high genus. It is conjectured that the PSHT T_κ is the Benjamini-Schramm limit of the uniform triangulation with n vertices of a surface of genus $\lfloor \theta n \rfloor$, for some $\theta = \theta(\kappa)$ (see e.g. [198] for precise definitions of maps on general surfaces). In our upcoming paper [20] we prove that all one-ended unimodular random rooted planar triangulations are also Benjamini-Schramm limits of finite triangulations. If the triangulation has expected degree greater than 6, then the finite approximating triangulations necessarily have genus linear in their size.

In the context of circle packing, it may be particularly interesting to take the Benjamini-Schramm limit (T,ρ) of the uniform *simple* triangulation with n vertices of the $\lfloor \theta n \rfloor$-holed torus T_n. This limit (which we conjecture exists) should be a simple variant of the PSHT. Letting ρ_n be a uniformly chosen root of T_n, it should also be the case that $\mathbb{E}[\deg(\rho_n)] \to \mathbb{E}[\deg(\rho)] > 6$ and $\mathbb{E}[\deg(\rho)^2] < \infty$, so that our results would be applicable to the circle packing of (T,ρ).

Delaunay triangulations of the hyperbolic plane. Start with a Poisson point process in the hyperbolic plane with intensity λ times the hyperbolic area measure, and add a root point at the origin. Consider now the **Delaunay triangulation** with this point process as its vertex set, where three vertices u, v, w form a triangle if the circle through u, v, w contains no other points of the process. This triangulation, known as the Poisson-Delaunay triangulation, is naturally embedded in the hyperbolic plane with hyperbolic geodesic edges. These triangulations, studied in [37, 49], are unimodular when rooted at the point at the origin. They are known to have anchored expansion [37] and are therefore invariantly non-amenable. (We also get a new proof of non-amenability from Theorem 7.1.1, as one can show the expected degree to be greater than six by transporting angles as in the proof of Theorem 7.1.1.) The Poisson-Delaunay triangulations are also simple and one-ended, and the degree of the root has finite second moment, so that our results apply directly to their circle packings.

7.3 Background and definitions

7.3.1 Unimodular random graphs and maps

Unimodularity of graphs (both fixed and random) has proven to be a useful and natural property in a number of settings. We give here the required definitions and some of their consequences, and refer the reader to [7, 173] for further background.

A **rooted graph** (G, ρ) is a graph $G = (V, E)$ with a distinguished vertex ρ called the **root**. We will allow our graphs to contain self-loops and multiple edges, and refer to graphs without either as **simple**. A graph is said to be **one-ended** if the removal of any finite set of vertices leaves precisely one infinite connected component. A graph isomorphism between two rooted graphs is a rooted graph isomorphism if it preserves the root.

A **map** is a proper (that is, with non-intersecting edges) embedding of a connected graph into a surface, viewed up to orientation preserving homeomorphisms of the surface, so that all connected components of the complement (called **faces**) are topological discs.[8] The map is **planar** if the surface is homeomorphic to an open subset of the sphere, and is **simply connected** if the surface is homeomorphic to the sphere or the plane. A map is a **triangulation** if every face is incident to exactly three edges. Note that an infinite planar triangulation is simply connected if and only if it is one-ended.

Every connected graph G can be made into a metric space by endowing it with the shortest path metric d_G. By abuse of notation, we use the ball $B_n(G, u)$ to refer both to the set of vertices $\{v \in V : d_G(u, v) \leq n\}$ and the induced subgraph on this set, rooted at u. The balls in a map inherit a map structure from the full map.

The **local topology** on the space of rooted connected graphs (introduced in [50]) is the topology induced by the metric

$$d_{\mathrm{loc}}\big((G,\rho),(G',\rho')\big) = e^{-R} \quad \text{where} \quad R = \sup\big\{n \geq 0 : B_n(G,\rho) \cong B_n(G',\rho')\big\}.$$

The local topology on rooted maps is defined similarly by requiring the isomorphism of the balls to be an isomorphism of rooted maps. We denote by $\mathcal{G}_\bullet$ and $\mathcal{M}_\bullet$ the spaces of isomorphism classes of rooted connected graphs and of maps with their respective local topologies. **Random rooted graphs and maps** are Borel random variables taking values in these spaces.

Several variants of these spaces will also be of use. A (countably) **marked** graph is a graph together with a mark function $m : V \cup E \to M$ which gives every edge and vertex a mark in some countable set M. A graph isomorphism between marked graphs is an isomorphism of marked graphs if it preserves the marks. The local topologies on rooted marked graphs and maps is defined in the obvious way. These spaces are denoted $\mathcal{G}_\bullet^M$ and $\mathcal{M}_\bullet^M$. Sometimes we will consider maps with marks only on vertices or only on edges; these fit easily into our framework. Marked graphs are

[8] There is an additional constraint regarding boundaries of faces of infinite degree. However, this condition is automatically satisfied for triangulations and for simply connected maps, so that we need not worry about it in this paper.

special cases of what Aldous and Lyons [7] call **networks**, for which the marks may take values in any separable complete metric space.

Similarly, we define $\mathcal{G}_{\bullet\bullet}$ (resp. $\mathcal{M}_{\bullet\bullet}$) to be the spaces of doubly rooted (that is, with a distinguished ordered pair of vertices) connected graphs (resp. maps) (G, u, v). These spaces, along with their marked versions, are equipped with natural variants of the local topology. All such spaces we consider are Polish.

A **mass transport** is a non-negative Borel function $f : \mathcal{G}_{\bullet\bullet} \to \mathbb{R}_+$. A random rooted graph (G, ρ) is said to be **unimodular** if it satisfies the **mass transport principle**: for any mass transport f,

$$\mathbb{E}\left[\sum_{v \in V(G)} f(G, \rho, v)\right] = \mathbb{E}\left[\sum_{v \in V(G)} f(G, v, \rho)\right].$$

In other words,

'Expected mass out equals expected mass in.'

This definition generalises naturally to define unimodular marked graphs and maps. Importantly, any finite graph G with a uniformly chosen root vertex ρ satisfies the mass transport principle.

The laws of unimodular random rooted graphs form a weakly closed, convex subset of the space of probability measures on $\mathcal{G}_\bullet$, so that weak limits of unimodular random graphs are unimodular. In particular, a weak limit of finite graphs with uniformly chosen roots is unimodular: such a limit of finite graphs is referred to as a **Benjamini-Schramm limit**. It is a major open problem to determine whether all unimodular random rooted graphs arise as Benjamini-Schramm limits of finite graphs [7, §10]. As mentioned in Section 7.2, we provide a positive solution to this problem in the planar case in the upcoming work [20], proving that every simply connected unimodular random rooted planar map is a Benjamini-Schramm limit of finite maps.

A common use of the mass transport principle to obtain proofs by contradiction is the following. If (G, ρ) is a unimodular random rooted graph and f is a mass transport such that the mass sent out from each vertex $\sum_v f(G, u, v) \leq M$ is uniformly bounded almost surely, then almost surely there are no vertices that receive infinite mass: if vertices receiving infinite mass were to exist with positive probability, the root would be such a vertex with positive probability [7, Lemma 2.3], contradicting the mass transport principle.

7.3.2 Random walk, reversibility and ergodicity

Recall that the **simple random walk** on a graph is the Markov chain that chooses X_{n+1} from among the neighbours of X_n weighted by the number of shared edges. Define $\mathcal{G}_\leftrightarrow$ (resp. $\mathcal{M}_\leftrightarrow$) to be spaces of isomorphism classes of graphs (resp. maps) equipped with a bi-infinite path $(G, (x_n)_{n\in\mathbb{Z}})$, which we endow with a natural variant of the local topology. When (G, ρ) is a random graph or map, we let $(X_n)_{n\geq 0}$ and $(X_{-n})_{n\geq 0}$ be two independent simple random walks started from ρ and consider $(G, (X_n)_{n\in\mathbb{Z}})$ to be a random element of $\mathcal{G}_\leftrightarrow$ or $\mathcal{M}_\leftrightarrow$ as appropriate.

A random rooted graph (G, ρ) is **stationary** if $(G, \rho) \overset{d}{=} (G, X_1)$ and **reversible** if $(G, \rho, X_1) \overset{d}{=} (G, X_1, \rho)$ as doubly rooted graphs. Equivalently, (G, ρ) is reversible if and only if $(G, (X_n)_{n\in\mathbb{Z}})$ is stationary with respect to the shift:

$$(G, (X_n)_{n\in\mathbb{Z}}) \overset{d}{=} (G, (X_{n+k})_{n\in\mathbb{Z}}) \text{ for every } k \in \mathbb{Z}.$$

To see this, it suffices to prove that if (G, ρ) is a reversible random graph then $(X_1, \rho, X_{-1}, X_{-2}, \ldots)$ has the law of a simple random walk started from X_1. But $(\rho, X_{-1}, \ldots)$ is a simple random walk started from ρ independent of X_1 and, conditional on (G, X_1), reversibility implies that ρ is uniformly distributed among the neighbours of X_1, so $(X_1, \rho, X_{-1}, X_{-2}, \ldots)$ has the law of a simple random walk as desired.

We remark that if (G, ρ) is stationary but not necessarily reversible, it is still possible to extend the walk to a doubly infinite path $(X_n)_{n\in\mathbb{Z}}$ so that G is stationary along the path. The difference is that in the reversible case the past $(X_n)_{n\leq 0}$ is itself a simple random walk with the same law as the future.

Reversibility is related to unimodularity via the following bijection, which is implicit in [7] and proven explicitly in [38]: if (G, ρ) is reversible, then biasing by $\deg(\rho)^{-1}$ (i.e. reweighing the law of (G, ρ) by the Radon-Nikodym derivative $\deg(\rho)^{-1}/\mathbb{E}[\deg(\rho)^{-1}]$) gives an equivalent unimodular random rooted graph, and conversely if (G, ρ) is a unimodular random rooted graph with finite expected degree, then biasing by $\deg(\rho)$ gives an equivalent reversible random rooted graph. Thus, the laws of reversible random rooted graphs are in bijection with the laws of unimodular random rooted graphs for which the root degree has finite expectation.

An event $A \subset \mathcal{G}_{\leftrightarrow}$ is said to be **invariant** if $(G, (X_n)_{n\in\mathbb{Z}}) \in A$ implies $(G, (X_{n+k})_{n\in\mathbb{Z}}) \in A$ for each $k \in \mathbb{Z}$. A reversible or unimodular random graph is said to be **ergodic** if the law of $(G, (X_n)_{n\in\mathbb{Z}})$ gives each invariant event probability either zero or one. An event $A \subseteq \mathcal{G}_\bullet$ is **rerooting-invariant** if $(G, \rho) \in A$ implies $(G, v) \in A$ for every vertex v of G.

Theorem 7.3.1 (Characterisation of ergodicity [7, §4])**.** *Let (G, ρ) be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$ (resp. a reversible random rooted graph). The following are equivalent.*

1. *(G, ρ) is ergodic.*
2. *Every rerooting-invariant event $A \subseteq \mathcal{G}_\bullet$ has probability in $\{0, 1\}$.*
3. *The law of (G, ρ) is an extreme point of the weakly closed convex set of laws of unimodular (resp. reversible) random rooted graphs.*

(The equivalence of items 2 and 3 holds for unimodular random rooted graphs without the assumption of finite expected degree.) A consequence of the extremal characterisation is that every unimodular random rooted graph is a *mixture* of ergodic unimodular random rooted graphs, meaning that it may be sampled by first sampling a random law of an ergodic unimodular random rooted graph, and then sampling from this randomly chosen law - this is known as an ergodic decomposition and its existence is a consequence of Choquet's Theorem. In particular, whenever we want to prove that a unimodular random rooted graph with some almost sure property also has some other almost sure property, it suffices to consider the ergodic case. The same comment applies for reversible random rooted graphs.

7.3.3 Invariant amenability

We begin with a brief review of general amenability, before combining it with unimodularity for the notion of invariant amenability. We refer the reader to [173, §6] for further details on amenability in general, and [7, §8] for invariant amenability.

A **weighted** graph is a graph together with a weight function $w : E \to \mathbb{R}_+$. Unweighted multigraphs may always be considered as weighted graphs by setting $w \equiv 1$. The weight function is extended to vertices by $w(x) = \sum_{e \ni x} w(e)$, and (with a slight abuse of notation) to sets of edges or vertices by additivity. The simple random walk $X = (X_n)_{n \geq 0}$ on a weighted graph is the Markov chain on V with transition probabilities $p(x, y) = w(x, y)/w(x)$. Here, our graphs are allowed to have infinite degree provided $w(v)$ is finite for every vertex.

The (edge) **Cheeger constant** of an infinite weighted graph is defined to be

$$\mathbf{i}_E(G) = \inf\left\{\frac{w(\partial_E W)}{w(W)} : \emptyset \neq W \subset V \text{ finite}\right\}$$

where $\partial_E W$ denotes the set of edges with exactly one end in W. A graph is said to be **amenable** if its Cheeger constant is zero and **non-amenable** if it is positive.

The Markov operator associated to simple random walk on G is the bounded, self-adjoint operator from $L^2(V, w)$ to itself defined by $(Pf)(u) = \sum p(u, v)f(v)$. The norm of this operator is commonly known as the **spectral radius** of the graph. If $u, v \in V$ then the transition probabilities are given by $p_n(u, v) = \left\langle P^n \mathbb{1}_v, \mathbb{1}_u / w(u)\right\rangle_w$, so that, by Cauchy-Schwarz,

$$p_n(u, v) \leq \sqrt{\frac{w(v)}{w(u)}} \|P\|_w^n \tag{7.3.1}$$

and in fact $\|P\|_w = \limsup_{n \to \infty} p_n(u, v)^{1/n}$. A fundamental result, originally proved for Cayley graphs by Kesten [152], is that the spectral radius of a weighted graph is less than one if and only if the graph is non-amenable (see [173, Theorem 6.7] for a modern account). As an immediate consequence, non-amenable graphs are transient for simple random walk.

Invariant amenability

There are natural notions of amenability and expansion for unimodular random networks due to Aldous and Lyons [7]. A **percolation** on a unimodular random rooted graph (G, ρ) is a random assignment of $\omega : E \cup V \to \{0, 1\}$ such that the marked graph (G, ρ, ω) is unimodular. We think of ω as a random subgraph of G consisting of the 'open' edges and vertices $\omega(e) = 1$, $\omega(v) = 1$, and may assume without loss of generality that if an edge is open then so are both of its endpoints.

The **cluster** $K_\omega(v)$ at a vertex v is the connected component of v in ω, i.e. the set of vertices for which there is a path of open edges to v (by convention, if $\omega(v) = 0$ or if there are no open edges touching v we put $K_\omega(v) = \{v\}$). A percolation is said to be **finitary** if all of its clusters are finite almost surely. The **invariant Cheeger constant** of an ergodic unimodular random rooted graph (G, ρ) is defined to be

$$\mathbf{i}^{\mathrm{inv}}((G,\rho)) = \inf\left\{\mathbb{E}\left[\frac{|\partial_E K_\omega(\rho)|}{|K_\omega(\rho)|}\right] : \omega \text{ a finitary percolation on } (G,\rho)\right\}. \tag{7.3.2}$$

The invariant Cheeger constant is closely related to another quantity: mean degrees in finitary percolations. Let $\deg_\omega(\rho)$ denote the degree of ρ in ω (seen as a subgraph; if $\rho \notin \omega$ we set $\deg_\omega(\rho) = 0$) and let

$$\alpha((G,\rho)) = \sup\left\{\mathbb{E}\left[\deg_\omega(\rho)\right] : \omega \text{ a finitary percolation on } (G,\rho)\right\}.$$

An easy application of the mass transport principle [7, Lemma 8.2] shows that, for any finitary percolation ω,

$$\mathbb{E}[\deg_\omega(\rho)] = \mathbb{E}\left[\frac{\sum_{v\in K_\omega(\rho)} \deg_\omega(v)}{|K_\omega(\rho)|}\right].$$

It follows that

$$\mathbb{E}[\deg(\rho)] = \mathbf{i}^{\mathrm{inv}}((G,\rho)) + \alpha((G,\rho))$$

so that if $\mathbb{E}[\deg(\rho)] < \infty$ then $\mathbf{i}^{\mathrm{inv}}((G,\rho))$ is positive if and only if $\alpha((G,\rho))$ is strictly smaller than $\mathbb{E}[\deg(\rho)]$.

We say that an ergodic unimodular random rooted graph (G,ρ) is **invariantly amenable** if $\mathbf{i}^{\mathrm{inv}}((G,\rho)) = 0$ and **invariantly non-amenable** otherwise. Note that this is a property of the *law* of (G,ρ) and not of an individual graph. We remark that what we are calling invariant amenability was called amenability when it was introduced by Aldous and Lyons [7, §8]. We qualify it as invariant to distinguish it from the more classical notion, which we also use below. While any invariantly amenable graph is trivially amenable, the converse is generally false. An example is a 3-regular tree where each edge is replaced by a path of independent length with unbounded distribution; see [7] for a more detailed discussion.

An important property of invariantly non-amenable graphs was first proved for Cayley graphs by Benjamini, Lyons and Schramm [45]. Aldous and Lyons [7] noted that the proof carried through with minor modifications to the case of invariantly non-amenable unimodular random rooted graphs, but did not provide a proof. As this property is crucial to our arguments, we provide a proof for completeness, which the reader may wish to skip. When (G,ρ) is an ergodic unimodular random rooted graph, we say that a percolation ω on G is **ergodic** if (G,ρ,ω) is ergodic as a unimodular random rooted marked graph. The following is stated slightly differently from both Theorem 3.2 in [45] and Theorem 8.13 in [7].

Theorem 7.3.2. *Let (G, ρ) be an invariantly non-amenable ergodic unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then G admits an ergodic percolation ω so that $\mathbf{i}_E(\omega) > 0$ and vertices in ω have uniformly bounded degrees in G.*

Let us stress that the condition of uniformly bounded degrees is for the degrees in the full graph G, and not the degrees in the percolation.

Remark 7.3.3. This theorem plays the same role for invariant non-amenability as Virág's oceans and islands construction [227] does for anchored expansion [16, 37, 227]. In particular, it gives us a percolation ω such that the induced network $\bar{\omega}$ is non-amenable (see the proof of Lemma 7.5.1).

Proof. Let ω_0 be the percolation induced by vertices of G of degree at most M and the edges connecting any two such vertices. By monotone convergence, and since $\alpha((G, \rho)) < \mathbb{E} \deg(\rho)$, we can take M to be large enough that $\mathbb{E} \deg_{\omega_0}(\rho) > \alpha((G, \rho))$. This gives a percolation with bounded degrees. We shall modify it further to get non-amenability as follows. Fix $\delta > 0$ by

$$3\delta = \mathbb{E}[\deg_{\omega_0}(\rho)] - \alpha((G, \rho)).$$

Construct inductively a decreasing sequence of site percolations ω_n as follows. Given ω_n, let η_n be independent Bernoulli(1/2) site percolations on ω_n, and for each set of vertices W let $\partial_E^{\omega_n} W$ denote the set of edges of ω_n in the boundary of W. If K is a finite connected cluster of η_n, with small boundary in ω_n, we remove it to construct ω_{n+1}. More precisely, let $\omega_{n+1} = \omega_n \setminus \gamma_n$, where γ_n is the subgraph of ω_n induced by the vertex set

$$\bigcup \big\{ K : \ K \text{ a finite cluster in } \eta_n \text{ with } |\partial_E^{\omega_n}(K)| < \delta |K| \big\}.$$

Let $\omega = \cap \omega_n$ be the limit percolation, which is clearly ergodic. We shall show below that $\omega \neq \emptyset$. Any finite connected set in ω appears as a connected cluster in η_n for infinitely many n. If such a set S has $|\partial_E^{\omega} S| < \delta |S|$ then it would have been removed at some step, and so ω has $|\partial_E^{\omega} S| \geq \delta |S|$ for all finite connected S. Since degrees are bounded by M, this implies $\mathbf{i}_E(\omega) \geq \delta / M > 0$.

It remains to show that $\omega \neq \emptyset$. For some n, and any vertex u, let $K(u)$ be its cluster in η_n. Consider the mass transport

$$f_n(u, v) = \begin{cases} \deg_{\omega_n}(v)/|K(u)| & u \in \gamma_n \text{ and } v \in K(u), \\ E(v, K(u))/|K(u)| & u \in \gamma_n \text{ and } v \in \omega_n \setminus K(u), \\ 0 & u \notin \gamma_n. \end{cases}$$

Here $E(v, K(u))$ is the number of edges between v and $K(u)$. We have that the total mass into v is the difference $\deg_{\omega_n}(v) - \deg_{\omega_{n+1}}(v)$ (where the degree is 0 for vertices not in the percolation) while the mass sent from a vertex $v \in \gamma_n$ is twice the number of edges with either end in $K(v)$, divided by $|K(v)|$. Applying the mass transport principle we get

$$\mathbb{E}[\deg_{\omega_n}(\rho) - \deg_{\omega_{n+1}}(\rho)] = \mathbb{E}\left[\frac{\sum_{v\in K(\rho)} \deg_{\gamma_n}(v) + 2|\partial_E^{\omega_n}(K(\rho))|}{|K(\rho)|}\mathbb{1}_{\rho\in\gamma_n}\right]. \tag{7.3.3}$$

By a second transport, of $\deg_{\gamma_n}(u)/|K(u)|$ from every $u \in \gamma_n$ to each $v \in K(u)$, we see that

$$\mathbb{E}\left[\frac{\sum_{v\in K(\rho)} \deg_{\gamma_n}(v)}{|K(\rho)|}\mathbb{1}_{\rho\in\gamma_n}\right] = \mathbb{E}[\deg_{\gamma_n}(\rho)]. \tag{7.3.4}$$

Additionally, on the event $\{\rho \in \gamma_n\}$, we have by definition that

$$|\partial_E^{\omega_n}(K(\rho))|/|K(\rho)| \leq \delta.$$

Plugging these two in (7.3.3) gives

$$\mathbb{E}[\deg_{\omega_n}(\rho) - \deg_{\omega_{n+1}}(\rho)] \leq \mathbb{E}[\deg_{\gamma_n}(\rho)] + 2\delta\mathbb{P}(\rho \in \gamma_n).$$

Let $\gamma = \cup_{n\geq 1}\gamma_n$, which is a percolation since it is defined as a measurable, automorphism invariant function of (G, ρ) and the i.i.d. sequence of Bernoulli percolations (η_n). Note the percolations γ_n are disjoint, so that the event $\rho \in \gamma_n$ can occur for at most one n and that γ is a finitary percolation. Thus $\mathbb{E}[\deg_\gamma(\rho)] = \sum_n \mathbb{E}[\deg_{\gamma_n}(\rho)] \leq \alpha((G, \rho))$. Also, $\sum_n \mathbb{P}(\rho \in \gamma_n) \leq 1$. Summing over n gives

$$\mathbb{E}[\deg_{\omega_0}(\rho) - \deg_\omega(\rho)] \leq \alpha((G, \rho)) + 2\delta.$$

The definition of δ leaves $\mathbb{E}[\deg_\omega(\rho)] \geq \delta$. Thus ω is indeed non-empty as claimed, completing the proof. □

7.3.4 Circle packings and vertex extremal length

Recall that a **circle packing** $\mathcal{C}$ is a collection of discs of disjoint interior in the plane $\mathbb{C}$. Given a circle packing $\mathcal{C}$, we define its **tangency map** as the map whose embedded vertex set V corresponds to the centres of circles in $\mathcal{C}$ and whose edges are given by straight lines between the centres of tangent circles. If $\mathcal{C}$ is a packing whose tangency map is isomorphic to G, we call $\mathcal{C}$ a packing of G.

Theorem 7.3.4 (Koebe-Andreev-Thurston Circle Packing Theorem [156, 221])**.** *Every finite simple planar map arises as the tangency map of a circle packing. If the map is a triangulation, the packing is unique up to Möbius transformations of the sphere.*

The **carrier** of a circle packing is the union of all the discs in the packing together with the curved triangular regions enclosed between each triplet of circles corresponding to a face (the *interstices*). Given some planar domain D, we say that a circle packing is **in** D if its carrier is D.

Theorem 7.3.5 (Rigidity for Infinite Packings, Schramm [206]). *Let G be a triangulation, circle packed in either $\mathbb{C}$ or $\mathbb{D}$. Then the packing is unique up to Möbius transformations preserving of $\mathbb{C}$ or $\mathbb{D}$ respectively.*

It is often fruitful to think of packings in $\mathbb{D}$ as being circle packings in (the Poincaré disc model of) the hyperbolic plane. The uniqueness of the packing in $\mathbb{D}$ up to Möbius transformations may then be stated as uniqueness of the packing in the hyperbolic plane up to isometries of the hyperbolic plane.

The **vertex extremal length**, defined in [121], from a vertex to infinity on an infinite graph G is defined to be

$$\mathsf{VEL}_G(v,\infty) = \sup_m \frac{\inf_{\gamma:v\to\infty} m(\gamma)^2}{\|m\|^2}, \tag{7.3.5}$$

where the supremum is over measures m on $V(G)$ such that $\|m\|^2 = \sum m(u)^2 < \infty$, and the infimum is over paths from v to ∞ in G. A connected graph is said to be **VEL parabolic** if $\mathsf{VEL}(v \to \infty) = \infty$ for some vertex v (and hence for any vertex) and **VEL hyperbolic** otherwise. The VEL type is monotone in the sense that subgraphs of VEL parabolic graphs are also VEL parabolic. A simple random walk on any VEL hyperbolic graph is transient. For graphs with bounded degrees the converse also holds: Transient graphs with bounded degrees are VEL hyperbolic [121].

Theorem 7.3.6 (He-Schramm [121, 122]). *Let G be a one-ended, infinite, simple planar triangulation. Then G may be circle packed in either the plane $\mathbb{C}$ or the unit disc $\mathbb{D}$, according to whether it is VEL parabolic or hyperbolic respectively.*

The final classical fact about circle packing we will need is the following quantitative version, due to Hansen [116], of the Ring Lemma of Rodin and Sullivan [200], which will allow us to control the radii along a random walk.

Theorem 7.3.7 (The Sharp Ring Lemma [116]). *Let u and v be two adjacent vertices in a circle packed triangulation, and $r(u), r(v)$ the radii of the corresponding circles. There exists a universal positive constant C such that*

$$\frac{r(v)}{r(u)} \leq e^{C\deg(v)}.$$

Measurability of Circle Packing

At several points throughout the paper, we will want to define mass transports in terms of circle packings. In order for these to be measurable functions of the graph, we require measurability of the circle packing. Let (G,u,v) be a doubly rooted triangulation and let $\mathcal{C}(G,u,v)$ be the unique circle packing of G in $\mathbb{D}$ or $\mathbb{C}$ such that the circle corresponding to u is centred at 0, the circle corresponding to v is centered on the positive real line and, in the parabolic case, the root circle has radius one.

Let G_k be an exhaustion of G by finite induced subgraphs with no cut-vertices and such that the complements $G \setminus G_k$ are connected. Such an exhaustion exists by the assumption that G is a

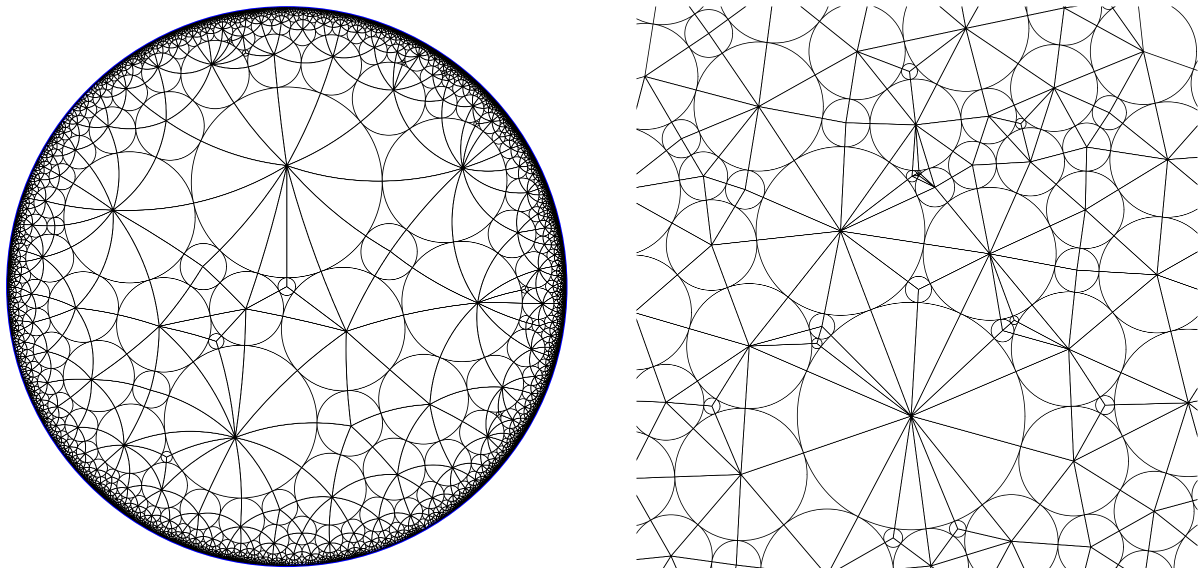

Figure 7.2: Circle packing induces an embedding of a triangulation with either hyperbolic or Euclidean geodesics, depending on CP type. By rigidity (Theorem 7.3.5), the angles between pairs of adjacent edges do not depend on the choice of packing.

one-ended triangulation. Form a finite triangulation G_k^* by adding an extra vertex ∂_k and an edge from ∂_k to each boundary vertex of G_k.

Consider first the case when G is CP hyperbolic. By applying a Möbius transformation to some circle packing of G_k^*, we find a unique circle packing $\mathcal{C}_k^*$ of G_k^* in $\mathbb{C}^\infty$ such that the circle corresponding to u is centred at the origin, the circle corresponding to v is centred on the positive real line and ∂_k corresponds to the unit circle $\partial\mathbb{D}$. In the course of the proof of the He-Schramm Theorem, it is shown that this sequence of packings converges to the unique packing of G in $\mathbb{D}$, normalised so that the circle corresponding to u is centred at 0 and the circle corresponding to v is centred on the positive real line.

As a consequence, the centres and radii of the circles of $\mathcal{C}(G, u, v)$ are limits as $r \to \infty$ of the centre and radius of a graph determined by the ball of radius r around u. In particular, they are pointwise limits of continuous functions (with respect to the local topology on graphs) and hence are measurable.

The hyperbolic radii are particularly nice to consider here. Since the circle packing in $\mathbb{D}$ is unique up to isometries of the hyperbolic plane (Möbius maps), the hyperbolic radii do not depend on the choice of packing, and we find that $r_h(v)$ is a function of (G, v).

In the CP parabolic case, the same argument works except that the packing $\mathcal{C}_k^*$ of G_k^* must be chosen to map u to the unit circle and ∂_k to a larger circle also centred at 0.

7.4 Characterisation of the CP type

Proof of Theorem 7.1.1. Since (G, ρ) is ergodic, and since the CP type does not depend on the choice of root, the CP type of G is not random. We first relate the circle packing type to the average degree. Suppose (G, ρ) is CP hyperbolic and consider a circle packing of G in the unit disc. Embed G in $\mathbb{D}$ by drawing the hyperbolic geodesics between the hyperbolic centres of the circles in its packing, so that each triangle of G is represented by a hyperbolic triangle (see Figure 7.2). It is easy to see that this is a proper embedding of G. By rigidity of the circle packing (Theorem 7.3.5), this drawing is determined by the isomorphism class of G, up to isometries of the hyperbolic plane.

Define a mass transport as follows. For each face (u, v, w) of the triangulation with angle β at u, transport β from u to each of u, v, w. If u and v are adjacent, the transport from u to v has contributions from both faces containing the edge, and the transport from u to itself has a term for each face containing u. By rigidity (Theorem 7.3.5), these angles are independent of the choice of circle packing, so that the mass sent from u to v is a measurable function of (G, u, v).

For each face f of G, let $\theta(f)$ denote the sum of the internal angles in f in the drawing. The sum of the angles of a hyperbolic triangle is π minus its area, so $\theta(f) < \pi$ for each face f. Each vertex u sends each angle 3 times, for a total mass out of exactly 6π. A vertex receives mass

$$\sum_{f\colon u\in f} \theta(f) < \pi \deg(u).$$

Applying the mass transport principle,

$$6\pi < \pi\mathbb{E}[\deg(\rho)].$$

Thus if G is CP hyperbolic then $\mathbb{E}[\deg(\rho)] > 6$.

In the CP parabolic case, we may embed G in $\mathbb{C}$ by drawing straight lines between the centres of the circles in its packing in the plane. By rigidity, this embedding is determined up to translation and scaling, and in particular all angles are determined by G. Since the sum of angles in a Euclidean triangle is π, the same transport as above applied in the CP parabolic case shows that $\mathbb{E}[\deg(\rho)] = 6$.

We now turn to amenability. Euler's formula implies that the average degree of any finite simple planar graph is at most 6. It follows that

$$\alpha((G,\rho)) = \sup\left\{\mathbb{E}\left[\frac{\sum_{v\in K_\omega(\rho)} \deg_\omega(v)}{|K_\omega(\rho)|}\right] : \omega \text{ a finitary percolation}\right\} \leq 6.$$

If G is CP hyperbolic then $\mathbb{E}[\deg(\rho)] > 6$, so that $\alpha((G, \rho)) < \mathbb{E}[\deg(\rho)]$ and (G, ρ) is invariantly non-amenable.

Conversely, suppose G is invariantly non-amenable. By Theorem 7.3.2, G almost surely admits a percolation ω which has positive Cheeger constant and bounded degrees. Such an ω is transient and since it has bounded degree it is also VEL hyperbolic. By monotonicty of the vertex extremal length, G is almost surely VEL hyperbolic as well. The He-Schramm Theorem then implies that G is almost surely CP hyperbolic. $\square$

Remark 7.4.1. In the hyperbolic case, let Area(u) be the total area of the triangles surrounding u in its drawing. Since the angle sum in a hyperbolic triangle is π minus its area, the mass transport that gives average degree greater than 6 in the hyperbolic case also gives

$$\mathbb{E}[\deg(\rho)] = 6 + \frac{1}{\pi}\mathbb{E}[\mathrm{Area}(\rho)],$$

which relates the expected degree to the density of the circle packing.

7.4.1 Completing the proof of the Benjamini-Schramm Theorem

In this section we complete our new proof of the following theorem of Benjamini and Schramm.

Theorem 7.4.2 ([50]). *Let (G,ρ) be a weak local limit of finite planar graphs G_n and suppose that G has bounded degrees almost surely. Then (G,ρ) is almost surely recurrent.*

Recall that the number of **ends of a graph** G is the supremum over finite sets K of the number of infinite connected components of $G \setminus K$. As explained in [50], it suffices to prove Theorem 7.4.2 when the graphs G_n are simple triangulations. In this case Proposition 7.1.2 implies a special case of the Benjamini-Schramm Theorem: If (G,ρ) is a simple one-ended triangulation that is a Benjamini-Schramm limit of finite planar triangulations of uniformly bounded degree, then (G,ρ) is recurrent almost surely. At the time, this was the most difficult case.

Thus, to complete the proof of Theorem 7.4.2 we need to consider the case in which the limit (G,ρ) has multiple ends. We describe below two different methods to handle this case.

Method 1. This proof considers separately three cases, depending on the number of ends of G. First, by combining Proposition 6.10 and Theorem 8.13 of [7], we have the following.

Proposition 7.4.3 ([7]). *Let (G,ρ) be an ergodic unimodular random rooted graph. Then G has one, two or infinitely many ends almost surely. If (G,ρ) has infinitely many ends almost surely, it is invariantly non-amenable.*

We rule out the case of infinitely many ends by showing that local limits of finite planar graphs are invariantly amenable. Recall the celebrated Lipton-Tarjan Planar Separator Theorem [167, Theorem 2] (which can also be proved using circle packing theory [182]).

Theorem 7.4.4 ([167]). *There exists a universal constant C such that for every m and every finite planar graph G, there exists a set $S \subset V(G)$ of size at most $Cm^{-1/2}|G|$ such that every connected component of $G \setminus S$ contains at most m vertices.*

Corollary 7.4.5. *Let (G,ρ) be the local limit of a sequence of finite planar maps G_n and suppose $\mathbb{E}[\deg(\rho)] < \infty$. Then (G,ρ) is invariantly amenable and hence has at most two ends.*

Proof. Let ω_n^m be a subset of $V(G_n)$ such that $G_n \setminus \omega_n^m$ has size at most $Cm^{-1/2}|G_n|$ and every connected component of ω_n^m has size at most m. The sequence $(G_n, \rho_n, \omega_n^m)$ is tight and therefore has a subsequence converging to (G, ρ, ω^m) for some finitary percolation ω^m on (G, ρ). Since it is a limit of percolations on finite graphs with a uniform root, the limit is unimodular.

We have that

$$\mathbb{P}(\rho \in \omega^m) \geq 1 - Cm^{-1/2} \xrightarrow[m\to\infty]{} 1.$$

Similarly, $\mathbb{P}(X_1 \in \omega^m) \to 1$. By integrability of $\deg(\rho)$, we have that

$$\mathbb{E}[\deg_{\omega^m}(\rho)] = \mathbb{E}[\mathbb{1}(\rho, X_1 \in \omega^m)\deg(\rho)] \to \mathbb{E}[\deg(\rho)].$$

Thus $\alpha((G, \rho)) = \mathbb{E}[\deg(\rho)]$ and hence (G, ρ) is invariantly amenable. □

Finally, we deal with the two-ended case.

Proposition 7.4.6. *Let (G, ρ) be a unimodular random rooted graph with exactly two ends almost surely and suppose $\mathbb{E}[\deg(\rho)] < \infty$. Then G is recurrent almost surely.*

Proof. We prove the equivalent statement for (G, ρ) reversible. We may also assume that (G, ρ) is ergodic. Say that a finite set S **disconnects** G if $G \setminus S$ has two infinite components. Since G is two-ended almost surely, such a set S exists and each infinite component of $G \setminus S$ is necessarily one-ended. We call these two components G_1 and G_2. Suppose for contradiction that G is transient almost surely. In this case, a simple random walk X_n eventually stays in one of the G_i, and hence the subgraph induced by this G_i must be transient.

Now, since G is two-ended almost surely, there exist R and M such that, with positive probability, the ball $B_R(X_n)$ disconnects G and $|B_R(X_n)| \leq M$. By the Ergodic Theorem this occurs for infinitely many n almost surely. On the event that X_n eventually stays in G_i, since G_i is one-ended, this yields an infinite collection of disjoint cutsets of size at most M separating ρ from infinity in G_i. Thus, G_i is recurrent by the Nash-Williams criterion [173], a contradiction. □

Theorem 7.4.2 now follows by combining Theorem 7.1.1, Corollary 7.4.5, and Proposition 7.4.6.

Method 2. This proof reduces Theorem 7.4.2 to Theorem 7.1.1 by taking universal covers. Given a (not necessarily planar) map M, a **cover** of M is a map $\tilde{M}$ together with a surjective graph homomorphism $\pi : \tilde{M} \to M$, such that for each vertex v, the homomorphism π maps the edges adjacent to v bijectively to the edges adjacent to $\pi(v)$ and preserves their cyclic ordering, and such that for each face f, π maps the edges adjacent to f bijectively to the edges adjacent to $\pi(f)$. The **universal cover** of M is a cover $\pi : \tilde{M} \to M$ such that $\tilde{M}$ is simply connected. If M is drawn on a surface S, the universal cover $\tilde{M}$ of M may be constructed by taking every lift of every edge of M in S to the universal cover $\tilde{S}$ of S (see e.g. [119] for the topological notions of universal cover and path lifting). Alternatively, the universal cover $\tilde{M}$ may be constructed directly as in [215]. The universal cover is unique in the sense that if $\pi' : \tilde{M}' \to M$ is also a universal cover of M then there exists an isomorphism of maps $f : \tilde{M}' \to \tilde{M}$ such that $\pi' = \pi \circ f$. Note that if a cover $\pi : \tilde{M} \to M$ is a cover of a map M and $\tilde{M}$ is recurrent, the projection $X_n = \pi(\tilde{X}_n)$ of a simple random walk $\tilde{X}_n$ on $\tilde{M}$ is a simple random walk on M, and it follows that M is also recurrent.

Let (M, ρ) be a unimodular random rooted map with universal cover $\pi : \tilde{M} \to M$. Let $\tilde{\rho}$ be chosen arbitrarily from the preimage $\pi^{-1}(\rho)$; The isomorphism class of the rooted map $(\tilde{M}, \tilde{\rho})$ does not depend on this choice. We claim that the random rooted map $(\tilde{M}, \tilde{\rho})$ is unimodular. To see this, recall that a random rooted graph is unimodular if and only if it is *involution invariant* [7, Proposition 2.2], meaning that

$$\mathbb{E} \sum_{v \in V} f(G, \rho, v) = \mathbb{E} \sum_{v \in V} f(G, v, \rho)$$

whenever f is a mass-transport such that $f(G, u, v)$ is zero unless u and v are adjacent in G. The equivalence of unimodularity and involution invariance extends immediately to random rooted maps. Given such an $f : \mathcal{M}_{\bullet\bullet} \to [0, \infty]$, let $g : \mathcal{M}_{\bullet\bullet} \to [0, \infty]$ be defined to be

$$g(M, u, v) = \sum_{e :\, e^- = u,\, e^+ = v} f(\tilde{M}, \tilde{e}^-, \tilde{e}^+),$$

where $\pi : \tilde{M} \to M$ is the universal cover of M and $\tilde{e}$ is an arbitrary element of $\pi^{-1}(e)$ for each oriented edge e of M (by uniqueness of the universal cover, the value of g does not depend on this choice). Then g is a mass transport and, letting $\tilde{V}$ denote the vertex set of $\tilde{M}$, we have

$$\sum_{v \in V} g(M, \rho, v) = \sum_{\tilde{v} \in \tilde{V}} f(\tilde{M}, \tilde{\rho}, \tilde{v}) \quad \text{and} \quad \sum_{v \in V} g(M, v, \rho) = \sum_{\tilde{v} \in \tilde{V}} f(\tilde{M}, \tilde{v}, \tilde{\rho}),$$

so that we deduce involution invariance of $(\tilde{M}, \tilde{\rho})$ from involution invariance of (M, ρ). Furthermore, if (M, ρ) is ergodic then $(\tilde{M}, \tilde{\rho})$ is also ergodic. Indeed, for every invariance event $A \subseteq \mathcal{M}_\bullet$, the event $\{(G, \rho) \in \mathcal{M}_\bullet : (\tilde{G}, \tilde{\rho}) \in \mathcal{M}_\bullet\}$ is also invariant to changing the root, and it follows that if (M, ρ) is ergodic then $(\tilde{M}, \tilde{\rho})$ is also ergodic.

Alternative proof of Theorem 7.4.2. Let (G, ρ) be a simple, bounded degree, ergodic unimodular random rooted triangulation with $\mathbb{E}[\deg(\rho)] = 6$. The universal cover $(\tilde{G}, \tilde{\rho})$ of (G, ρ) has all these properties and is also one-ended, so that $\tilde{G}$ is CP parabolic almost surely by Theorem 7.1.1. By the He-Schramm Theorem [121], $\tilde{G}$ is recurrent almost surely, and so G is also recurrent almost surely, completing the proof of the Benjamini-Schramm Theorem. □

7.5 Boundary theory

Recall that given a G and a vertex v we write P_v^G and E_v^G to denote the probability and expectation with respect to random walk $(X_n)_{n \geq 0}$ on G started from v.

7.5.1 Convergence to the boundary

Let (G, ρ) be a one-ended, simple, CP hyperbolic reversible random triangulation. Recall that for a CP hyperbolic G with circle packing $\mathcal{C}$ in $\mathbb{D}$, we write $r(v)$ and $z(v)$ for the Euclidean radius and centre of the circle corresponding to the vertex v in $\mathcal{C}$ and $z_h(v), r_h(v)$ for the hyperbolic centre and radius.

Our first goal is to show that the Euclidean radii $r(X_n)$ decay exponentially along a random walk (X_n). We initially prove only a bound, and will prove the existence of the limit rate of decay stated in Theorem 7.1.4 only after we have proven the exit measure is non-atomic.

Lemma 7.5.1. *Let (G, ρ) be a CP hyperbolic reversible random rooted triangulation with $\mathbb{E}[\deg(\rho)] < \infty$ and let $\mathcal{C}$ be a circle packing of G in the unit disc. Let $(X_n)_{n\geq 0}$ be a simple random walk on G started from ρ. Then almost surely*

$$\limsup_{n\to\infty} \frac{\log r(X_n)}{n} < 0.$$

Proof. We may assume that (G, ρ) is ergodic, else we may take an ergodic decomposition. By Theorem 7.1.1 (G, ρ) is invariantly non-amenable. By Theorem 7.3.2, there is an ergodic percolation ω on G such that $\deg(v)$ is bounded by some M for all $v \in \omega$ and $\mathbf{i}_E(\omega) > 0$ almost surely.

Recall the notion of an **induced random walk** on ω: let N_m be the mth time X is in ω (that is, $N_0 = \inf\{n \geq 0 : X_n \in \omega\}$ and inductively $N_{m+1} = \inf\{n > N_m : X_n \in \omega\}$). The **induced network** $\bar{\omega}$ is defined to be the weighted graph on the vertices of ω with edge weights given by

$$\bar{w}(u, v) = \deg(u) P_u^G(X_{N_1} = v)$$

so that X_{N_m} is the random walk on the weighted graph $\bar{\omega}$. Note that $\bar{\omega}$ may have non-zero weights between vertices which are not adjacent in G, so that $\bar{\omega}$ is no longer a percolation on G, and may not even be planar.

We first claim that $\bar{\omega}$ with the weights $\bar{w}$ of the induced random walk also has positive Cheeger constant. Indeed, the weight of a vertex $v \in \bar{\omega}$ is just its degree in G and so is between 1 and M for any vertex. The edge boundary in $\bar{\omega}$ of a set K is at least the number of edges connecting K to $V \setminus K$ in ω. Thus $\mathbf{i}_E(\bar{\omega}) \geq \mathbf{i}_E(\omega)/M > 0$. It follows that the induced random walk on ω has spectral radius less than one [173].

Now, as in (7.3.1), Cauchy-Schwarz gives that, for some $c > 0$,

$$P_\rho^G(X_{N_m} = v) \leq M^{1/2} \exp(-cm)$$

for every vertex v almost surely. Since the total area of all circles in the packing is at most π, with

c as above, there exists at most $e^{cm/2}$ circles of radius greater than $e^{-cm/4}$ for each m. Hence

$$\begin{aligned} P_\rho^G\left(r(X_{N_m}) \geq e^{-cm/4}\right) &= \sum_{v:r(v)\geq e^{-cm/4}} P_\rho^G\left(X_{N_m} = v\right) \\ &\leq \left|\left\{v : r(v) \geq e^{-cm/4}\right\}\right| \cdot M^{1/2} e^{-cm} \\ &\leq M^{1/2} e^{-cm/2}. \end{aligned}$$

These probabilities are summable, and so Borel-Cantelli implies that almost surely for large enough m,

$$r(X_{N_m}) \leq e^{-cm/4}.$$

That is, we have exponential decay of the radii for the induced walk:

$$\limsup_{m\to\infty} \frac{\log r(X_{N_m})}{m} \leq -\frac{c}{4}. \tag{7.5.1}$$

It remains to prove that the exponential decay is maintained between visits to ω. By stationarity and ergodicity of (G, ρ, ω), the density of visits to ω is $\mathbb{P}(\rho \in \omega) \neq 0$. That is,

$$\lim_{m\to\infty} \frac{N_m}{m} = \mathbb{P}(\rho \in \omega)^{-1}$$

almost surely. In particular,

$$\limsup_{m\to\infty} \frac{\log r(X_{N_m})}{N_m} < 0. \tag{7.5.2}$$

Given n, let m be the number of visits to ω up to time n, so that $N_m \leq n < N_{m+1}$. Since m/N_m converges, $n/N_m \to 1$. By the Sharp Ring Lemma (Theorem 7.3.7),

$$\frac{r(X_n)}{r(X_{N_m})} \leq \exp\left(C \sum_{i=N_m}^{n} \deg(X_i)\right)$$

so that

$$\frac{\log r(X_n)}{n} \leq \frac{\log r(X_{N_m})}{n} + \frac{C}{n} \sum_{i=N_m}^{n} \deg(X_i). \tag{7.5.3}$$

Again, the Ergodic Theorem gives us the almost sure limit

$$\lim \frac{1}{n} \sum_{i=0}^{n} \deg(X_i) = \mathbb{E}[\deg(\rho)].$$

Thus almost surely

$$\lim \frac{1}{n} \sum_{i=N_m}^{n} \deg(X_i) = \lim \frac{1}{n} \sum_{i=0}^{n} \deg(X_i) - \lim \frac{N_m}{n} \lim \frac{1}{N_m} \sum_{i=0}^{N_m} \deg(X_i)$$
$$= \mathbb{E}[\deg(\rho)] - \mathbb{E}[\deg(\rho)] = 0.$$

Combined with (7.5.2) and (7.5.3) we get

$$\limsup_{n\to\infty} \frac{\log r(X_n)}{n} = \limsup_{n\to\infty} \frac{\log r(X_{N_m})}{n}$$
$$= \lim_{n\to\infty} \frac{N_m}{n} \limsup_{n\to\infty} \frac{\log r(X_{N_m})}{N_m} < 0. \qquad \square$$

Proof of Theorem 7.1.3, item 1. We prove the equivalent statement for (G, ρ) reversible with $\mathbb{E}[\deg(\rho)] < \infty$, putting us in the setting of Lemma 7.5.1. The path formed by drawing straight lines between the Euclidean centres of the circles along the random walk path has length $r(\rho) + 2\sum_{i\geq 1} r(X_i)$, which is almost surely finite by Lemma 7.5.1. It follows that the sequence of Euclidean centres is Cauchy almost surely and hence converges to some point, necessarily in the boundary. Because the radii of the circles $r(X_n)$ converge to zero almost surely, the hyperbolic centres must also converge to the same point. $\square$

7.5.2 Full support and non-atomicity of the exit measure

We now prove item 2 of Theorem 7.1.3, which states that the exit measure on the unit circle has full support and no atoms almost surely. We start with a general observation regarding atoms in boundaries of stationary graphs.

We say that two metrics d_1 and d_2 on the vertex set V of a graph G are **compatible** if the identity map from V to itself extends to an isomorphism between the completions of the metric spaces (V, d_1) and (V, d_2) (or, equivalently, if the same sequences are Cauchy for d_1 and d_2). For example, the Euclidean distances between centres of circles corresponding to vertices in different circle packings of a one-ended planar triangulation in either the full plane or the unit disc are compatible by Theorem 7.3.5. We define a **compatible family of metrics** to be a Borel function $d = d_u^G(v, w)$ from the space of triply rooted graphs (i.e., the set of isomoprhism classes of graphs with an ordered triple of distinguished vertices, equipped with an appropriate variant of the local topology) to the positive reals such that for every locally finite, connected graph $G = (V, E)$,

1. $d_u^G(\cdot, \cdot)$ is a metric on V for every vertex u of G, and
2. the metrics d_u^G and d_v^G are compatible for each two vertices u and v of G.

Given a compatible family of metrics d and a rooted graph (G, ρ), the completion $\bar{V}$ of V with respect to d_ρ^G has a topology that does not depend on the choice of root vertex ρ. Such a completion is called an **invariant completion**. Compatible families of metrics for maps are defined similarly.

In some cases of interest, a compatible family of metrics d might only be defined for graphs or maps in some rerooting-invariant class (e.g. circle packings are defined for the class of one-ended simple planar triangulations). In this case, we may extend d arbitrarily to all graphs or maps by setting it to be the discrete metric where it is not defined.

Lemma 7.5.2. *Let d be a compatible family of metrics, let (G,ρ) be a stationary random rooted graph or map, and let $\bar{V}$ be the completion of V with respect to d_ρ^G. Suppose that the random walk on G converges almost surely to a point in the boundary $\partial V = \bar{V} \setminus V$. Then the exit measure on ∂V is either trivial (concentrated on a single point) or non-atomic almost surely.*

For each CP hyperbolic one-ended simple planar triangulation G, take a circle packing of G in $\mathbb{D}$, normalized so that the circle corresponding to u is centred at the origin, and let $d = d_u^G(v,w)$ be the Euclidean distance between the Euclidean centres of the circles corresponding to v and w. By circle packing rigidity (Theorem 7.3.5), this circle packing is unique up to rotations, so that the metric d_u^G is well defined. The metrics d_u^G and d_v^G are also compatible for every pair of vertices u and v in G. Thus, after an arbitrary extension to other maps, d is an compatible family of metrics.

Another natural example, defined for all graphs, is the Martin compactification. As a consequence of this lemma, for any stationary random graph, the exit measure on the Martin boundary (which can be defined as the completion of V with respect to a compatible family of metrics, see e.g. [229]) is almost surely either non-atomic or trivial. In particular, this gives an alternative proof of a recent result of Benjamini, Paquette and Pfeffer [37], which states that for every stationary random graph, the space of bounded harmonic functions on the graph is either one dimensional or infinite dimensional almost surely. A straightforward extension of this lemma applies to random families of metrics.

Proof. Condition on (G,ρ). For each atom ξ of the exit measure, define the harmonic function $h_\xi(v) = P_v^G(\lim X_n = \xi)$. By Lévy's 0-1 law,

$$h_\xi(X_n) \xrightarrow{a.s.} \mathbb{1}(\lim X_n = \xi)$$

for each atom ξ and

$$P_{X_n}^G(\lim X_n \text{ is an atom}) = \sum_\xi h_\xi(X_n) \xrightarrow{a.s.} \mathbb{1}(\lim X_n \text{ is an atom}).$$

Define $M(G,v) = \max_\xi h_\xi(v)$ to be the maximal atom size. Since the topology of $\bar{V}$ does not depend on the choice of root, the sequence $M(G,X_n)$ is stationary. Combining the two above limits, we find the almost sure limit

$$M(G,X_n) \xrightarrow{a.s.} \mathbb{1}(\lim X_n \text{ is an atom}).$$

Since $M(G,X_n)$ is a stationary sequence with limit in $\{0,1\}$, it follows that $M(G,\rho) \in \{0,1\}$ almost surely. That is, either there are no atoms in the exit measure or there is a single atom with weight

1 almost surely. □

Proof of Theorem 7.1.3, item 2. We may assume that (G, ρ) is ergodic. Applying Lemma 7.5.2 to the $\deg(\rho)$-biasing of (G, ρ), we deduce that the exit measure has at most a single atom. Next, we rule out having a single atom. Suppose for contradiction that there is a single atom $\xi = \xi(\mathcal{C})$ almost surely for some (and hence every) circle packing $\mathcal{C}$ of G in $\mathbb{D}$. Applying the Möbius transformation

$$\Phi(z) = -i\frac{z+\xi}{z-\xi},$$

which maps $\mathbb{D}$ to the upper half-plane $\mathbb{H} = \{\Im(z) > 0\}$ and ξ to ∞, gives a circle packing of G in $\mathbb{H}$ such that the random walk tends to ∞ almost surely. Since circle packings in $\mathbb{H}$ are unique up to Möbius transformations and the boundary point ∞ is determined by the graph G, such a circle packing in $\mathbb{H}$ is unique up to Möbius transformations of the upper half-plane that fix ∞, namely $az + b$ with real $a \geq 0$ and b (translations and dilations).

Inverting around the atom has therefore given us a way of canonically endowing G with Euclidean geometry: if we draw G in $\mathbb{H}$ using straight lines between the Euclidean centres of the circles in the half-plane packing, the angles at the corners around each vertex u are independent of the original choice of packing $\mathcal{C}$. Transporting each angle from u to each of the three vertices forming the corresponding face f as in the proof of Theorem 7.1.1 implies that $\mathbb{E}[\deg(\rho)] = 6$. This contradicts Theorem 7.1.1 and the assumption that (G, ρ) is CP hyperbolic almost surely. This completes the proof that the exit measure is non-atomic.

To finish, we show that the exit measure has support $\partial\mathbb{D}$. Suppose not. We will define a mass transport on G in which each vertex sends a mass of at most one but some vertices receive infinite mass, contradicting the mass transport principle.

Consider the complement of the support of the exit measure, which is a union of disjoint open intervals $\bigcup_{i\in I}(\theta_i, \psi_i)$ in $\partial\mathbb{D}$. Since the exit measure is non-atomic, $\theta_i \neq \psi_i \mod 2\pi$ for all i.

For each such interval (θ_i, ψ_i), let γ_i be the hyperbolic geodesic from $e^{i\theta_i}$ to $e^{i\psi_i}$. That is, γ_i is the intersection with $\mathbb{D}$ of the circle passing through both $e^{i\theta_i}$ and $e^{i\psi_i}$ that intersects $\partial\mathbb{D}$ at right angles. Let A_i be the set of vertices such that the circle corresponding to v is contained in the region to the right of γ_i, i.e. bounded between γ_i and the boundary interval (θ_i, ψ_i) (see Figure 7.3(a)).

Each vertex is contained in at most one such A_i. For each vertex u in A_i, consider the hyperbolic geodesic ray γ_u from the hyperbolic centre $z_h(u)$ to $e^{i\theta_i}$. Define a mass transport by sending mass one from $u \in A_i$ to the vertex v corresponding to the first circle intersected by both γ_u and γ_i. There may be no such circle, in which case no mass is sent from u. Since the transport is defined in terms of the hyperbolic geometry and the support of the exit measure, it is a function of the isomorphism class of (G, u, v) by Theorem 7.3.5.

Let $\phi \in (\theta_i, \psi_i)$ and consider the set of vertices whose corresponding circles intersect both γ_i and the geodesic γ_ϕ from $e^{i\phi}$ to $e^{i\theta_i}$. As ϕ increases from θ_i to ψ_i, this set is increasing. It follows that for each fixed v for which the circle corresponding to v intersects γ_i, the set B_v of $\phi \in (\theta_i, \psi_i)$

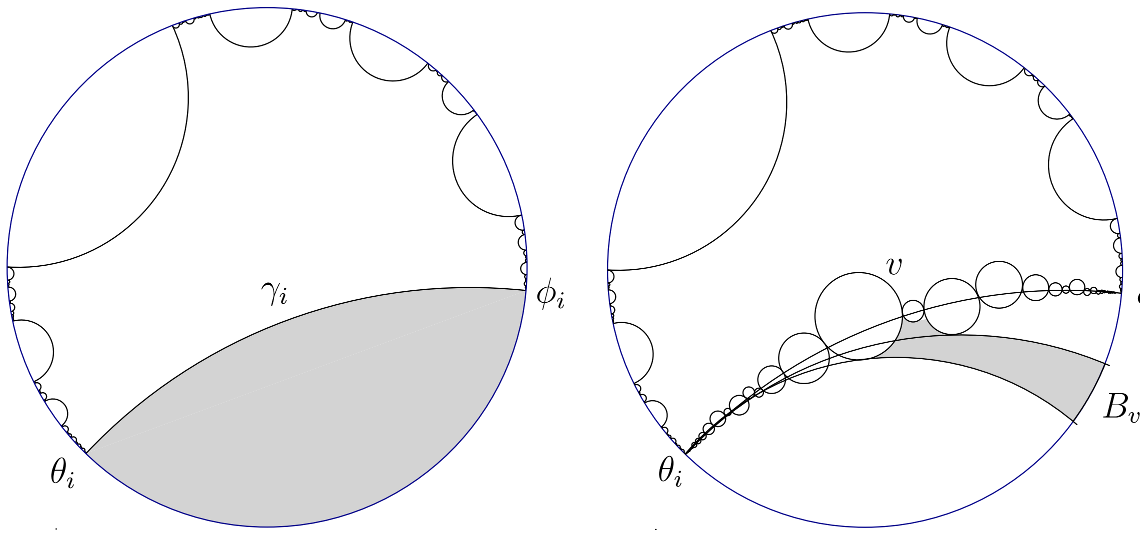

(a) A geodesic γ_i is drawn over each component of the complement of the support. Circles contained in the shaded area are in A_i.

(b) The vertex v receives mass from circles with hyperbolic centres in the shaded area.

Figure 7.3: An illustration of the mass transport used to show the exit measure has full support.

for which the circle corresponding to v is the first circle intersected by γ_ϕ that also intersects γ_i is an interval (see Figure 7.3(b)).

Since there are only countably many vertices, B_v must have positive length for some v. Thus there is an open neighbourhood of the boundary in which all the circles send mass to this vertex. This vertex therefore receives infinite mass, contradicting the mass transport principle. $\square$

7.5.3 The unit circle is the Poisson boundary

The Poisson-Furstenberg boundary [91–93] (or simply the Poisson boundary) of a graph (or more generally, of a Markov chain) is a formal way to encode the asymptotic behaviour of random walks on G. We refer the reader to [142, 173, 195] for more detailed introductions.

Recall that a function $h : V(G) \to \mathbb{R}$ is said to be **harmonic** if

$$h(v) = \frac{1}{\deg(v)} \sum_{u \sim v} h(u)$$

for all $v \in V(G)$ — or, equivalently, if $h(X_n)$ is a martingale. Let $\bar{G} = G \cup \partial G$ be a compactification of G so that the random walk X_n converges almost surely. For each $v \in V(G)$ we let P_v^G denote the law of the limit of the random walk started at v. Every bounded Borel function g on ∂G extends to a harmonic function

$$h(v) := E_v^G\big[g(\lim X_n)\big]$$

on G. Such a compactification is called a **realisation of the Poisson boundary** of G if *every* bounded harmonic funtion h on G may be represented as an extension of a boundary function in this way.

Harmonic functions can be used to encode asymptotic behaviour of the random walk as follows. Let $G^{\mathbb{N}}$ be the space of sequences in G. The shift operator on $G^{\mathbb{N}}$ is defined by $\theta(x_0, x_1, \dots) = (x_1, x_2, \dots)$, and we write $\mathcal{I}$ for the σ-algebra of shift-invariant events $A = \theta A$. Be careful to note the distinction between invariant events for the random walk on G, just defined, and invariant events for the sequence $(G, (X_{n+k})_{n\in\mathbb{Z}})_{k\in\mathbb{Z}}$ as defined in Section 7.3.2.

There is an isomorphism between the space of bounded harmonic functions on G and $L^\infty(G^{\mathbb{N}}, \mathcal{I})$ given by

$$h \mapsto g(x_1, x_2, \dots) = \lim_{n\to\infty} h(x_n), \qquad\qquad g \mapsto h(v) = E_v^G\big[g(v, X_1, X_2, \dots)\big].$$

The limit here exists P_v^G-almost surely by the bounded Martingale Convergence Theorem, while the fact that these two mappings are inverses of one another is a consequence of Lévy's 0-1 Law: If $h(v) = E_v^G[g(X)]$ is the harmonic extension of some invariant function g, then

$$h(X_n) \xrightarrow{a.s.} g(\rho, X_1, X_2, \dots). \tag{7.5.4}$$

As a consequence of this isomorphism, and since the span of simple functions is dense in L^∞, the topological boundary ∂G is a realisation of the Poisson boundary of G if and only if for every invariant event A there exists a Borel set $B \subset \partial G$ such that the symmetric difference $A \Delta \{\lim X_n \in B\}$ is P_v^G-null.

For example, the Poisson boundary of a tree may be realised as its space of ends, and the one-point compactification of a transient graph G gives rise to a realisation of the Poisson boundary if and only if G is Liouville (i.e. the only bounded harmonic functions on G are constant). The Poisson boundary of any graph may be realised as the graph's Martin boundary [229], but this is not always the most natural construction.

Our main tools for controlling harmonic functions will be Lévy's 0-1 Law and the following consequence of the Optional Stopping Theorem. For a set $W \subset V$ of vertices, let T_W be the first time the random walk visits W. If h is a positive, bounded harmonic function, the Optional Stopping Theorem implies

$$h(v) \geq E_v^G[h(X_{T_W})\mathbb{1}_{T_W<\infty}] \geq P_v^G(\text{Hit } W)\inf\{h(u) : u \in W\}. \tag{7.5.5}$$

Lemma 7.5.3. *Let (G, ρ) be a CP hyperbolic reversible random rooted triangulation with $\mathbb{E}[\deg(\rho)] < \infty$, and let $(X_n)_{n\in\mathbb{Z}}$ be the reversible bi-infinite random walk. Then almost surely*

$$P_{X_n}^G\left(\mathit{hit}\ \{X_{-1}, X_{-2}, \dots\}\right) \to 0.$$

Proof. Let $\mathcal{C}$ be a circle packing of G in $\mathbb{D}$. Recall from Theorem 7.1.3, item 1 that for a random walk X_n, almost surely $\Xi := \lim z(X_n)$ exists, and its law is non-atomic and of full support on $\partial\mathbb{D}$. Since the exit measure is non-atomic, the limit points $\Xi_+ := \lim z(X_n)$ and $\Xi_- := \lim z(X_{-n})$ are almost surely distinct.

Let $\{U_i\}_{i\in I}$ be a countable basis for the topology of $\partial\mathbb{D}$ (say, intervals with rational endpoints) and for each i let h_i be the harmonic function

$$h_i(v) = P_v^G\left(\Xi \in U_i\right).$$

By Lévy's 0-1 law, $h_i(X_n) \to \mathbb{1}(\Xi_+ \in U_i)$ for every i almost surely. Thus there exists some i_0 with $\Xi_- \in U_{i_0}$ and $\Xi_+ \notin U_{i_0}$. In particular there is almost surely some bounded harmonic function $h = h_{i_0} \geq 0$ with $h(X_n) \xrightarrow[n\to\infty]{} 0$ and

$$a := \inf\{h(X_{-m}) : m > 0\} > 0.$$

By (7.5.5)

$$h(X_n) \geq a \cdot P_{X_n}^G\left(\text{hit } \{X_{-1}, X_{-2}, \dots\}\right).$$

Since $h(X_n) \to 0$, we almost surely have

$$P_{X_n}^G\left(\text{Hit } \{X_{-1}, X_{-2}, \dots\}\right) \to 0. \qquad \square$$

Proof of Theorem 7.1.3, item 3. We prove the equivalent statement for (G,ρ) reversible with $\mathbb{E}[\deg(\rho)] < \infty$, and may assume that (G,ρ) is ergodic.

We need to prove that for every invariant event A for the simple random walk on G with $P_\rho^G(A) > 0$, there is a Borel set $B \subset \partial\mathbb{D}$ such that

$$P_\rho^G\left(A \,\Delta\, \{\Xi_+ \in B\}\right) = 0,$$

where $\Xi_+ = \lim z(X_n)$. Let h be the harmonic function $h(v) = P_v^G(A)$, and let B be the set of $\xi \in \partial\mathbb{D}$ such that there exists a path $(\rho, v_1, v_2, \dots)$ in G such that for some $c > 0$,

$$h(v_i) \to 1, \quad z(v_i) \to \xi, \quad \text{and} \quad |\xi - z(v_i)| < 2e^{-ci},$$

where $|\cdot|$ denotes Euclidean length. The condition on exponential decay of $|\xi - z(v_i)|$ can be omitted by invoking the theory of universally measurable sets. We are spared from this by Lemma 7.5.1. With an explicit rate of convergence, it is straightforward to see that B is Borel: Let $B_{c,m,\varepsilon,n}$ be the open set of $\xi \in \partial\mathbb{D}$ such that there exists a path $\rho, v_1, \dots, v_n$ in G such that $h(v_i) > 1 - \varepsilon$ for every $i \geq m$, and with $|\xi - z(v_i)| < 2e^{-ci}$. Then $B = \bigcup_c \bigcap_\varepsilon \bigcup_m \bigcap_n B_{c,m,\varepsilon,n}$, where m, n are integers and c, ε are positive rationals, and it follows that B is Borel.

If the random walk has $(\rho, X_1, \dots) \in A$ then, by Lévy's 0-1 law and Lemma 7.5.1, the limit point Ξ_+ is in B almost surely. In particular, if $P_\rho^G(A) > 0$, then the exit measure of B is positive. It remains to show that $(\rho, X_1, \dots) \in A$ almost surely on the event that $\Xi_+ \in B$.

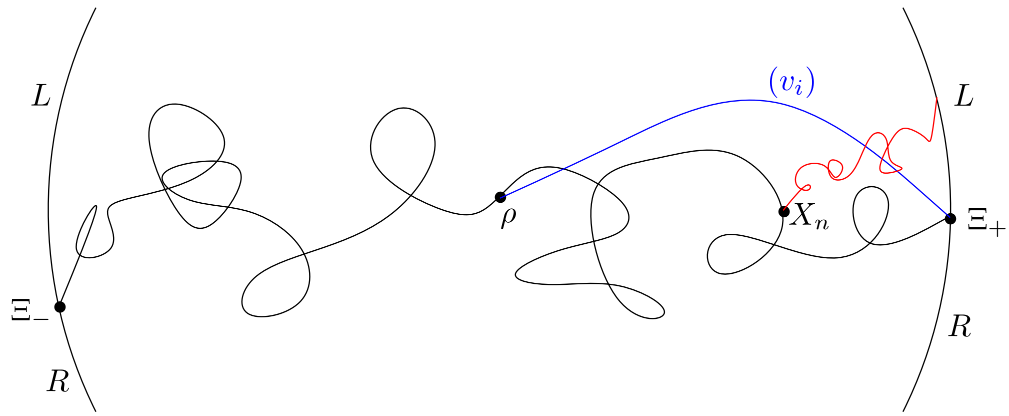

Figure 7.4: For infinitely many n, a new random walk (red) started from X_n has probability at least $1/3$ of hitting each of L and R, and probability at least $1/4$ of hitting the path (v_i) (blue).

Consider the two intervals L and R separating the almost surely distinct limit points Ξ_+ and Ξ_-. Let p_L^n and p_R^n be the probabilities that a *new, independent* random walk started from X_n hits the boundary in the interval L or R respectively. Since the exit measure is non-atomic almost surely, the event

$$E_n = \{\min(p_L^n, p_R^n) > 1/3\}$$

has positive probability (in fact, it is not hard to see that each of the random variables p_L^n is uniformly distributed on $[0,1]$ so that E_n has probability $1/3$). Moreover, the value of p_L^n does not depend on the choice of circle packing and is therefore a function of $(G, (X_{n+k})_{k\in\mathbb{Z}})$. By the stationarity and ergodicity of $(G, (X_n)_{n\in\mathbb{Z}})$, the events E_n happen infinitely often almost surely (see Figure 7.4).

Now condition on $\Xi_+ \in B$. Since the exit measure of B is positive, the events E_n still happen infinitely often almost surely after conditioning. Let $(v_i)_{i\geq 0}$ be a path from ρ in G such that $z(v_i) \to \Xi_+ \in B$ and $h(v_i) \to 1$. In particular,

$$\inf\{h(v_i) : i \geq 1\} > 0.$$

The path $(\dots, X_{-2}, X_{-1}, \rho, v_1, v_2, \dots)$ disconnects X_n from at least one of the intervals L or R and so

$$P^G_{X_n}(\text{hit } \{\dots, X_{-1}, \rho, v_1, \dots\}) \geq \min(p_L^n, p_R^n) \tag{7.5.6}$$

which is greater than $1/3$ infinitely often almost surely. We stress that the expression refers to the probability that an independent random walk started from X_n hits the path $(\dots, X_{-1}, \rho, v_1, \dots)$, and that this bound holds trivially if X_n is on the path (v_m). By Lemma 7.5.3,

$$P^G_{X_n}(\text{hit } \{\dots, X_{-1}, \rho\}) \xrightarrow{a.s.} 0, \tag{7.5.7}$$

and hence $P^G_{X_n}(\text{hit } \{v_1, v_2, \dots\}) > 1/4$ infinitely often almost surely (see Figure 7.4). Note that, since the choice of v_i could depend on the whole trajectory of X, we have *not* shown that X hits the path (v_i) infinitely often. Nevertheless, by (7.5.5), almost surely infinitely often

$$h(X_n) > \frac{1}{4}\inf\{h(v_i) : i \geq 1\} > 0. \tag{7.5.8}$$

By Lévy's 0-1 law, $\lim_{n\to\infty} h(X_n) = 1$ almost surely as desired. □

7.6 Hyperbolic speed and decay of radii

We now use the fact that the exit measure is almost surely non-atomic to strengthen Lemma 7.5.1 and deduce that the limit rate of decay of the Euclidean radii along the random walk exists. The key idea is to use a circle packing in the upper half-plane normalised by the limits of two independent random walks.

Fix some circle packing $\mathcal{C}$ in $\mathbb{D}$, so that, by Theorem 7.1.3, the limit points $\Xi_\pm = \lim_{n\to\pm\infty} z(X_n)$ exist and are distinct almost surely. Let Φ_X be a Möbius transformation that maps $\mathbb{D}$ to the upper half-plane $\mathbb{H}$ and sends Ξ_+ to 0 and Ξ_- to ∞. We consider the upper half-plane packing $\widehat{\mathcal{C}} = \Phi_X(\mathcal{C})$.

Similarly to the proof of non-atomicity in Section 7.5.2, we now have two boundary points 0 and ∞ fixed by the graph G and the path (X_n), so that the resulting circle packing is unique up to scaling. Now, however, the packing depends on both G and the random walk, so that this new situation is not paradoxical (as it was in Section 7.5.2 where we ruled out the possibility that the exit measure has a single atom).

Proof of Theorem 7.1.4. We prove the equivalent statement for (G,ρ) reversible with $\mathbb{E}[\deg(\rho)] < \infty$, and may assume that (G,ρ) is ergodic. We fix a circle packing $\widehat{\mathcal{C}} = \Phi_X(\mathcal{C})$ in $\mathbb{H}$ as above, with the doubly infinite random walk from ∞ to 0. Let $\hat{r}(v)$ be the Euclidean radius of the circle corresponding to v in $\widehat{\mathcal{C}}$. The ratio of radii $\hat{r}(X_n)/\hat{r}(X_{n-1})$ does not depend on the choice of $\widehat{\mathcal{C}}$, so these ratios form a stationary ergodic sequence. By the Sharp Ring Lemma, $\mathbb{E}\big[|\log(\hat{r}(X_1)/\hat{r}(\rho))|\big] \leq C\mathbb{E}\big[\deg(\rho)\big] < \infty$, so that the Ergodic Theorem implies that

$$-\frac{1}{n}\log\frac{\hat{r}(X_n)}{\hat{r}(\rho)} = -\frac{1}{n}\sum_{1}^{n}\log\frac{\hat{r}(X_i)}{\hat{r}(X_{i-1})} \xrightarrow[n\to\infty]{a.s.} -\mathbb{E}\left[\log\frac{\hat{r}(X_1)}{\hat{r}(\rho)}\right]. \tag{7.6.1}$$

Now, since $\widehat{\mathcal{C}}$ is the image of $\mathcal{C}$ through the Möbius map Φ_X, and since Φ_X is conformal at Ξ_+,

$$\frac{\hat{r}(X_n)}{r(X_n)} \to \left|\Phi'_X(\Xi_+)\right| > 0. \tag{7.6.2}$$

Therefore

$$\lim\frac{-\log r(X_n)}{n} = \mathbb{E}\left[-\log\frac{\hat{r}(X_1)}{\hat{r}(\rho)}\right] \tag{7.6.3}$$

and by Lemma 7.5.1 this limit must be positive. This establishes the rate of decay of the radii.

Next, we relate this to the distance of $z(X_n)$ from $\partial\mathbb{D}$. By the triangle inequality, $1 - |z(X_n)|$

is at most the length of the path formed by drawing straight lines between the Euclidean centres of the circles along the random walk path starting at X_n:

$$1 - |z_h(X_n)| \leq 1 - |z(X_n)| \leq \sum_{i \geq n} 2r(X_i).$$

Since the radii decay exponentially, taking the limits of the logarithms,

$$\liminf \frac{-\log\big(1 - |z_h(X_n)|\big)}{n} \geq \lim \frac{-\log r(X_n)}{n}. \tag{7.6.4}$$

To get a corresponding upper bound, note that, since every circle neighbouring X_n is contained in the open unit disc, $1 - |z_h(X_n)|$ is at least the radius of the smallest neighbour of X_n. Applying the Sharp Ring Lemma, we have

$$1 - |z_h(X_n)| \geq r(X_n) \exp(-C \deg(X_n)).$$

Taking logarithms and passing to the limit,

$$\begin{aligned} \limsup \frac{-\log\big(1 - |z_h(X_n)|\big)}{n} &\leq \lim \frac{-\log r(X_n)}{n} + \lim \frac{C \deg(X_n)}{n} \\ &= \lim \frac{-\log r(X_n)}{n}, \end{aligned} \tag{7.6.5}$$

where the almost sure limit $\deg(X_n)/n \to 0$ follows from $\mathbb{E}[\deg(\rho)] < \infty$ and Borel-Cantelli. Combining (7.6.4) and (7.6.5) gives the almost sure limit

$$\lim \frac{-\log\big(1 - |z_h(X_n)|\big)}{n} = \lim \frac{-\log r(X_n)}{n}.$$

Finally, to relate this to the speed in the hyperbolic metric, recall that distances from the origin in the hyperbolic metric are given by

$$d_{hyp}(0, z) = 2 \tanh^{-1} |z|$$

and hence

$$\begin{aligned} \lim \frac{1}{n} d_{hyp}(z_h(\rho), z_h(X_n)) &= \lim \frac{1}{n} d_{hyp}(0, z_h(X_n)) \\ &= \lim \frac{2}{n} \tanh^{-1} |z_h(X_n)| \\ &= \lim -\frac{1}{n} \log(1 - |z_h(X_n)|). \end{aligned} \qquad \square$$

7.7 Extensions

We now discuss two basic extensions of our main results beyond simple triangulations. These are to weighted and to non-simple triangulations. The latter are of particular interest since the PSHT is not simple. Some of our results hold for much more general planar maps, which are treated in [20].

Weighted networks. Suppose (G, ρ, w) is a unimodular random rooted weighted triangulation. As in the unweighted case, if $\mathbb{E}[w(\rho)]$ is finite then biasing by $w(\rho)$ gives an equivalent random rooted weighted triangulation which is reversible for the weighted simple random walk [7, Theorem 4.1]. Our arguments generalise with no change to recover all our main results in the weighted setting provided the following conditions are satisfied.

1. $\mathbb{E}[w(\rho)] < \infty$. This allows us to bias to get a reversible random rooted weighted triangulation.
2. $\mathbb{E}[w(\rho)\deg(\rho)] < \infty$. After biasing by $w(\rho)$, the expected degree is finite, allowing us to apply the Ring Lemma together with the Ergodic Theorem as in the proofs of Lemma 7.5.1 and Theorem 7.1.4.
3. A version of Theorem 7.3.2 holds. That is, there exists a percolation ω such that the induced network $\bar{\omega}$ has positive Cheeger constant almost surely. Two natural situations in which this occurs are
   (a) when all the weights are non-zero almost surely. In this situation, we may adapt the proof of Theorem 7.3.2 by first deleting all edges of weight less than $1/M$ and all vertices of total weight greater than M before continuing the construction as before.
   (b) when the subgraph formed by the edges of non-zero weight is connected and is itself invariantly non-amenable. This occurs when we circle pack planar maps that are not triangulations by adding edges of weight 0 in non-triangular faces to triangulate them.

Non-simple triangulations. Suppose G is a one-ended planar map. The endpoints of any double edge or loop in G disconnect G into connected components exactly one of which is infinite. The **simple core** of G, denoted $\mathrm{core}(G)$, is defined by deleting the finite component contained within each double edge or loop of G before gluing the double edges together or deleting the loop as appropriate. See Figure 7.5 for an example, and [24] for a more detailed description. When G is a triangulation, so is its core. The core can be seen as a subgraph of G, with some vertices removed, and multiple edges replaced by a single edge. The induced random walk on the core, is therefore a random walk on a weighted simple triangulation.

In general, it is possible that all of G is deleted by this procedure, but in this case there are infinitely many disjoint vertex cut-sets of size 2 separating each vertex from infinity, implying that G is VEL parabolic and hence invariantly amenable. When G is invariantly non-amenable, the conclusions of Theorem 7.1.3 hold with the necessary modifications.

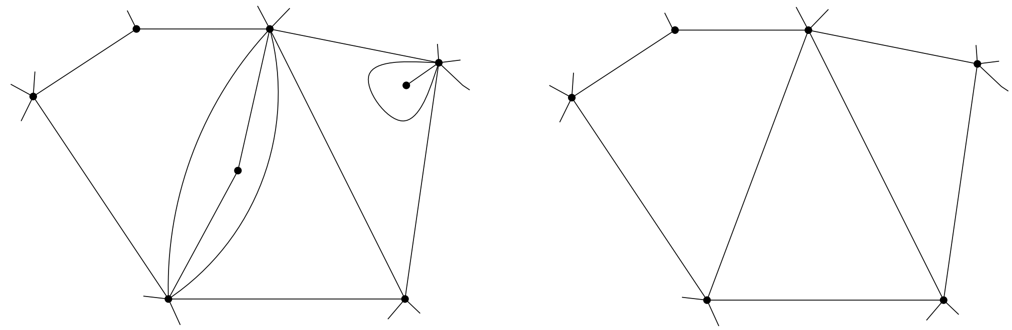

Figure 7.5: Extracting the core of a non-simple map. Left: part of a map. Right: corresponding part of its core.

Theorem 7.7.1. *Let (G, ρ) be an invariantly non-amenable, one-ended, unimodular random rooted planar triangulation with $\mathbb{E}[\deg^2(\rho)] < \infty$. Then $\mathrm{core}(G)$ is CP hyperbolic. Let $\mathcal{C}$ be a circle packing of $\mathrm{core}(G)$ in $\mathbb{D}$, and let $(Y_n)_{n\in\mathbb{N}}$ be the induced random walk on $\mathrm{core}(G)$. The following hold conditional on (G, ρ) almost surely:*

1. *$z(Y_n)$ and $z_h(Y_n)$ both converge to a (random) point denoted $\Xi \in \partial\mathbb{D}$,*
2. *The law of Ξ has full support and no atoms.*
3. *$\partial\mathbb{D}$ is a realisation of the Poisson boundary of G. That is, for every bounded harmonic function h on G there exists a bounded measurable function $g : \partial\mathbb{D} \to \mathbb{R}$ such that*

$$h(v) = E_v^G\big[g(\Xi)\big].$$

Since the additional components needed to prove this are straightforward, we omit some of the details.

Sketch of proof. First, $(\mathrm{core}(G), \rho)$ is unimodular when sampled conditional on $\rho \in \mathrm{core}(G)$: essentially, a mass transport on $\mathrm{core}(G)$ gives a mass transport on G which is 0 for all deleted vertices. The mass transport principle for G implies the principle for $\mathrm{core}(G)$.

Second, $\mathrm{core}(G)$ is CP hyperbolic. Since (G, ρ) is invariantly non-amenable, it is VEL hyperbolic (see the proof of Theorem 7.1.1). Because the infimum over paths in the definition of the vertex extremal length is the same as the infimum over paths in the core, the vertex extremal length from $v \in \mathrm{core}(G)$ to ∞ is the same in G and $\mathrm{core}(G)$. (Alternatively, one could deduce non-amenability of $\mathrm{core}(G)$ from non-amenability of G, and apply Theorem 7.1.1.)

Now, since $\mathrm{core}(G)$ is a weighted CP hyperbolic unimodular simple triangulation (and the second moment of the degree of the root is finite), by Theorem 7.1.3 the random walk on $\mathrm{core}(G)$ converges to a point in the boundary, the exit measure has full support and no atoms, and $\partial\mathbb{D}$ is a realisation of the Poisson boundary of $\mathrm{core}(G)$.

Finally, by the Optional Stopping Theorem, the bounded harmonic functions on $\mathrm{core}(G)$ are in one-to-one correspondence with the bounded harmonic functions on G by restriction and extension:

$$h_G \mapsto h_{\mathrm{core}(G)} = h_G|_{\mathrm{core}(G)}, \quad h_{\mathrm{core}(G)} \mapsto h_G(v) = E_v^G[h_{\mathrm{core}(G)}(X_{N_0})].$$

Thus, the realisation of $\partial\mathbb{D}$ as the Poisson boundary of $\mathrm{core}(G)$ extends to G. □

7.8 Open problems

Problem 7.8.1. Can the identification of the Poisson and geometric boundaries be strengthened to an identification of the Martin boundary? This was done in [17] for CP hyperbolic triangulations with bounded degrees. Specifically, we believe the following.

Conjecture 7.8.2. *Let (G, ρ) be an infinite simple, one-ended, CP hyperbolic unimodular random rooted planar triangulation with $\mathbb{E}[\deg(\rho)] < \infty$, and let $\mathcal{C}$ be a circle packing of G in the unit disc. Then almost surely for every point $\xi \in \partial\mathbb{D}$ there exists a unique positive harmonic function h_ξ on G such that $h_\xi(\rho) = 1$ and h_ξ is bounded on $\{v : |z(v) - \xi| \geq \varepsilon\}$ for every $\varepsilon > 0$. Moreover, the function $\xi \mapsto h_\xi$ almost surely extends to a homeomorphism from $z(V) \cup \partial\mathbb{D}$ to the Martin compactification of G.*

Problem 7.8.3 (Hölder continuity of the exit measures)**.** In the setting of Theorem 7.1.3, do there exist positive constants c and C such that

$$P_\rho^G(\Xi \in I) \leq C|I|^c$$

for every interval $I \subset \partial\mathbb{D}$?

Problem 7.8.4 (Dirichlet energy of z)**.** In the bounded degree case, by applying the main theorem of [12], convergence to the boundary may be shown by observing that the Dirichlet energy of the centres function z is finite:

$$\mathcal{E}(z) = \sum_{u \sim v} (z(u) - z(v))^2 \leq \sum 2\deg(v) r(v)^2 \leq 2\max\{\deg(v)\}.$$

Is the Dirichlet energy of z almost surely finite for a unimodular random rooted CP hyperbolic triangulation? This may provide a route to weakening the moment assumption in our results.

Problem 7.8.5 (Other embeddings)**.** How does the canonical embedding of the Poisson-Delaunay triangulation differ from the embedding given by the circle packing? Is there a circle packing so that $d_{hyp}(v, z_h(v))$ is stationary?

The **conformal embedding** of a triangulation is defined by forming a Riemann surface by gluing equilateral triangles according to the combinatorics of the triangulation before mapping the resulting surface conformally to $\mathbb{D}$ or $\mathbb{C}$. Is it possible to control the large scale distortion between the conformal embedding and the circle packing? In general the answer is no, but in the unimodular case there is hope.

Regardless of the answer to this question, our methods should extend without too much difficulty to establish analogues of Theorem 7.1.3 for these other embeddings, the main obstacle being to show almost sure convergence of the random walk to a point in the boundary $\partial\mathbb{D}$.

Problem 7.8.6. Reduce the moment assumption on $\deg(\rho)$ in Theorems 7.1.3 and 7.1.4. Finite expectation is needed to switch to a reversible distribution on rooted maps, but perhaps the second moment is not needed.

Chapter 8

Boundaries of planar graphs: a unified approach

Summary. We give a new proof that the Poisson boundary of a planar graph coincides with the boundary of its square tiling and with the boundary of its circle packing, originally proven by Georgakopoulos [100] and Angel, Barlow, Gurel-Gurevich and Nachmias [17] respectively. Our proof is robust, and also allows us to identify the Poisson boundaries of graphs that are rough-isometric to planar graphs.

We also prove that the boundary of the square tiling of a bounded degree plane triangulation coincides with its Martin boundary. This is done by comparing the square tiling of the triangulation with its circle packing.

8.1 Introduction

Square tilings of planar graphs were introduced by Brooks, Smith, Stone and Tutte [60], and are closely connected to random walk and potential theory on planar graphs. Benjamini and Schramm [48] extended the square tiling theorem to infinite, *uniquely absorbing* plane graphs (see Section 8.2.2). These square tilings take place on the cylinder $\mathbb{R}/\eta\mathbb{Z} \times [0,1]$, where η is the effective conductance to infinity from some fixed root vertex ρ of G. They also proved that the random walk on a transient, bounded degree, uniquely absorbing plane graph converges to a point in the boundary of the cylinder $\mathbb{R}/\eta\mathbb{Z} \times \{1\}$, and that the limit point of a random walk started at ρ is distributed according to the Lebesgue measure on the boundary of the cylinder.

Benjamini and Schramm [48] applied their convergence result to deduce that every transient, bounded degree planar graph admits *non-constant bounded harmonic functions*. Recall that a function $h : V \to \mathbb{R}$ on the state space of a Markov chain (V, P) is **harmonic** if

$$h(u) = \sum_{v \sim u} P(u,v)h(v)$$

for every vertex $u \in V$, or equivalently if $\langle h(X_n)\rangle_{n\geq 0}$ is a martingale when $\langle X_n\rangle_{n\geq 0}$ is a trajectory of the Markov chain. If G is a transient, uniquely absorbing, bounded degree plane graph, then for each bounded Borel function $f : \mathbb{R}/\eta\mathbb{Z} \to \mathbb{R}$, we define a harmonic function h on G by setting

$$h(v) = E_v \left[f \left(\lim_{n\to\infty} \theta(X_n) \right) \right]$$

for each $v \in V$, where E_v denotes the expectation with respect to a random walk $\langle X_n \rangle_{n\geq 0}$ started at v and $\theta(v)$ is the horizontal coordinate associated to the vertex v by the square tiling of G (see Section 8.2.3). Georgakopoulos [100] proved that moreover *every* bounded harmonic function on G may be represented this way, answering a question of Benjamini and Schramm [48]. In other words, Georgakopoulos's theorem identifies the geometric boundary of the square tiling of G with the *Poisson boundary* of G (see Section 8.3). Probabilistically, this means that the tail σ-algebra of the random walk $\langle X_n \rangle_{n\geq 0}$ is trivial conditional on the limit of $\theta(X_n)$.

In this paper, we give a new proof of Georgakopoulos's theorem. We state our result in the natural generality of plane networks. Recall that a **network** (G, c) is a connected, locally finite graph $G = (V, E)$, possibly containing self-loops and multiple edges, together with a function $c : E \to (0, \infty)$ assigning a positive **conductance** to each edge of G. The conductance $c(v)$ of a vertex v is defined to be the sum of the conductances of the edges emanating from v, and for each pair of vertices u, v the conductance $c(u, v)$ is defined to be the sum of the conductances of the edges connecting u to v. The random walk on the network is the Markov chain with transition probabilities $p(u, v) = c(u, v)/c(u)$. Graphs without specified conductances are considered networks by setting $c(e) \equiv 1$. We will usually suppress the notation of conductances, and write simply G for a network. Instead of square tilings, general plane networks are associated to *rectangle tilings*, see Section 8.2.3. See Section 8.2.2 for detailed definitions of plane graphs and networks. For each vertex v of G, $I(v) \subseteq \mathbb{R}/\eta\mathbb{Z}$ is an interval associated to v by the rectangle tiling of G.

Theorem 8.1.1 (Identification of the Poisson boundary)**.** *Let G be a plane network and let $\mathcal{S}_\rho$ be the rectangle tiling of G in the cylinder $\mathbb{R}/\eta\mathbb{Z} \times [0, 1]$. Suppose that $\theta(X_n)$ converges to a point in $\mathbb{R}/\eta\mathbb{Z}$ and that* $\mathrm{length}(I(X_n))$ *converges to zero almost surely as n tends to infinity. Then for every bounded harmonic function h on G, there exists a bounded Borel function $f : \mathbb{R}/\eta\mathbb{Z} \to \mathbb{R}$ such that*

$$h(v) = E_v \left[f \left(\lim_{n\to\infty} \theta(X_n) \right) \right].$$

for every $v \in V$. That is, the geometric boundary of the rectangle tiling of G coincides with the Poisson boundary of G.

The convergence theorem of Benjamini and Schramm [48] implies that the hypotheses of Theorem 8.1.1 are satisfied when G has bounded degrees.

8.1.1 Circle packing

An alternative framework in which to study harmonic functions on planar graphs is given by the circle packing theorem. A **circle packing** is a collection $\mathcal{C}$ of non-overlapping (but possibly

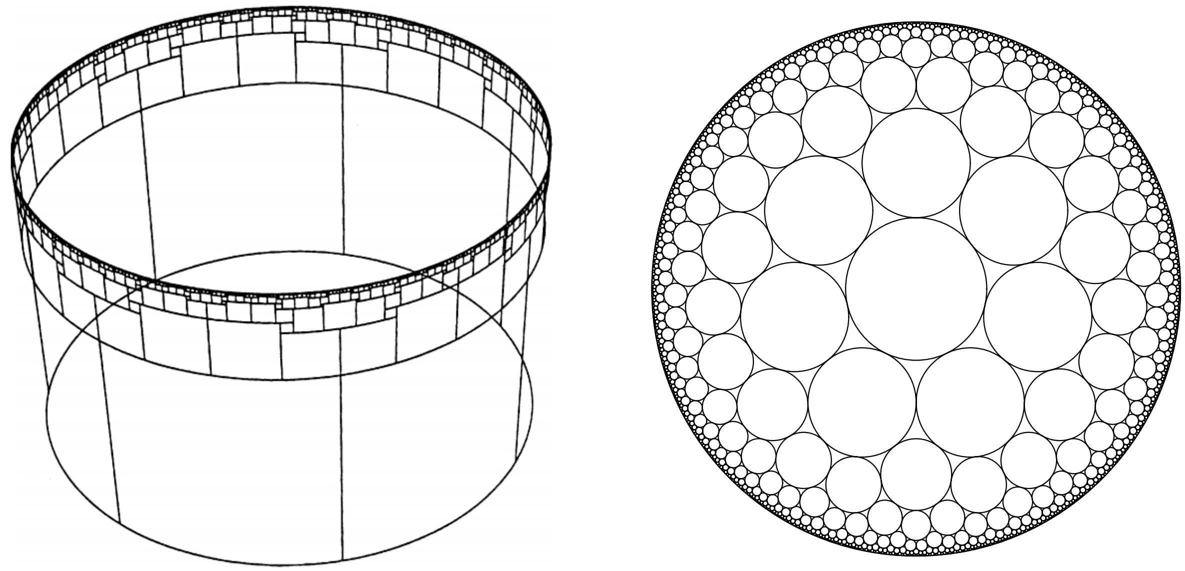

Figure 8.1: The square tiling and the circle packing of the 7-regular hyperbolic triangulation.

tangent) discs in the plane. Given a circle packing $\mathcal{C}$, the **tangency graph** of $\mathcal{C}$ is defined to be the graph with vertices corresponding to the discs of $\mathcal{C}$ and with two vertices adjacent if and only if their corresponding discs are tangent. The tangency graph of a circle packing is clearly planar, and can be drawn with straight lines between the centres of tangent discs in the packing. The Koebe-Andreev-Thurston Circle Packing Theorem [156, 221] states conversely that every finite, simple (i.e., containing no self-loops or multiple edges), planar graph may be represented as the tangency graph of a circle packing. If the graph is a triangulation (i.e., every face has three sides), its circle packing is unique up to Möbius transformations and reflections. We refer the reader to [215] and [202] for background on circle packing.

The **carrier** of a circle packing is defined to be the union of all the discs in the packing together with the bounded regions that are disjoint from the discs in the packing and are enclosed by the some set of discs in the packing corresponding to a face of the tangency graph. Given some planar domain D, we say that a circle packing is **in** D if its carrier is D.

The circle packing theorem was extended to infinite planar graphs by He and Schramm [121], who proved that every proper plane triangulation admits a locally finite circle packing in the plane or the disc, but not both. We call a triangulation of the plane **CP parabolic** if it can be circle packed in the plane and **CP hyperbolic** otherwise. He and Schramm also proved that a bounded degree simple triangulation of the plane is CP parabolic if and only if it is recurrent for the simple random walk.

Benjamini and Schramm [47] proved that, when a bounded degree, CP hyperbolic triangulation is circle packed in the disc, the simple random walk converges to a point in the boundary of the disc and the law of the limit point is non-atomic and has full support. Angel, Barlow, Gurel-Gurevich and Nachmias [17] later proved that, under the same assumptions, the boundary of the disc is a realisation of the Poisson boundary of the triangulation. These results were extended to unimodular random rooted triangulations of unbounded degree by Angel, Hutchcroft, Nachmias and Ray [21]. Our proof of Theorem 8.1.1 is adapted from the proof of [21], and also yields a new proof of the Poisson boundary result of [17], which follows as a special case of both Theorems 8.1.2 and 8.1.3 below.

8.1.2 Robustness under rough isometries

The proof of Theorem 8.1.1 is quite robust, and also allows us to characterise the Poisson boundaries of certain non-planar networks.

Benjamini and Schramm [47] proved that every transient network of bounded local geometry that is rough isometric to a planar graph admits non-constant bounded harmonic functions. In general, however, rough isometries do not preserve the property of admitting non-constant bounded harmonic functions [47, Theorem 3.5], and consequently do not preserve Poisson boundaries.

Our next theorem establishes that, for a bounded degree graph G roughly isometric to a bounded degree proper plane graph G', the Poisson boundary of G coincides with the geometric boundary of a suitably chosen embedding of G', so that the same embedding gives rise to a realisation of the Poisson boundaries of both G and G'. See Section 8.2.2 for the definition of an embedding of a planar graph.

Theorem 8.1.2 (Poisson boundaries of roughly planar networks)**.** *Let G be a transient network with bounded local geometry such that there exists a proper plane graph G' with bounded degrees and a rough isometry $\phi : G \to G'$. Let $\langle X_n \rangle_{n\geq 0}$ be a random walk on G. Then there exists an embedding z of G' into $\mathbb{D}$ such that $z \circ \phi(X_n)$ converges to a point in $\partial\mathbb{D}$ and the law of the limit point is non-atomic. Moreover, for every such embedding z and for every bounded harmonic function h on G, there exists a bounded Borel function $f : \partial\mathbb{D} \to \mathbb{R}$ such that*

$$h(v) = E_v \left[f\left(\lim_{n\to\infty} z \circ \phi(X_n) \right) \right].$$

for every $v \in V$. That is, the geometric boundary of the disc coincides with the Poisson boundary of G.

The part of Theorem 8.1.2 concerning the existence of an embedding is implicit in [47].

A further generalisation of Theorem 8.1.1 concerns embeddings of possibly irreversible planar Markov chains: The only changes required to the proof of Theorem 8.1.1 in order to prove the following are notational.

Theorem 8.1.3. *Let (V, P) be a Markov chain such that the graph*

$$G = \Big(V, \{(u, v) \in V^2 : P(u, v) > 0 \text{ or } P(v, u) > 0\}\Big)$$

is planar. Suppose further that there exists a vertex $\rho \in V$ such that for every $v \in V$ there exists n such that $P^n(\rho, v) > 0$, and let $\langle X_n \rangle_{n\geq 0}$ be a trajectory of the Markov chain. Let z be a (not *necessarily proper) embedding of G into the unit disc $\mathbb{D}$ such that $\langle z(X_n)\rangle_{n\geq 0}$ converges to a point in $\partial\mathbb{D}$ almost surely and the law of the limit point is non-atomic. Then for every bounded harmonic function h on (V,P), there exists a bounded Borel function $f:\partial\mathbb{D}\to\mathbb{R}$ such that*

$$h(v) = E_v\left[f\left(\lim_{n\to\infty} z(X_n)\right)\right].$$

for every $v\in V$.

8.1.3 The Martin boundary

In [17] it was also proven that the *Martin boundary* of a bounded degree CP hyperbolic triangulation can be identified with the geometric boundary of its circle packing. Recall that a function $g : V \to \mathbb{R}$ on a network G is **superharmonic** if

$$g(u) \geq \frac{1}{c(u)}\sum_{v\sim u} c(u,v)g(v)$$

for every vertex $u \in V$. Let ρ be a fixed vertex of G and consider the space $\mathfrak{S}^+$ of positive superharmonic functions g on G such that $g(\rho) = 1$, which is a convex, compact subset of the space of functions $V \to \mathbb{R}$ equipped with the product topology (i.e. the topology of pointwise convergence). For each $v \in V$, let P_v be the law of the random walk on G started at v. We can embed V into $\mathfrak{S}^+$ by sending each vertex u of G to its **Martin kernel**

$$M_u(v) := \frac{E_v\left[\#(\text{visits to } u)\right]}{E_\rho\left[\#(\text{visits to } u)\right]} = \frac{P_v(\text{hit } u)}{P_\rho(\text{hit } u)}.$$

The **Martin compactification** $\mathcal{M}(G)$ of the network G is defined as the closure of $\{M_u : u \in V\}$ in $\mathfrak{S}^+$, and the **Martin boundary** $\partial\mathcal{M}(G)$ of the network G is defined to be the complement of the image of V,

$$\partial\mathcal{M}(G) := \mathcal{M}(G) \setminus \{M_u : u \in V\}.$$

See [84, 201, 229] for background on the Martin boundary.

Our next result is that, for a triangulation of the plane with bounded local geometry, the geometric boundary of the square tiling coincides with the Martin boundary.

Theorem 8.1.4 (Identification of the Martin boundary)**.** *Let T be a transient, simple, proper plane triangulation with bounded local geometry. Let $\mathcal{S}_\rho$ be a square tiling of T in a cylinder $\mathbb{R}/\eta\mathbb{Z}\times[0,1]$. Then*

1. *A sequence of vertices $\langle v_n\rangle_{n\geq 0}$ in T converges to a point in the Martin boundary of T if and only if $y(v_n)\to 1$ and $\theta(v_n)$ converges to a point in $\mathbb{R}/\eta\mathbb{Z}$.*

2. *The map*

$$M : \theta \longmapsto M_\theta := \lim_{n\to\infty} M_{v_n} \text{ where } \big(\theta(v_n), y(v_n)\big) \to (\theta, 1),$$

which is well-defined by (1)*, is a homeomorphism from* $\mathbb{R}/\eta\mathbb{Z}$ *to the Martin boundary* $\partial\mathcal{M}(T)$ *of* T*.*

That is, the geometric boundary of the rectangle tiling of G *coincides with the Martin boundary of* G*.*

This will be deduced from the analogous statement for circle packings [17, Theorem 1.2] together with the following theorem, which states that for a bounded degree triangulation T, the square tiling and circle packing of T define equivalent compactifications of T.

Theorem 8.1.5 (Comparison of square tiling and circle packing)**.** *Let* T *be a transient, simple, proper plane triangulation with bounded local geometry. Let* $\mathcal{S}_\rho$ *be a rectangle tiling of* T *in the cylinder* $\mathbb{R}/\eta\mathbb{Z} \times [0,1]$ *and let* $\mathcal{C}$ *be a circle packing of* T *in* $\mathbb{D}$ *with associated embedding* z*. Then*

1. *A sequence of vertices* $\langle v_n \rangle_{n\geq 0}$ *in* T *converges to a point in* $\partial\mathbb{D}$ *if and only if* $y(v_n) \to 1$ *and* $\theta(v_n)$ *converges to a point in* $\mathbb{R}/\eta\mathbb{Z}$*.*

2. *The map*

$$\xi \mapsto \lim_{i\to\infty} \theta(v_i) \text{ where } z(v_i) \to \xi,$$

which is well-defined by (1)*, is a homeomorphism from* $\partial\mathbb{D}$ *to* $\mathbb{R}/\eta\mathbb{Z}$*.*

Theorem 8.1.5 also allows us to deduce the Poisson boundary results of [100] and [17] from each other in the case of bounded degree triangulations.

Theorem 8.1.4 has the following immediate corollaries by standard properties of the Martin compactification.

Corollary 8.1.6 (Continuity of harmonic densities)**.** *Let* T *be a transient proper simple plane triangulation with bounded local geometry, and let* $\mathcal{S}_\rho$ *be a rectangle tiling of* T *in the cylinder* $\mathbb{R}/\eta\mathbb{Z} \times [0,1]$*. For each vertex* v *of* T*, let* ω_v *denote the harmonic measure from* v*, defined by*

$$\omega_v(\mathscr{A}) := P_v\Big(\lim_{n\to\infty} \theta(X_n) \in \mathscr{A}\Big)$$

for each Borel set $\mathscr{A} \subseteq \mathbb{R}/\eta\mathbb{Z}$*, and let* $\lambda = \omega_\rho$ *denote the Lebesgue measure on* $\mathbb{R}/\eta\mathbb{Z}$*. Then for every* v *of* T*, the density of the harmonic measure from* v *with respect to Lebesgue is given by*

$$\frac{d\omega_v}{d\lambda}(\theta) = M_\theta(v)$$

which is continuous with respect to θ*.*

Corollary 8.1.7 (Representation of positive harmonic functions)**.** *Let* T *be a transient proper simple plane triangulation with bounded local geometry, and let* $\mathcal{S}_\rho$ *be a rectangle tiling of* T *in the cylinder* $\mathbb{R}/\eta\mathbb{Z} \times [0,1]$*. Then for every positive harmonic function* h *on* T*, there exists a unique measure* μ *on* $\mathbb{R}/\eta\mathbb{Z}$ *such that*

$$h(v) = \int_{\mathbb{R}/\eta\mathbb{Z}} M_\theta(v)\,\mathrm{d}\mu(\theta).$$

for every $v \in V$.

8.2 Background

8.2.1 Notation

We use e to denote both oriented and unoriented edges of a graph or network. An oriented edge e is oriented from its tail e^- to its head e^+. Given a network G and a vertex v of G, we write P_v for the law of the random walk on G started at v and E_v for the associated expectation operator.

8.2.2 Embeddings of planar graphs

Let $G = (V, E)$ be a graph. For each edge e, choose an orientation of e arbitrarily and let $I(e)$ be an isometric copy of the interval $[0, 1]$. The metric space $\mathbb{G} = \mathbb{G}(G)$ is defined to be the quotient of the union $\bigcup_e I(e) \cup V$, where we identify the endpoints of $I(e)$ with the vertices e^- and e^+ respectively, and is equipped with the path metric. An **embedding of** G into a surface S is a continuous, injective map $z : \mathbb{G} \to S$. The embedding is **proper** if every compact subset of S intersects at most finitely many edges and vertices of $z(\mathbb{G})$. A graph is planar if and only if it admits an embedding into $\mathbb{R}^2$. A **plane graph** is a planar graph together with an embedding $\mathbb{G}(G) \to \mathbb{R}^2$ or some other surface homeomorphic to $\mathbb{R}^2$ such as the open disc; it is a **proper plane graph** if the embedding z is proper. A **(proper) plane network** is a (proper) plane graph together with an assignment of conductances $c : E \to (0, \infty)$. A **proper plane triangulation** is a proper plane network in which every face (i.e connected component of $\mathbb{R}^2 \setminus z(\mathbb{G})$) has three sides.

A set of vertices $W \subseteq V$ is said to be **absorbing** if with positive probability the random walk $\langle X_n \rangle_{n \geq 0}$ on G is contained in W for all n greater than some random N. A plane graph G is said to be **uniquely absorbing** if for every finite subgraph G_0 of G, there is exactly one connected component D of $\mathbb{R}^2 \setminus \bigcup\{z(I(e)) : e \in G_0\}$ such that the set of vertices $\{v \in V : z(v) \in D\}$ is absorbing.

8.2.3 Square tiling

Let G be a transient, uniquely absorbing plane network and let ρ be a vertex of G. For each $v \in V$ let $y(v)$ denote the probability that the random walk on G started at v never visits ρ, and let

$$\eta := \sum_{u \sim \rho} c(\rho, u) y(u).$$

Let $\mathbb{R}/\eta\mathbb{Z}$ be the circle of length η. Then there exists a set

$$\mathcal{S}_\rho = \{S(e) : e \in E\}$$

such that:

1. For each oriented edge e of G such that $y(e^+) \geq y(e^-)$, $S(e) \subseteq \mathbb{R}/\eta\mathbb{Z} \times [0,1)$ is a rectangle of the form
$$S(e) = I(e) \times \Big[y(e^-), y(e^+)\Big]$$
where $I(e) \subseteq \mathbb{R}/\eta\mathbb{Z}$ is an interval of length
$$\text{length}\big(I(e)\big) := c(e)\Big(y(e^+) - y(e^-)\Big).$$
If e is such that $y(e^+) < y(e^-)$, we define $I(e) = I(-e)$ and $S(e) = S(-e)$. In particular, the aspect ratio of $S(e)$ is equal to the conductance $c(e)$ for every edge $e \in E$.

2. The interiors of the rectangles $S(e)$ are disjoint, and the union $\bigcup_e S(e) = \mathbb{R}/\eta\mathbb{Z} \times [0,1)$.

3. For every vertex $v \in V$, the set $I(v) = \bigcup_{e^-=v} I(e)$ is an interval and is equal to $\bigcup_{e^+=v} I(e)$.

4. For almost every $\theta \in C$ and for every $t \in [0,1)$, the line segment $\{\theta\} \times [0,t]$ intersects only finitely many rectangles of $\mathcal{S}_\rho$.

Note that the rectangle corresponding to an edge through which no current flows is degenerate, consisting of a single point. The existence of the above tiling was proven by Benjamini and Schramm [48]. Their proof was stated for the case $c \equiv 1$ but extends immediately to our setting, see [100].

Let us also note the following property of the rectangle tiling, which follows from the construction given in [48].

(5) For each two edges e_1 and e_2 of G, the interiors of the vertical sides of the rectangles $S(e_1)$ and $S(e_2)$ have a non-trivial intersection only if e_1 and e_2 both lie in the boundary of some common face f of G.

For each $v \in V$, we let $\theta(v)$ be a point chosen arbitrarily from $I(v)$.

Let G be a uniquely absorbing proper plane network. Benjamini and Schramm [48] proved that if G has bounded local geometry and $\langle X_n \rangle_{n\geq 0}$ is a random walk on G started at ρ, then $\theta(X_n)$ converges to a point in $\mathbb{R}/\eta\mathbb{Z}$ and the law of the limit point is Lebesgue (their proof is given for bounded degree plane graphs but extends immediately to this setting). An observation of Georgakopoulos [100, Lemma 6.2] implies that, more generally, whenever G is such that $\theta(X_n)$ converges to a point in $\mathbb{R}/\eta\mathbb{Z}$ and $\text{length}(I(X_n))$ converges to zero almost surely, the law of the limit point is Lebesgue.

8.3 The Poisson boundary

Let (V, P) be a Markov chain. Harmonic functions on (V, P) encode asymptotic behaviours of a trajectory $\langle X_n \rangle_{n\geq 0}$ as follows. Let Ω denote the path space
$$\Omega = \Big\{\langle x_i \rangle_{i\geq 0} \in V^{\mathbb{N}} : p(x_i, x_{i+1}) > 0 \quad \forall i \geq 0\Big\}$$

and let $\mathcal{B}$ denote the Borel σ-algebra for the product topology on Ω. Let $\mathcal{I}$ denote the invariant σ-algebra

$$\mathcal{I} = \left\{\mathscr{A} \in \mathcal{B} : \langle x_i \rangle_{i\geq 0} \in \mathscr{A} \iff \langle x_{i+1} \rangle_{i\geq 0} \in \mathscr{A} \quad \forall \langle x_i \rangle_{i\geq 0} \in \Omega \right\}.$$

Assume that there exists a vertex $\rho \in V$ from which all other vertices are reachable:

$$\forall v \in V \ \exists k \geq 0 \text{ such that } p_k(\rho, v) > 0$$

which is always satisfied when P is the transition operator of the random walk on a network G. Then there exists an invertible linear transformation $\mathbf{H}$ between $L^\infty(\Omega, \mathcal{I}, P_\rho)$ and the space of bounded harmonic functions on G defined as follows [55, 173, Proposition 14.12]:

$$\begin{aligned} \mathbf{H} : f \longmapsto \mathbf{H}f(v) &= E_v\left[f\left(\langle X_n \rangle_{n\geq 0}\right)\right] \\ \mathbf{H}^{-1} : h \longmapsto \tilde{h}\left(\langle x_i \rangle_{i\geq 0}\right) &= \lim_{i\to\infty} h(x_i) \end{aligned} \tag{8.3.1}$$

where the above limit exists for P_ρ-a.e. sequence $\langle x_i \rangle_{i\geq 0}$ by the martingale convergence theorem.

Proposition 8.3.1 (Path-hitting criterion for the Poisson boundary)**.** *Let (V, P) be a Markov chain and let $\langle X_n \rangle_{n\geq 0}$ be a trajectory of the Markov chain. Suppose that $\psi : V \to \mathbb{M}$ is a function from V to a metric space $\mathbb{M}$ such that $\psi(X_n)$ converges to a point in $\mathbb{M}$ almost surely. For each $k \geq 0$, let $\langle Z_j^k \rangle_{j\geq 0}$ be a trajectory of the Markov chain started at X_k that is conditionally independent of $\langle X_n \rangle_{n\geq 0}$ given X_k, and let $\mathbb{P}$ denote the joint distribution of $\langle X_n \rangle_{n\geq 0}$ and each of the $\langle Z_m^k \rangle_{m\geq 0}$. Suppose that almost surely, for every path $\langle v_i \rangle_{i\geq 0} \in \Omega$ started at ρ such that $\lim_{i\to\infty} \psi(v_i) = \lim_{n\to\infty} \psi(X_n)$, we have that*

$$\limsup_{k\to\infty} \mathbb{P}\left(\langle Z_m^k \rangle_{m\geq 0} \text{ hits } \{v_i : i \geq 0\} \,\Big|\, \langle X_n \rangle_{n\geq 0}\right) > 0. \tag{8.3.2}$$

Then for every bounded harmonic function h on G, there exists a bounded Borel function $f : \mathbb{M} \to \mathbb{R}$ such that

$$h(v) = E_v\left[f\left(\lim_{n\to\infty} \psi(X_n)\right)\right]$$

for every $v \in V$.

Note that it suffices to define f on the support of the law of $\lim_{n\to\infty} \psi(X_n)$, which is contained in the set of accumulation points of $\{\psi(v) : v \in V\}$.

A consequence of the correspondence (8.3.1) is that, to prove Proposition 8.3.1, it suffices to prove that for every invariant event $\mathscr{A} \in \mathcal{I}$, there exists a Borel set $\mathscr{B} \subseteq \mathbb{M}$ such that

$$P_v\left(\mathscr{A} \triangle \{\lim_{n\to\infty} \psi(X_n) \in \mathscr{B}\}\right) = 0 \text{ for every vertex } v \in V.$$

Proof of Proposition 8.3.1. Let $\mathscr{A} \in \mathcal{I}$ be an invariant event and let h be the harmonic function

$$h(v) := P_v(\langle X_n \rangle_{n\geq 0} \in \mathscr{A}).$$

Lévy's 0-1 law implies that

$$h(X_n) \xrightarrow[n\to\infty]{a.s.} \mathbb{1}\left(\langle X_n\rangle_{n\geq 0} \in \mathscr{A}\right)$$

and so it suffices to exhibit a Borel set $\mathscr{B} \subseteq \mathbb{M}$ such that

$$P_\rho\left(\mathscr{A} \triangle \big\{\lim_{n\to\infty} \psi(X_n) \in \mathscr{B}\big\}\right)$$
$$= P_\rho\left(\big\{\limsup_{n\to\infty} h(X_n) > 0\big\} \triangle \big\{\lim_{n\to\infty} \psi(X_n) \in \mathscr{B}\big\}\right) = 0.$$

We may assume that $P_\rho(\mathscr{A}) > 0$, otherwise the claim is trivial.

Let $d_{\mathbb{M}}$ denote the metric of $\mathbb{M}$. For each natural number $m > 0$, let $N(m)$ be the smallest natural number such that

$$P_\rho\left(\exists n \geq N(m) \text{ such that } d_{\mathbb{M}}\left(\psi(X_n), \lim_{k\to\infty}\psi(X_k)\right) \geq \frac{1}{m}\right) \leq 2^{-m}.$$

For each n, let $m(n)$ be the largest m such that $n \geq N(m)$, so that $m(n) \to \infty$ as $n \to \infty$. Borel-Cantelli implies that, for all but finitely many m, there does not exist an $n \geq N(m)$ such that

$$d_{\mathbb{M}}\left(\psi(X_n), \lim_{k\to\infty}\psi(X_k)\right) > \frac{1}{m},$$

and it follows that

$$d_{\mathbb{M}}\left(\psi(X_n), \lim_{k\to\infty}\psi(X_k)\right) \leq \frac{1}{m(n)}$$

for all but finitely many n almost surely.

Define the set $\mathscr{B} \subseteq \mathbb{M}$ by

$$\mathscr{B} := \left\{ x \in \mathbb{M} : \begin{array}{l} \exists \text{ a path } \langle v_i\rangle_{i\geq 0} \text{ in } G \text{ with } v_0 = \rho \text{ such that } d_{\mathbb{M}}(\psi(v_i), x) \leq 1/m(i) \\ \text{for all but finitely many } i \text{ and } \inf_{i\geq 0} h(v_i) > 0 \end{array} \right\}.$$

To see that $\mathscr{B}$ is Borel, observe that it may be written in terms of closed subsets of $\mathbb{M}$ as follows

$$\mathscr{B} = \bigcup_{k\geq 0}\bigcup_{j\geq 0}\bigcap_{I\geq j} \left\{ x \in \mathbb{M} : \begin{array}{l} \exists \text{ a path } \langle v_i\rangle_{i=0}^{I} \text{ in } G \text{ with } v_0 = \rho \text{ such that} \\ d_{\mathbb{M}}(\psi(v_i), x) \leq 1/m(i) \text{ for all } j \leq i \leq I \text{ and} \\ h(v_i) \geq 1/k \text{ for all } 0 \leq i \leq I \end{array} \right\}.$$

It is immediate that $\lim_{n\to\infty}\psi(X_n) \in \mathscr{B}$ almost surely on the event that $h(X_n)$ converges to 1: simply take $\langle v_i\rangle_{i\geq 0} = \langle X_i\rangle_{i\geq 0}$ as the required path. In particular, the event $\{\lim_{n\to\infty}\psi(X_n) \in \mathscr{B}\}$

has positive probability.

We now prove conversely that $\liminf_{n\to\infty} h(X_n) > 0$ almost surely on the event that $\lim_{n\to\infty} \psi(X_n) \in \mathscr{B}$. Condition on this event, so that there exists a path $\langle v_i \rangle_{i\geq 0}$ in G starting at ρ such that $\lim_{i\to\infty} \psi(v_i) = \lim_{n\to\infty} \psi(X_n)$ and $\inf_{i\geq 0} h(v_i) > 0$. Fix one such path. Applying the optional stopping theorem to $\langle h(Z^k_m)\rangle_{m\geq 0}$, we have

$$h(X_k) \geq \mathbb{P}\left(\langle Z^k_m\rangle_{m\geq 0} \text{ hits } \{v_i : i \geq 0\} \,\middle|\, \langle X_n\rangle_{n\geq 0}\right) \cdot \inf\{h(v_i) : i \geq 0\}$$

and so, by our assumption (8.3.2), we have that

$$\limsup_{k\to\infty} h(X_k) \geq \limsup_{k\to\infty} \mathbb{P}\left(\langle Z^k_m\rangle_{m\geq 0} \text{ hits } \{v_i : i \geq 0\} \,\middle|\, \langle X_n\rangle_{n\geq 0}\right) \cdot \inf\{h(v_i) : i \geq 0\}$$

is positive almost surely. □

We remark that controlling the rate of convergence of the path in the definition of $\mathscr{B}$ can be avoided by invoking the theory of universally measurable sets.

We now apply the criterion given by Proposition 8.3.1 to prove Theorem 8.1.1 and Theorem 8.1.2.

Proof of Theorem 8.1.1. Let $\langle X_n\rangle_{n\geq 0}$ and $\langle Y_n\rangle_{n\geq 0}$ be independent random walks on G started at ρ, and for each $k \geq 0$ let $\langle Z^k_j\rangle_{j\geq 0}$ be a random walk on G started at X_k that is conditionally independent of $\langle X_n\rangle_{n\geq 0}$ and $\langle Y\rangle_{n\geq 0}$ given X_k. Let $\mathbb{P}$ denote the joint distribution of $\langle X_n\rangle_{n\geq 0}$, $\langle Y_n\rangle_{n\geq 0}$ and all of the random walks $\langle Z^k_m\rangle_{m\geq 0}$. Given two points $\theta_1, \theta_2 \in \mathbb{R}/\eta\mathbb{Z}$, we denote by $(\theta_1,\theta_2) \subset \mathbb{R}/\eta\mathbb{Z}$ the open arc between θ_1 and θ_2 in the counter-clockwise direction. For each such interval $(\theta_1,\theta_2) \in \mathbb{R}/\eta\mathbb{Z}$, let

$$q_{(\theta_1,\theta_2)}(v) := P_v\left(\lim_{n\to\infty} \theta(X_n) \in (\theta_1,\theta_2)\right).$$

be the probability that a random walk started at v converges to a point in the interval (θ_1,θ_2).

Since the law of $\lim_{n\to\infty}\theta(X_n)$ is Lebesgue and hence non-atomic, the two random variables $\theta^+ := \lim_{n\to\infty}\theta(X_n)$ and $\theta^- := \lim_{n\to\infty}\theta(Y_n)$ are almost surely distinct. We can therefore write $\mathbb{R}/\eta\mathbb{Z} \setminus \{\theta^+,\theta^-\}$ as the union of the two disjoint non-empty intervals $\mathbb{R}/\eta\mathbb{Z} \times \{1\} \setminus \{\theta^+,\theta^-\} = (\theta^+,\theta^-) \cup (\theta^-,\theta^+)$. Let

$$Q_k := q_{(\theta^-,\theta^+)}(X_k) = \mathbb{P}\left(\lim_{m\to\infty} \theta(Z^k_m) \in (\theta^-,\theta^+) \,\middle|\, \langle X_n\rangle_{n\geq 0}, \langle Y_n\rangle_{n\geq 0}\right)$$

be the probability that a random walk started at X_k, that is conditionally independent of $\langle X_n\rangle_{n\geq 0}$ and $\langle Y_n\rangle_{n\geq 0}$ given X_k, converges to a point in the interval (θ^-,θ^+).

We claim that the random variable Q_k is uniformly distributed on $[0,1]$ conditional on $\langle X_n\rangle_{n=0}^k$ and $\langle Y_n\rangle_{n\geq 0}$. Indeed, since the law of θ^+ given X_k is non-atomic, for each $s \in [0,1]$ there exists[9]

[9]In the present setting, θ^s is unique since the ;aw of θ^+ given X_k has full support.

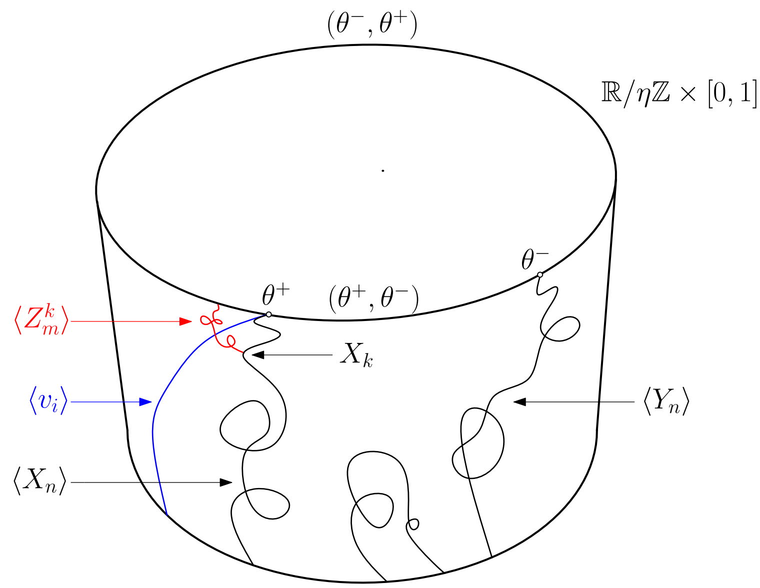

Figure 8.2: Illustration of the proof. Conditioned on the random walk $\langle X_n \rangle_{n \geq 0}$, there exists a random $\varepsilon > 0$ such that almost surely, for infinitely many k, a new random walk $\langle Z_m^k \rangle_{m \geq 0}$ (red) started from X_k has probability at least ε of hitting the path $\langle v_i \rangle_{i \geq 0}$ (blue).

$\theta^s = \theta^s(X_k, \theta^-) \in \mathbb{R}/\eta\mathbb{Z}$ such that

$$\mathbb{P}\left(\theta^+ \in (\theta^-, \theta^s) \,\middle|\, X_k, \theta^-\right) = s.$$

The claim follows by observing that

$$\mathbb{P}\left(Q_k \in [0,s] \,\middle|\, \langle X_n \rangle_{n=0}^k, \langle Y_n \rangle_{n \geq 0}\right) = \mathbb{P}\left(\theta^+ \in (\theta^-, \theta^s) \mid X_k, \theta^-\right) = s.$$

As a consequence, Fatou's lemma implies that for every $\varepsilon > 0$,

$$\begin{aligned}\mathbb{P}(Q_k \in [\varepsilon, 1-\varepsilon] \text{ infinitely often}) &= \mathbb{E}\left[\limsup_{k \to \infty} \mathbb{1}\left(Q_k \in [\varepsilon, 1-\varepsilon]\right)\right] \\ &\geq \limsup_{k \to \infty} \mathbb{P}\left(Q_k \in [\varepsilon, 1-\varepsilon]\right) = 1 - 2\varepsilon.\end{aligned}$$

and so

$$\limsup_{k \to \infty} \min\{Q_k, 1 - Q_k\} > 0 \text{ almost surely.} \tag{8.3.3}$$

Let $\langle v_i \rangle_{i \geq 0}$ be a path in G started at ρ such that $\lim_{i \to \infty} \theta(v_i) = \theta^+$ and the height $y(v_i)$ tends to 1. Observe (Figure 1) that for every $k \geq 0$, the union of the traces $\{v_i : i \geq 0\} \cup \{Y_n : n \geq 0\}$ either contains X_k or disconnects X_k from at least one of the two intervals (θ^-, θ^+) or (θ^+, θ^-). That is, for at least one of the intervals (θ^-, θ^+) or (θ^+, θ^-), any path in G started at X_k that converges to this interval must intersect $\{v_i : i \geq 0\} \cup \{Y_n : n \geq 0\}$. It follows that

$$\mathbb{P}\left(\langle Z^k_m \rangle_{m\geq 0} \text{ hits } \{v_i : i \geq 0\} \cup \{Y_n : n \geq 0\} \,\middle|\, \langle X_n \rangle_{n\geq 0},\, \langle Y_n \rangle_{n\geq 0} \right) \geq \min\{Q_k, 1-Q_k\}.$$

and so, applying (8.3.3),

$$\limsup_{k\to\infty} \mathbb{P}\left(\langle Z^k_m \rangle_{m\geq 0} \text{ hits } \{v_i : i \geq 0\} \cup \{Y_n : n \geq 0\} \,\middle|\, \langle X_n \rangle_{n\geq 0},\, \langle Y_n \rangle_{n\geq 0} \right) > 0 \tag{8.3.4}$$

almost surely. We next claim that

$$\mathbb{P}\left(\langle Z^k_m \rangle_{m\geq 0} \text{ hits } \{Y_n : n \geq 0\} \,\middle|\, \langle X_n \rangle_{n\geq 0},\, \langle Y_n \rangle_{n\geq 0} \right) \xrightarrow[k\to\infty]{a.s.} 0. \tag{8.3.5}$$

Indeed, we have that

$$\begin{aligned} &\mathbb{P}\left(\langle Z^k_m \rangle_{m\geq 0} \text{ hits } \{Y_n : n \geq 0\} \,\middle|\, \langle X_n \rangle_{n\geq 0},\, \langle Y_n \rangle_{n\geq 0} \right) \\ &\qquad = \mathbb{P}\left(\langle Z^k_m \rangle_{m\geq 0} \text{ hits } \{Y_n : n \geq 0\} \,\middle|\, \langle X_n \rangle_{n=0}^k,\, \langle Y_n \rangle_{n\geq 0} \right) \\ &\qquad\qquad = \mathbb{P}\left(\langle X_m \rangle_{m\geq k} \text{ hits } \{Y_n : n \geq 0\} \,\middle|\, \langle X_n \rangle_{n=0}^k,\, \langle Y_n \rangle_{n\geq 0} \right). \end{aligned}$$

The rightmost expression converges to zero almost surely by an easy application of Lévy's 0-1 law: For each $k_0 \leq k$, we have that

$$\begin{aligned} &\mathbb{P}\left(\langle Z^k_m \rangle_{m\geq 0} \text{ hits } \{Y_n : n \geq 0\} \,\middle|\, \langle X_n \rangle_{n\geq 0},\, \langle Y_n \rangle_{n\geq 0} \right) \\ &\qquad \leq \mathbb{P}\left(\langle X_m \rangle_{m\geq k_0} \text{ hits } \{Y_n : n \geq 0\} \,\middle|\, \langle X_n \rangle_{n=0}^k,\, \langle Y_n \rangle_{n\geq 0} \right) \\ &\qquad\qquad \xrightarrow[k\to\infty]{a.s.} \mathbb{1}(\langle X_m \rangle_{m\geq k_0} \text{ hits } \{Y_n : n \geq 0\}), \end{aligned}$$

and the claim follows since $\mathbb{1}(\langle X_m \rangle_{m\geq k_0} \text{ hits } \{Y_n : n \geq 0\}) \to 0$ a.s. as $k_0 \to \infty$.

Combining (8.3.4) and (8.3.5), we deduce that

$$\begin{aligned} &\limsup_{k\to\infty} \mathbb{P}\left(\langle Z^k_m \rangle_{m\geq 0} \text{ hits } \{v_i : i \geq 0\} \,\middle|\, \langle X_n \rangle_{n\geq 0} \right) \\ &\qquad = \limsup_{k\to\infty} \mathbb{P}\left(\langle Z^k_m \rangle_{m\geq 0} \text{ hits } \{v_i : i \geq 0\} \,\middle|\, \langle X_n \rangle_{n\geq 0},\, \langle Y_n \rangle_{n\geq 0} \right) > 0 \end{aligned}$$

almost surely. Applying Proposition 8.3.1 with $\psi = (\theta, y) : V \to \mathbb{R}/\eta\mathbb{Z} \times [0,1]$ completes the proof. □

8.3.1 Proof of Theorem 8.1.2

Proof. Let $G = (V, E)$ be a transient network with bounded local geometry such that there exists a proper plane graph $G' = (V', G')$ with bounded degrees and a rough isometry $\phi : V \to V'$. Then G' is also transient [173, §2.6]. We may assume that G' is simple, since ϕ remains a rough isometry if we modify G' by deleting all self-loops and identifying all multiple edges between each pair of vertices. In this case, there exists a simple, bounded degree, proper plane triangulation T' containing G' as a subgraph [47, Lemma 4.3]. The triangulation T' is transient by Rayleigh monotonicty, and since it has bounded degrees it is CP hyperbolic by the He-Schramm Theorem. Let $\mathcal{C}$ be a circle packing of T' in the disc $\mathbb{D}$ and let z be the embedding of G' defined by this circle packing, so that for each $v \in V'$, $z(v')$ is the center of the circle of $\mathcal{C}$ corresponding to v', and each edge of G' is embedded as a straight line between the centers of the circles corresponding to its endpoints.

Recall that the **Dirichlet energy** of a function $f : V \to \mathbb{R}$ is defined by

$$\mathcal{E}_G(f) = \frac{1}{2} \sum_e c(e) \left| f(e^+) - f(e^-) \right|^2.$$

The following lemma is implicit in [212].

Lemma 8.3.2. *Let $\phi : V \to V'$ be a rough isometry from a network G with bounded local geometry to a network G' with bounded local geometry. Then there exists a constant C such that*

$$\mathcal{E}_G(f \circ \phi) \leq C \mathcal{E}_{G'}(f). \tag{8.3.6}$$

for all functions $f : V' \to \mathbb{R}$.

For each pair of adjacent vertices u and v in G, we have that $d'(\phi(u), \phi(v)) \leq \alpha + \beta$, and so there exists a path in G' from $\phi(u)$ to $\phi(v)$ of length at most $\alpha + \beta$. Fix one such path $\Phi(u, v)$ for each u and v.

Proof. Since there at most $\alpha + \beta$ edges of G in each path $\Phi(u, v)$, and the conductances of G are bounded, there exists a constant C_1 such that

$$\begin{aligned} \mathcal{E}_G(f \circ \phi) &= \sum_e c(e) \Big(f \circ \phi(e^+) - f \circ \phi(e^-) \Big)^2 = \sum_e c(e) \Big(\sum_{e' \in \Phi(e)} f(e'^+) - f(e'^-) \Big)^2 \\ &\leq C_1 \sum_e \sum_{e' \in \Phi(e)} \Big(f(e'^+) - f(e'^-) \Big)^2. \end{aligned} \tag{8.3.7}$$

Let e' be an edge of G'. If e_1 and e_2 are two edges of G such that $\Phi(e_1^-, e_1^+)$ and $\Phi(e_1^-, e_1^+)$ both contain e', then $d'(\phi(e_1^-), \phi(e_2^-)) \leq 2(\alpha + \beta)$ and so

$$d(e_1^-, e_2^-) \leq \alpha \left(d' \Big(\phi(e_1^-), \phi(e_2^-) \Big) + \beta \right) \leq \alpha(2\alpha + 3\beta).$$

Thus, the set of edges e of G such that $\Phi(e^-, e^+)$ contains e' is contained a ball of radius $\alpha(2\alpha+3\beta)$ in G. Since G has bounded degrees, the number of edges contained in such a ball is bounded by some constant C_2. Combining this with the assumption that the conductances of G' are bounded below by some constant C_3, we have that

$$\mathcal{E}_G(f \circ \phi) \leq C_1 \sum_e \sum_{e' \in \Phi(e)} \Big(f(e'^+) - f(e'^-)\Big)^2$$
$$\leq C_1 C_2 \sum_{e'} \Big(f(e'^+) - f(e'^-)\Big)^2 \leq C_1 C_2 C_3 \mathcal{E}_{G'}(f). \qquad \square$$

The proofs of Lemmas 8.3.3 and 8.3.4 below are adapted from [17].

Lemma 8.3.3. *Let $\langle X_n \rangle_{n \geq 0}$ denote the random walk on G. Then $z \circ \phi(X_n)$ converges to a point in $\partial\mathbb{D}$ almost surely.*

Proof. Ancona, Lyons and Peres [12] proved for every function f of finite Dirichlet energy on a transient network G, the sequence $f(X_n)$ converges almost surely as $n \to \infty$. Thus, it suffices to prove that each coordinate of $z \circ \phi$ has finite Dirichlet energy. Applying the inequality (8.3.6), it suffices to prove that each coordinate of z has finite energy. For each vertex v' of G', let $r(v')$ denote the radius of the circle corresponding to v' in $\mathcal{C}$. Then, letting $z = (z_1, z_2)$,

$$\mathcal{E}_{G'}(z_1) + \mathcal{E}_{G'}(z_2) = \sum_{e'} \Big|z(e^+) - z(e^-)\Big|^2 = C \sum_{e'} \Big(r(e^+) + r(e^-)\Big)^2$$
$$\leq 2 \max\big(\deg(v')\big)\, C \sum_{v'} r(v)^2 \leq 2C \max\big(\deg(v')\big) < \infty$$

since $\sum \pi r(v)^2 \leq \pi$ is the total area of all the circles in the packing. $\square$

Lemma 8.3.4. *The law of* $\lim_{n\to\infty} z \circ \phi(X_n)$ *does not have any atoms.*

Proof. Let $B_k(\rho)$ be the set of vertices of G at graph distance at most k from ρ, and let G_k be the subnetwork of G induced by $B_k(\rho)$ (i.e. the induced subgraph together with the conductances inherited from G). Recall that the **free effective conductance** between a set A and a set B in an infinite graph G is given by

$$\mathscr{C}^{\mathrm{F}}_{\mathrm{eff}}(A \leftrightarrow B\,;\,G) = \min\big\{\mathcal{E}(F) : F(a) = 1\ \forall a \in A,\ \ F(b) = 0\ \forall b \in B\big\}.$$

The same variational formula also defines the effective conductance between two sets A and B in a finite network G, denoted $\mathscr{C}_{\mathrm{eff}}(A \leftrightarrow B; G)$. A related quantity is the effective conductance from a vertex v to infinity in G

$$\mathscr{C}_{\mathrm{eff}}(v \to \infty; G) := \lim_{k\to\infty} \mathscr{C}^{\mathrm{F}}_{\mathrm{eff}}(v \leftrightarrow V \setminus B_k(\rho); G) = \lim_{k\to\infty} \mathscr{C}^{\mathrm{F}}_{\mathrm{eff}}(v \leftrightarrow V \setminus B_k(\rho); G_{k+1}),$$

which is positive if and only if G is transient. See [173, §2 and §9] for background on electrical networks. The inequality (8.3.6) above implies that there exists a constant C such that

$$\mathscr{C}^{\mathrm{F}}_{\mathrm{eff}}(A \leftrightarrow B\,;\,G) \leq C\mathscr{C}^{\mathrm{F}}_{\mathrm{eff}}\big(\phi(A) \leftrightarrow \phi(B)\,;\,G'\big) \tag{8.3.8}$$

for each two sets of vertices A and B in G.

Let $\rho \in V$ and $\xi \in \partial\mathbb{D}$ be fixed, and let

$$A_\varepsilon(\xi) := \{v \in V : |z \circ \phi(v) - \xi| \le \varepsilon\}.$$

In [17, Corollary 5.2], it is proven that

$$\lim_{\varepsilon\to 0} \mathscr{C}^{\mathrm{F}}_{\mathrm{eff}}\left(\phi(\rho) \leftrightarrow \phi(A_\varepsilon(\xi))\,;\,T'\right) = 0.$$

Applying (8.3.8) and Rayleigh monotonicity, we have that

$$\begin{aligned}\mathscr{C}^{\mathrm{F}}_{\mathrm{eff}}\left(\rho \leftrightarrow A_\varepsilon(\xi)\,;\,G\right) &\le C\mathscr{C}^{\mathrm{F}}_{\mathrm{eff}}\left(\phi(\rho) \leftrightarrow \phi(A_\varepsilon(\xi))\,;\,G'\right)\\ &\le C\mathscr{C}^{\mathrm{F}}_{\mathrm{eff}}\left(\phi(\rho) \leftrightarrow \phi(A_\varepsilon(\xi))\,;\,T'\right) \xrightarrow[\varepsilon\to 0]{} 0.\end{aligned} \tag{8.3.9}$$

The well-known inequality of [173, Exercise 2.34] then implies that

$$\begin{aligned}P_\rho(\text{hit } A_\varepsilon(\xi) \text{ before } V \setminus B_k(\rho)) &\le \frac{\mathscr{C}_{\mathrm{eff}}\left(\rho \leftrightarrow A_\varepsilon(\xi) \cap B_k(\rho); G_{k+1}\right)}{\mathscr{C}_{\mathrm{eff}}\left(\rho \leftrightarrow A_\varepsilon(\xi) \cup V \setminus B_k(\rho); G_{k+1}\right)}\\ &\le \frac{\mathscr{C}_{\mathrm{eff}}\left(\rho \leftrightarrow A_\varepsilon(\xi) \cap B_k(\rho); G_{k+1}\right)}{\mathscr{C}_{\mathrm{eff}}\left(\rho \leftrightarrow V \setminus B_k(\rho); G_{k+1}\right)}.\end{aligned}$$

Applying the exhaustion characterisation of the free effective conductance [173, §9.1], which states that

$$\mathscr{C}^{\mathrm{F}}_{\mathrm{eff}}(A \leftrightarrow B; G) = \lim_{k\to\infty} \mathscr{C}_{\mathrm{eff}}(A \leftrightarrow B; G_k),$$

we have that

$$\begin{aligned}P_\rho\left(\text{hit } A_\varepsilon(\xi)\right) &= \lim_{k\to\infty} P_\rho(\text{hit } A_\varepsilon(\xi) \text{ before } V \setminus B_k(\rho))\\ &\le \lim_{k\to\infty} \frac{\mathscr{C}_{\mathrm{eff}}\left(\rho \leftrightarrow A_\varepsilon(\xi) \cap B_k(\rho); G_{k+1}\right)}{\mathscr{C}_{\mathrm{eff}}\left(\rho \leftrightarrow V \setminus B_k(\rho); G_{k+1}\right)} = \frac{\mathscr{C}^{\mathrm{F}}_{\mathrm{eff}}\left(\rho \leftrightarrow A_\varepsilon(\xi); G\right)}{\mathscr{C}_{\mathrm{eff}}\left(\rho \to \infty; G\right)}.\end{aligned} \tag{8.3.10}$$

Combining (8.3.9) and (8.3.10) we deduce that

$$P_\rho\left(\lim_{n\to\infty} z \circ \phi(X_n) = \xi\right) \le \lim_{\varepsilon\to 0} P_\rho\left(\text{hit } A_\varepsilon(\xi)\right) = 0. \qquad \square$$

This concludes the part of Theorem 8.1.2 concerning the existence of an embedding.

Remark 8.3.5. The statement that $\mathbb{P}(\text{hit } A_\varepsilon(\xi))$ converges to zero as ε tends to zero for every $\xi \in \partial\mathbb{D}$ can also be used to deduce convergence of $z \circ \phi(X_n)$ without appealing to the results of [12]: Suppose for contradiction that with positive probability $z \circ \phi(X_n)$ does not converge, so that there exist two boundary points $\xi_1, \xi_2 \in \partial\mathbb{D}$ in the closure of $\{z \circ \phi(X_n) : n \ge 0\}$. It is not hard to see that in this case the closure of $\{z \circ \phi(X_n) : n \ge 0\}$ must contain one of the boundary intervals $[\xi_1, \xi_2]$ or $[\xi_2, \xi_1]$. However, if the closure of $\{z \circ \phi(X_n) : n \ge 0\}$ contains an interval of positive length with positive probability, then there exists a point $\xi \in \partial\mathbb{D}$ that is contained in the closure of $\{z \circ \phi(X_n) : n \geq 0\}$ with positive probability. This contradicts the convergence of $\mathbb{P}(\text{hit } A_\varepsilon(\xi))$ to zero.

Proof of Theorem 8.1.2, identification of the Poisson boundary. Let $\langle X_n \rangle_{n\geq 0}$ and $\langle Y_n \rangle_{n\geq 0}$ be independent random walks on G started at ρ, and for each $k \geq 0$ let $\langle Z^k_j \rangle_{j\geq 0}$ be a random walk on G started at X_k that is conditionally independent of $\langle X_n \rangle_{n\geq 0}$ and $\langle Y \rangle_{n\geq 0}$ given X_k. Let $\mathbb{P}$ denote the joint distribution of $\langle X_n \rangle_{n\geq 0}$, $\langle Y_n \rangle_{n\geq 0}$ and all of the random walks $\langle Z^k_m \rangle_{m\geq 0}$.

Suppose that z is an embedding of G' in $\mathbb{D}$ such that $z\circ\phi(X_n)$ converges to a point in $\partial\mathbb{D}$ almost surely and the law of the limit is non-atomic. Since the law of $\lim_{n\to\infty} z\circ\phi(X_n)$ is non-atomic, the random variables $\xi^+ = \lim_{n\to\infty} z\circ\phi(X_n)$ and $\xi^- = \lim_{n\to\infty} z\circ\phi(Y_n)$ are almost surely distinct. Let

$$Q_k := \mathbb{P}\left(\lim_{m\to\infty} z\circ\phi(Z^k_m) \in (\xi^-, \xi^+) \,\middle|\, \langle X_n \rangle_{n\geq 0}, \langle Y_n \rangle_{n\geq 0}\right).$$

Using the non-atomicity of the law of $\lim_{n\to\infty} z\circ\phi(X_n)$, the same argument as in the proof of Theorem 8.1.1 also shows that

$$\limsup_{k\to\infty} \min\{Q_k, 1-Q_k\} > 0 \text{ almost surely.} \tag{8.3.11}$$

We now come to a part of the proof that requires more substantial modification. Let $\langle v_i \rangle_{i\geq 0}$ be a path in G starting at ρ such that $z\circ\phi(v_i) \to \xi^+$.

Given a path $\langle u_i \rangle_{i\geq 0}$ in G, let $\Phi(\langle u_i \rangle_{i\geq 0})$ denote the path in G' formed by concatenating the paths $\Phi(u_i, u_{i+1})$ for $i \geq 0$. For each $k \geq 0$, let τ_k be the first time t such that the path $\Phi(Z^k_{t-1}, Z^k_t)$ intersects the union of the traces of the paths $\Phi(\langle v_i \rangle_{i\geq 0})$ and $\Phi(\langle Y_n \rangle_{n\geq 0})$. Since the union of the traces of the paths $\Phi(\langle v_i \rangle_{i\geq 0})$ and $\Phi(\langle Y_n \rangle_{n\geq 0})$ either contains $\phi(X_k)$ or disconnects $\phi(X_k)$ from at least one of the two intervals (ξ^-, ξ^+) or (ξ^+, ξ^-), we have that

$$\mathbb{P}\left(\tau_k < \infty \,\middle|\, \langle X_n \rangle_{n\geq 0}, \langle Y_n \rangle_{n\geq 0}\right) \geq \min\{Q_k, 1-Q_k\}$$

and so

$$\limsup_{k\to\infty} \mathbb{P}\left(\tau_k < \infty \,\middle|\, \langle X_n \rangle_{n\geq 0}, \langle Y_n \rangle_{n\geq 0}\right) > 0 \tag{8.3.12}$$

almost surely by (8.3.11).

On the event that τ_k is finite, by definition of Φ, there exists a vertex $u \in \{v_i : i \geq 0\} \cup \{Y_n : n \geq 0\}$ such that $d'(\phi(Z_\tau), \phi(u)) \leq 2\alpha + 2\beta$, and consequently $d(Z_\tau, u) \leq \alpha(2\alpha + 3\beta)$ since ϕ is a rough isometry. Since G has bounded degrees and edge conductances bounded above and below, $c(e)/c(u) \geq \delta$ for some $\delta > 0$. Thus, by the strong Markov property,

$$\mathbb{P}\left(\langle Z^k_m \rangle_{m\geq 0} \text{ hits } \{v_i : i \geq 0\} \cup \{Y_n : n \geq 0\} \,\middle|\, \langle X_n \rangle_{n\geq 0},\, \langle Y_n \rangle_{n\geq 0},\, \{\tau_k < \infty\}\right) \geq \delta^{\alpha(2\alpha+3\beta)} > 0 \tag{8.3.13}$$

for all $k \geq 0$. Combining (8.3.12) and (8.3.13) yields

$$\limsup_{k\to\infty} \mathbb{P}\left(\langle Z_m^k\rangle_{m\geq 0} \text{ hits } \{v_i : i \geq 0\} \cup \{Y_n : n \geq 0\} \,\middle|\, \langle X_n\rangle_{n\geq 0}, \langle Y_n\rangle_{n\geq 0}\right) > 0 \tag{8.3.14}$$

almost surely. The same argument as in the proof of Theorem 8.1.1 also implies that

$$\mathbb{P}\left(\langle Z_m^k\rangle_{m\geq 0} \text{ hits } \{Y_n : n \geq 0\} \,\middle|\, \langle X_n\rangle_{n\geq 0}, \langle Y_n\rangle_{n\geq 0}\right) \xrightarrow[k\to\infty]{a.s.} 0. \tag{8.3.15}$$

We conclude by combining (8.3.14) with (8.3.15) and applying Proposition 8.3.1 to $\psi = z \circ \phi : V \to \mathbb{D} \cup \partial\mathbb{D}$. □

8.4 Identification of the Martin boundary

In this section, we prove Theorem 8.1.5 and deduce Theorem 8.1.4. We begin by proving that the rectangle tiling of a bounded degree triangulation with edge conductances bounded above and below does not have any accumulations of rectangles other than at the boundary circle $\mathbb{R}/\eta\mathbb{Z} \times \{1\}$.

Proposition 8.4.1. *Let T be a transient, simple, proper plane triangulation with bounded local geometry. Then for every vertex v of T and every $\varepsilon > 0$, there exist at most finitely many vertices u of T such that the probability that a random walk started at u visits v is greater than ε.*

Proof. Let $\mathcal{C}$ be a circle packing of T in the unit disc $\mathbb{D}$, with associated embedding z of T. Benjamini and Schramm [47, Lemma 5.3] proved that for every $\varepsilon > 0$ and $\kappa > 0$, there exists $\delta > 0$ such that for any $v \in \tilde{V}$ with $|z(v)| \geq 1 - \delta$, the probability that a random walk from v ever visits a vertex u such that $|z(u)| \leq 1 - \kappa$ is at most ε. (Their proof is given for $c \equiv 1$ but extends immediately to this setting.) By setting $\kappa = 1 - |z(v)|$, it follows that for every $\varepsilon > 0$, there exists $\delta > 0$ such that for every vertex u with $|z(u)| \geq 1 - \delta$, the probability that a random walk started at u hits v is at most ε. The claim follows since $|z(u)| \geq 1 - \delta$ for all but finitely many vertices u of T. □

Corollary 8.4.2. *Let T be a transient proper plane triangulation with bounded local geometry, and let $\mathcal{S}_\rho$ be a rectangle tiling of T. Then for every $t \in [0, 1)$, the cylinder $\mathbb{R}/\eta\mathbb{Z} \times [0, t]$ intersects only finitely many rectangles $S(e) \in \mathcal{S}_\rho$.*

We also require the following simple geometric lemma.

Lemma 8.4.3. *Let T be a transient proper plane triangulation with bounded local geometry, and let $\mathcal{S}_\rho$ be the square tiling of T in the cylinder $\mathbb{R}/\eta\mathbb{Z} \times [0, 1]$. Then for every sequence of vertices $\langle v_n\rangle_{n\geq 0}$ such that $y(v_n) \to 1$ and $\theta(v_n) \to \theta_0$ for some (θ_0, y_0) as $n \to \infty$, there exists a path $\langle \gamma_n\rangle_{n\geq 0}$ containing $\{v_n : n \geq 0\}$ such that $y(\gamma_n) \to 1$ and $\theta(\gamma_n) \to \theta_0$ as $n \to \infty$.*

Proof. Define a sequence of vertices $\langle v'_n\rangle_{n\geq 0}$ as follows. If the interval $I(v_n)$ is non-degenerate (i.e. has positive length), let $v'_n = v_n$. Otherwise, let $\langle \eta_j\rangle_{j\geq 0}$ be a path from v_n to ρ in T, let $j(n)$ be the smallest $j > 0$ such that the interval $I(\eta_{j(n)})$ is non-degenerate, and set $v'_n = \eta_{j(n)}$. Observe that

$y(v'_n) = y(v_n)$, that the interval $I(v'_n)$ contains the singleton $I(v_n)$, and that the path $\tilde{\eta}^n = \langle \eta^n_j \rangle_{j=0}^{j(n)}$ from v_n to v'_n in T satisfies $(I(\tilde{\eta}_j), y(\tilde{\eta}_j)) = (I(v_n), y(v_n))$ for all $j < j(n)$.

Let $\theta'(v'_n) \in I(v_n)$ be chosen so that the line segment ℓ_n between the points $(\theta'(v'_n), y(v'_n))$ and $(\theta'(v'_{n+1}), y(v'_{n+1}))$ in the cylinder $\mathbb{R}/\eta\mathbb{Z} \times [0,1]$ does not intersect any degenerate rectangles of $\mathcal{S}_\rho$ or a corner of any rectangles of $\mathcal{S}_\rho$ (this is a.s. the case if $\theta'(v'_n)$ is chosen uniformly from $I(v_n)$ for each n). The line segment ℓ_n intersects some finite sequence of squares corresponding to non-degenerate rectangles of $\mathcal{S}_\rho$, which we denote $e^n_1, \ldots, e^n_{l(n)}$. Let e^n_i be oriented so that $y(e^{n+}_i) > y(e^{n-}_i)$ for every i and n. For each $1 \leq i \leq l(n) - 1$, either

1. the vertical sides of the rectangles $S(e^n_i)$ and $S(e^n_{i+1})$ have non-disjoint interiors, in which case e^n_i and e^n_{i+1} lie in the boundary of a common face of T, or
2. the horizontal sides of the rectangles $S(e^n_i)$ and $S(e^n_{i+1})$ have non-disjoint interiors, in which case e^n_i and e^n_{i+1} share a common endpoint.

In either case, since T is a triangulation, there exists an edge in T connecting e^{n+}_i to e^{n-}_{i+1} for each i and n. We define a path γ^n by alternatingly concatenating the edges e^n_i and the edges connecting e^{n+}_i to e^{n-}_{i+1} as i increases from 1 to $l(n)$. Define the path γ by concatenating all of the paths

$$\tilde{\eta}^n \circ \gamma^n \circ (-\tilde{\eta}^{n+1})$$

where $-\tilde{\eta}^{n+1}$ denotes the reversal of the path $\tilde{\eta}^{n+1}$.

Let M be an upper bound for the degrees of T and for the conductances and resistances of the edges of T. For each vertex v of T, there are at most M rectangles of $\mathcal{S}_\rho$ adjacent to $I(v)$ from above, so that at least one of these rectangles has width at least $\operatorname{length}(I(v))/M$. This rectangle must have height at least $\operatorname{length}(I(v))/M^2$. It follows that

$$\operatorname{length}(I(v)) \leq M^{-2} \big(1 - y(v)\big) \tag{8.4.1}$$

for all $v \in T$, and so for every edge e of T,

$$|y(e^-) - y(e^+)| \leq \frac{1}{c(e)} \operatorname{length}\Big(I(e^-)\Big) \leq M^{-3}\Big(1 - y(e^-)\Big) \tag{8.4.2}$$

By construction, every vertex w visited by the path $\tilde{\eta}^n \circ \gamma^n \circ (-\tilde{\eta}^{n+1})$ has an edge emanating from it such that the associated rectangle intersects the line segment ℓ_n, and consequently a neighbouring vertex w' such that $y(w') \geq \min\{y(v_n), y(v_{n+1})\}$. Applying (8.4.2), we deduce that

$$y(w) \geq (1 + M^{-3}) \min\{y(v_n), y(v_{n+1})\} - M^{-3}$$

for all vertices w visited by the path $\tilde{\eta}^n \circ \gamma^n \circ (-\tilde{\eta}^{n+1})$, and consequently $y(\gamma_k) \to 1$ as $k \to \infty$. The estimate (8.4.2) then implies that $\operatorname{length}(I(\gamma_k)) \to 0$ as $k \to \infty$. Since $I(w)$ intersects the projection to the boundary circle of the line segment ℓ_n for each vertex w visited by the path $\tilde{\eta}^n \circ \gamma^n \circ (-\tilde{\eta}^{n+1})$, we deduce that $\theta(\gamma_k) \to \theta_0$ as $k \to \infty$. $\square$

We also have the following similar lemma for circle packings.

Lemma 8.4.4. *Let T be a CP hyperbolic plane triangulation, and let $\mathcal{C}$ be a circle packing of T in $\mathbb{D}$ with associated embedding z. Then for every sequence $\langle v_n \rangle_{n\geq 0}$ such that $z(v_n)$ converges as $n \to \infty$, there exists a path $\langle \gamma_n \rangle_{n\geq 0}$ in T containing $\{v_n : n \geq 0\}$ such that $z(\gamma_n)$ converges as $n \to \infty$.*

Proof. The proof is similar to that of Lemma 8.4.3, and we provide only a sketch. Draw a straight line segment in $\mathbb{D}$ between the centres of the circles corresponding to each consecutive pair of vertices v_n and v_{n+1}. The set of circles intersected by the line segment contains a path γ^n in T from v_n to v_{n+1}. (If this line segment is not tangent to any of the circles of $\mathcal{C}$, then the set of circles intersected by the segment is exactly a path in T.) The path γ is defined by concatenating the paths γ^n. For every $\varepsilon > 0$, there are at most finitely many v for which the radius of the circle corresponding to v is greater than ε, since the sum of the squared radii of all the circles in the packing is at most 1. Thus, for large n, all the circles corresponding to vertices used by the path γ^n are small. The circles that are intersected by the line segment between $z(v_n)$ and $z(v_{n+1})$ therefore necessarily have centers close to ξ_0 for large n. We deduce that $z(\gamma_i) \to \xi_0$ as $i \to \infty$. □

Proof of Theorem 8.1.5. Let $\langle X_n \rangle_{n\geq 0}$ be a random walk on T. Our assumptions guarantee that $\theta(X_n)$ and $z(X_n)$ both converge almost surely as $n \to \infty$ and that the laws of these limits are both non-atomic and have support $\mathbb{R}/\eta\mathbb{Z}$ and $\partial\mathbb{D}$ respectively.

Suppose that γ is a path in T that visits each vertex at most finitely often. We claim that $\theta(\gamma_i)$ converges if and only if $z(\gamma_i)$ converges if and only if X_n almost surely does not hit γ_i infinitely often. We prove that $\theta(\gamma_i)$ converges if and only if X_n almost surely does not hit γ_i infinitely often. The proof for $z(\gamma_i)$ is similar.

If $\theta(\gamma_i)$ converges, then X_n almost surely does not hit γ_i infinitely often, since otherwise $\lim_{i\to\infty} \theta(\gamma_i)$ would be an atom in the law of $\lim_{n\to\infty} \theta(X_n)$. Conversely, if $\theta(\gamma_i)$ does not converge, then there exist at least two distinct points $\theta_1, \theta_2 \in \mathbb{R}/\eta\mathbb{Z}$ such that $\theta_1 \times \{1\}$ and $\theta_2 \times \{1\}$ are in the closure of $\{(\theta(\gamma_i), y(\gamma_i)) : n \geq 0\}$. Let $\langle Y_n \rangle_{n\geq 0}$ be a random walks started at ρ independent of $\langle X_n \rangle_{n\geq 0}$. Since the law of $\lim_{n\to\infty} \theta(X_n)$ has full support, we have with positive probability that $\lim_{n\to\infty} \theta(X_n) \in (\theta_1, \theta_2)$ and $\lim_{n\to\infty} \theta(Y_n) \in (\theta_2, \theta_1)$. On this event, the union of the traces $\{X_n : n \geq 0\} \cup \{Y_n : n \geq 0\}$ disconnects $\theta_1 \times \{1\}$ from $\theta_2 \times \{1\}$, and consequently the path $\langle \gamma_i \rangle_{n\geq 0}$ must hit $\{X_n : n \geq 0\} \cup \{Y_n : n \geq 0\}$ infinitely often. By symmetry, there is a positive probability that $\langle X_n \rangle_{n\geq 0}$ hits the trace $\{\gamma_i : n \geq 0\}$ infinitely often as claimed.

We deduce that for every path γ in T that visits each vertex of T at most finitely often, $\theta(\gamma_i)$ converges if and only if $z(\gamma_i)$ converges. It follows from this and Lemmas 8.4.3 and 8.4.4 that for any sequence of vertices $\langle v_i \rangle_{i\geq 0}$ in T that includes each vertex of T at most finitely often, the sequence $\theta(v_i)$ converges if and only if $z(v_i)$ converges, and hence that the map $\xi \mapsto \theta(\xi) = \lim_{z(v)\to\xi} \theta(v)$ is well defined. To see that this map is a homeomorphism, suppose that ξ_n is a sequence of points in $\partial\mathbb{D}$ converging to ξ. For every n, there exists a vertex $v_n \in V$ such that $|\xi_n - z(v_n)| \leq 1/n$ and

$|\theta(\xi_n) - \theta(v_n)| \leq 1/n$. Thus, $z(v_n) \to \xi$ and we have

$$|\theta(\xi) - \theta(\xi_n)| \leq |\theta(\xi) - \theta(v_n)| + |\theta(v_n) - \theta(\xi_n)| \xrightarrow[n\to\infty]{} 0.$$

The proof of the continuity of the inverse is similar. $\square$

Proof of Theorem 8.1.4. By [17, Theorem 1.2], a sequence of vertices $\langle v_i \rangle_{i\geq 0}$ converges to a point in the Martin boundary of T if and only if $z(v_i)$ converges to a point in $\partial\mathbb{D}$, and the map

$$\xi \mapsto M_\xi := \lim_{\imath\to\infty} M_{v_i} \text{ where } z(v_i) \to \xi$$

is a homeomorphism from $\partial\mathbb{D}$ to $\partial\mathcal{M}(T)$. Combining this with Theorem 8.1.5 completes the proof. $\square$

Chapter 9

Hyperbolic and parabolic unimodular random maps

Summary. We show that for infinite planar unimodular random rooted maps, many global geometric and probabilistic properties are equivalent, and are determined by a natural, local notion of average curvature. This dichotomy includes properties relating to amenability, conformal geometry, random walks, uniform and minimal spanning forests, and Bernoulli bond percolation. We also prove that every simply connected unimodular random rooted map is sofic, that is, a Benjamini-Schramm limit of finite maps.

9.1 Introduction

In the classical theory of Riemann surfaces, the Uniformization Theorem states that every simply connected, non-compact Riemann surface is conformally equivalent to either the plane or the disc, which are inequivalent to each other by Liouville's theorem. The dichotomy provided by this theorem manifests itself in several different ways, relating to analytic, geometric and probabilistic properties of surfaces. In particular, if S is a simply connected, non-compact Riemann surface, then either

> S **is parabolic:** it is conformally equivalent to the plane, admits a compatible Riemannian metric of constant curvature 0, does not admit non-constant bounded harmonic functions, and is recurrent for Brownian motion,

or else

> S **is hyperbolic:** it is conformally equivalent to the disc, admits a compatible Riemannian metric of constant curvature -1, admits non-constant bounded harmonic functions, and is transient for Brownian motion.

In the 1990's, a discrete counterpart to this dichotomy began to develop in the setting of bounded degree planar graphs [35, 47, 48, 121, 122]. A milestone in this theory was the work of He and Schramm [121, 122], who studied *circle packings* of infinite triangulations of the plane. They proved that every infinite triangulation of the plane can be circle packed, either in the unit disc or in the plane, but not both. A triangulation is called **CP hyperbolic** or **CP parabolic** accordingly. He and Schramm also connected the circle packing type to isoperimetric and probabilistic properties of the triangulation, showing in particular that, in the bounded degree case, CP parabolicity is equivalent to the recurrence of simple random walk. Later, Benjamini and Schramm [47, 48] provided an analytic aspect to this dichotomy, showing that every bounded degree, infinite planar graph admits non-constant bounded harmonic functions if and only if it is transient for simple random walk, and that in this case the graph also admits non-constant bounded harmonic functions of finite Dirichlet energy. Most of this theory fails without the assumption of bounded degrees, as one can easily construct pathological counterexamples to the theorems above.

The goal of this paper is to develop a similar theory for *unimodular random rooted maps*, without the assumption of bounded degree. In our earlier work [21], we studied circle packings of, and random walks on, random plane triangulations of unbounded degree. In this paper, we study many further properties of unimodular random planar maps, which we do not assume to be triangulations. Our main result may be stated informally as follows; see Theorem 9.1.1 for a complete and precise statement.

Theorem (The Dichotomy Theorem)**.** *Every infinite, planar, unimodular random rooted planar map is either hyperbolic or parabolic. The map is hyperbolic if and only if its average curvature is negative and is parabolic if and only if its average curvature is zero. The type of a unimodular random rooted map determines many of its properties.*

The many properties we show to be determined by the type of the map are far-reaching, relating to aspects of the map including amenability, random walks, harmonic functions, spanning forests, Bernoulli bond percolation, and the conformal type of associated Riemann surfaces. The seeds of such a dichotomy were already apparent in [21], in which we proved that a unimodular random rooted plane triangulation is CP parabolic almost surely if and only if the expected degree of the root is six (which is equivalent to the average curvature being zero), if and only if the triangulation is **invariantly amenable** – a notion of amenability due to Aldous and Lyons [7] that is particularly suitable to unimodular random rooted graphs. A notable property that is *not* a part of Theorem 9.1.1 is recurrence of the random walk: while every hyperbolic unimodular random rooted map is transient, not every parabolic unimodular random rooted map is recurrent. (Theorem 9.1.1 can be combined with the work of Gurel-Gurevich and the third author [107] to deduce that a parabolic unimodular random planar map is recurrent under the additional assumption that the degree of the root has an *exponential tail*.)

A **map** is a *proper* (i.e., locally finite) embedding of a graph into an oriented surface viewed up to orientation preserving homeomorphisms of the surface. (Other definitions extend to non-orientable maps, but we shall not be concerned with those here.) A **rooted map** is a map together with a distinguished root vertex. The map is called **planar** if the surface is homeomorphic to an open subset of the sphere, and is **simply connected** if the surface is homeomorphic to the sphere or the plane. (In particular, every simply connected map is planar.) A random rooted map is said to be **unimodular** if it satisfies the **mass-transport principle**, which can be interpreted as meaning that 'every vertex of the map is equally likely to be the root'. See Section 10.2 for precise definitions of each of these terms.

The **curvature** of a map is a local geometric property, closely related to the Gaussian curvature of manifolds that may be constructed from the map; see Section 9.4 for a precise definition. For one natural manifold constructed from the map by gluing together regular polygons, the curvature at each vertex v is

$$\kappa(v) = 2\pi - \sum_{f \perp v} \frac{\deg(f) - 2}{\deg(f)}\pi,$$

where the sum is taken over faces of the map incident to v, and a face is counted with multiplicity if more than one of the corners of the face are located at v. We can define the average curvature of (M, ρ) to be the expectation $\mathbb{E}[\kappa(\rho)]$. Theorem 9.4.9 states that the average curvature is a canonical quantity associated to the random map, in the sense that any unimodular way of associating a manifold to the map will result in the same average curvature. Observe that the average curvature of a unimodular random triangulation is equal to $(6 - \mathbb{E}[\deg(\rho)])\pi/3$, so that a unimodular triangulation has expected degree greater than six if and only if it has negative average curvature. This relates the dichotomy described in [21] to that of Theorem 9.1.1.

Classical examples of unimodular random maps are provided by Voronoi diagrams of stationary point processes [37] and (slightly modified) Galton-Watson trees [7, Example 1.1], as well as lattices in the Euclidean and hyperbolic planes, and arbitrary local limits of finite maps. Many local modifications of maps, such as taking Bernoulli percolation or uniform or minimal spanning trees, preserve unimodularity, giving rise to many additional examples. Unimodular random maps, most notably the uniform infinite planar triangulation (UIPT) [24] and quadrangulation (UIPQ) [73, 157], have also been studied in the context of 2-dimensional quantum gravity; see the survey [98] and references therein. More recently, hyperbolic variants of the UIPT have been constructed [23, 75]. Many of these examples do not have uniformly bounded degrees, so that the deterministic theory is not applicable to them.

Our results also have consequences for unimodular random maps that are not planar, which we develop in Section 9.7.

9.1.1 The Dichotomy Theorem

Since many of the notions tied together in the following theorem are well known, we first state the theorem, and defer detailed definitions to individual sections dealing with each of the properties.

Theorem 9.1.1 (The Dichotomy Theorem)**.** *Let (M, ρ) be an infinite, ergodic, unimodular random rooted planar map and suppose that $\mathbb{E}[\deg(\rho)] < \infty$. Then the average curvature of (M, ρ) is non-positive and the following are equivalent:*

1. *(M, ρ) has average curvature zero.*
2. *(M, ρ) is invariantly amenable.*
3. *Every bounded degree subgraph of M is amenable almost surely.*
4. *Every subtree of M is amenable almost surely.*
5. *Every bounded degree subgraph of M is recurrent almost surely.*
6. *Every subtree of M is recurrent almost surely.*

7. *(M,ρ) is a Benjamini-Schramm limit of finite planar maps.*
8. *(M,ρ) is a Benjamini-Schramm limit of a sequence $\langle M_n\rangle_{n\geq 0}$ of finite maps such that*

$$\frac{\mathit{genus}(M_n)}{\#\{\mathit{vertices\ of\ } M_n\}} \xrightarrow[n\to\infty]{} 0.$$

9. *The Riemann surface associated to M is conformally equivalent to either the plane $\mathbb{C}$ or the cylinder $\mathbb{C}/\mathbb{Z}$ almost surely.*
10. *M does not admit any non-constant bounded harmonic functions almost surely.*
11. *M does not admit any non-constant harmonic functions of finite Dirichlet energy almost surely.*
12. *The laws of the free and wired uniform spanning forests of M coincide almost surely.*
13. *The wired uniform spanning forest of M is connected almost surely.*
14. *Two independent random walks on M intersect infinitely often almost surely.*
15. *The laws of the free and wired minimal spanning forests of M coincide almost surely.*
16. *Bernoulli(p) bond percolation on M has at most one infinite connected component for every $p\in[0,1]$ almost surely (in particular, $p_c = p_u$).*
17. *M is vertex extremal length parabolic almost surely.*

In light of this theorem, we call a unimodular random rooted map (M,ρ) with $\mathbb{E}[\deg(\rho)]<\infty$ **parabolic** if its average curvature is zero (and, in the planar case, clauses (1)–(17) all hold), and **hyperbolic** if its average curvature is negative (and, in the planar case, the clauses all fail).

As one might guess from the structure of Theorem 9.1.1, the proof consists of many separate arguments for the different implications. Some of the implications are already present in the literature, and part of this paper is spent surveying the earlier works that form the individual implications between the long list of equivalent items in Theorem 9.1.1. For the sake of completeness we also include proofs of several standard technical results from the ergodic theory literature using probabilistic terminology.

Some of the implications in Theorem 9.1.1 hold in any graph. For example, (10) implies (11) for any graph, and (14) and (13) are always equivalent (see [44]). Other implications hold for any planar graph. For example (10) is equivalent to (14), see [40]. We do not provide a comprehensive list of the assumptions needed for each implication, but some of this information is encoded in Figure 9.1.

Most of the paper is dedicated to proving the theorem under the additional assumption that M is simply connected; the multiply-connected case is easier and is handled separately in Section 9.7. The logical structure of the proof in the simply connected case is summarized in Figures 9.1 and 9.2. Figure 9.1 also shows which implications were already known and which are proved in the present paper.

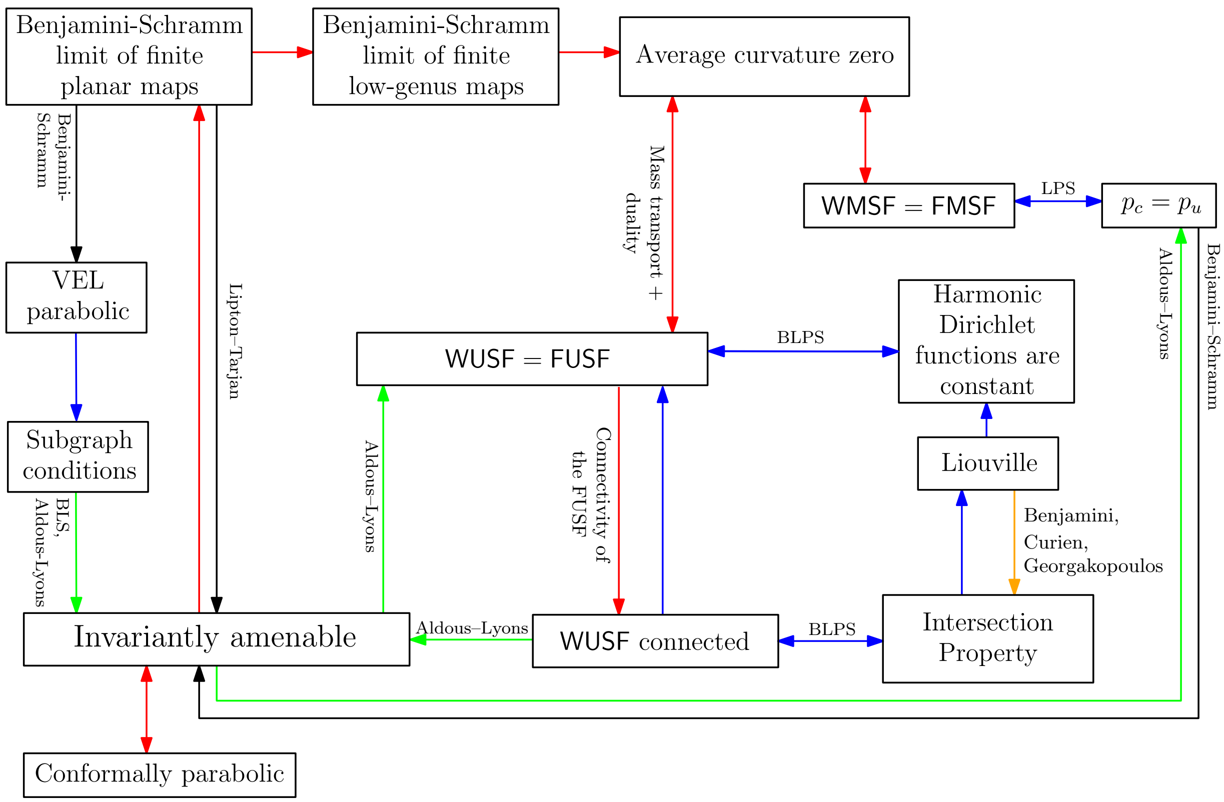

Figure 9.1: The logical structure for the proof of Theorem 9.1.1 in the simply connected case. Implications new to this paper are in red. Blue implications hold for arbitrary graphs; the orange implication holds for arbitrary planar graphs, and green implications hold for unimodular random rooted graphs even without planarity. A few implications between items that are known but not used in the proof are omitted.

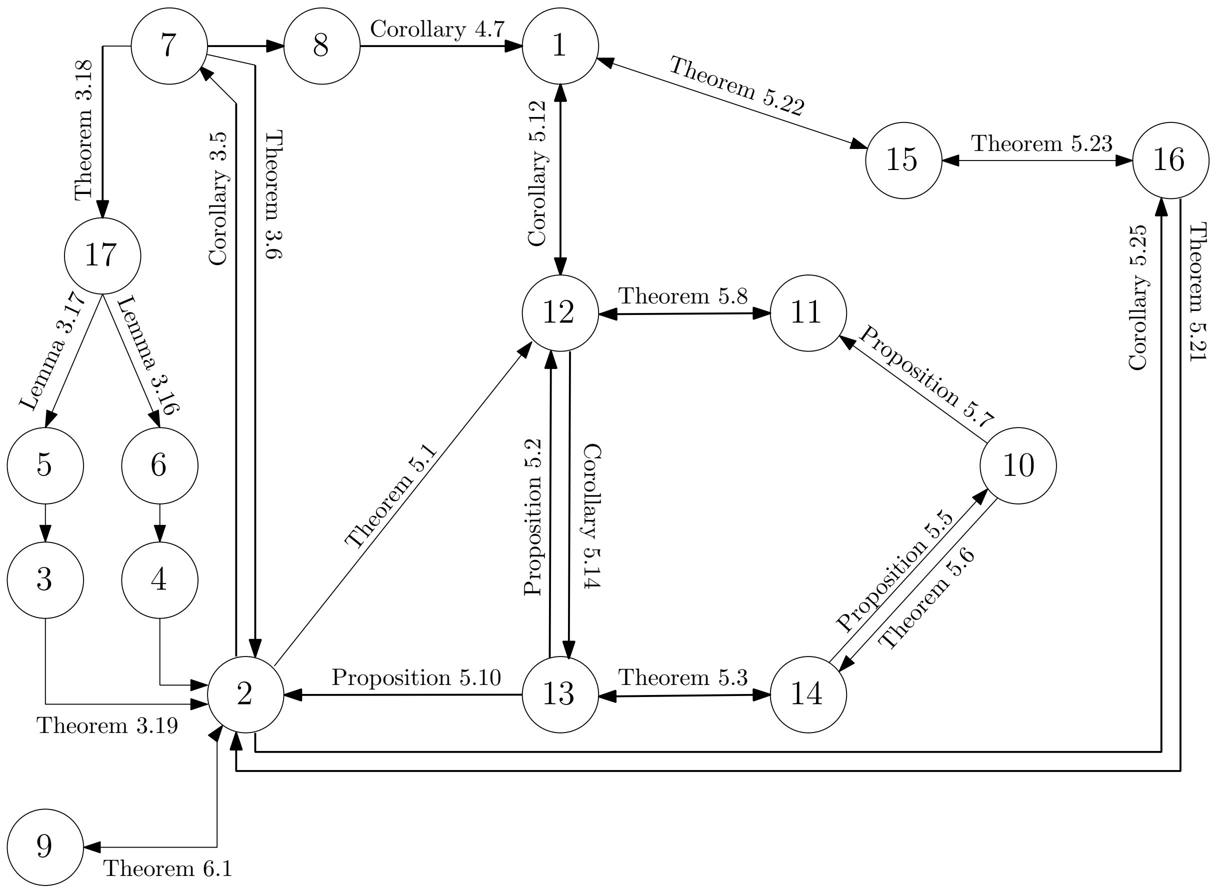

Figure 9.2: The numbers of the theorems, propositions, lemmas and corollaries forming the individual implications used to prove Theorem 9.1.1 in the simply connected case. Unlabelled implications are trivial.

9.1.2 Unimodular planar maps are sofic

Let $\langle G_n \rangle_{n \geq 0}$ be a sequence of (possibly random) finite graphs. We say that a random rooted graph (G, ρ) is the **Benjamini-Schramm limit** of the sequence $\langle G_n \rangle_{n \geq 0}$ if the random rooted graphs (G_n, ρ_n) converge in distribution to (G, ρ) with respect to the local topology on rooted graphs (see Section 9.2.2), where ρ_n is a uniform vertex of G_n. Benjamini-Schramm limits of finite maps and networks are defined similarly, except that the local topology takes into account the additional structure. When G_n is a uniformly chosen map of size n, this construction gives rise to the aforementioned UIPT and UIPQ.

It is easy to see that every (possibly random) finite graph with conditionally uniform root is unimodular. Moreover, unimodularity is preserved under distributional limits in the local topology. It follows that every Benjamini-Schramm limit of finite random graphs is unimodular. A random rooted graph that can be obtained in this way is called **sofic**. It is a major open problem to determine whether the converse holds, that is, whether every unimodular random rooted graph is sofic [7, Section 10]. The next theorem answers this question positively for simply connected unimodular random maps.

Theorem 9.1.2. *Every simply connected unimodular random rooted map is sofic.*

The proof of Theorem 9.1.2 relies on the corresponding result for trees, which is due to Bowen [56], Elek [85], and Benjamini, Lyons and Schramm [46]. As was observed by Elek and Lippner [86], it follows that *treeable* unimodular graphs (that is, unimodular graphs that exhibit an invariant spanning tree) are sofic, and so the key new step is proving the connectivity of the free uniform spanning forest.

Note that the finite maps converging to a given infinite unimodular random rooted map (M, ρ) need not be planar. Indeed, Theorem 9.1.1 characterises the Benjamini-Schramm limits of finite *planar* maps exactly as the parabolic unimodular random rooted maps. Moreover, if $\langle M_n \rangle_{n \geq 0}$ is a sequence of finite maps converging to an infinite hyperbolic unimodular random rooted map, then the approximating maps M_n must have genus comparable to their number of vertices as n tends to infinity.

9.2 Unimodular maps

9.2.1 Maps

We provide here a brief background to the concept of maps, and refer the reader to [160, Chapter 1.3] for a comprehensive treatment. Let $G = (V, E)$ be a connected graph, which may contain self-loops and multiple edges. An **embedding** of a graph in a surface S is a drawing of the graph in the surface with non-crossing edges. Given an embedding, the connected components of the complement of the image of G are called **faces**. An embedding is said to be **proper** if the following conditions hold:

1. it is locally finite (every compact set in S intersects finitely many edges),
2. every face is homeomorphic to an open disc, and
3. for every face f, if we consider the oriented edges of G that have their right hand side incident to f, and consider the permutation that maps each such oriented edge to the oriented edge following it in the clockwise order around the face, then this permutation has a single orbit. (Note that an edge can have the same face both to its right and its left.)

For example, the complete graph on three vertices can be properly embedded in the sphere but not in the plane. If S is simply connected or compact, then any embedding that satisfies (1) and (2) must also satisfy (3). For this reason, the condition (3) is not included in many references that deal primarily with finite maps. An example of an embedding that satisfies (1) and (2) but not (3) is given by drawing $\mathbb{Z}$ along a straight line in an infinite cylinder.

We define a (locally finite) **map** M to be a connected, locally finite graph G together with an equivalence class of proper embeddings of G into oriented surfaces, where two embeddings are equivalent if there is an orientation preserving homeomorphism between the two surfaces that sends one embedding to the other. If M is a map with underlying graph G, we refer to any proper embedding of G that falls into the equivalence class of embeddings corresponding to M as an embedding of M. A map is said to be **planar** if it is embedded into a surface homeomorphic to an open subset of the sphere, and is said to be **simply connected** if it is embedded into a simply connected surface (which is necessarily homeomorphic to either the sphere or the plane).

Let M be a map with underlying graph G and let z be a proper embedding of M into a surface S. If every face of M has finite degree, the **dual map** of M, denoted $M^\dagger$ is defined as follows. The underlying graph of $M^\dagger$, denoted $G^\dagger$, has the faces of M as vertices, and has an edge drawn between two faces of M for each edge in M that is incident to both of the faces. We define an embedding $z^\dagger$ of $G^\dagger$ into S by placing each vertex of $G^\dagger$ in the interior of the corresponding face of M and each edge of $G^\dagger$ so that it crosses the corresponding edge of G but no others. We define $M^\dagger$ to be the map with underlying graph G represented by the pair $(S, z^\dagger)$: Although the embedding $z^\dagger$ is not uniquely defined, every choice of $z^\dagger$ defines the same map. The construction gives a canonical bijection between edges of G and edges of $G^\dagger$. We write $e^\dagger$ for the edge of $G^\dagger$ corresponding to e. If e is an oriented edge, we let $e^\dagger$ be oriented so that it crosses e from right to left as viewed from the orientation of e.

Despite their topological definitions, maps and their duals can in fact be defined entirely combinatorially. Given any graph, we consider each edge as two oriented edges in opposite directions. We write $E^\rightarrow = E^\rightarrow(G)$ for the set of oriented edges of a graph. For each directed edge e, we have a head e^+ and tail e^-, and write $-e$ for the reversal of e. Given a map M and a vertex v of M, let $\sigma_v = \sigma_v(M)$ be the cyclic permutation of the set $\{e \in E^\rightarrow : e^- = v\}$ of oriented edges emanating from v corresponding to counter-clockwise rotation in S. This procedure defines a bijection between maps and graphs labelled by cyclic permutations.

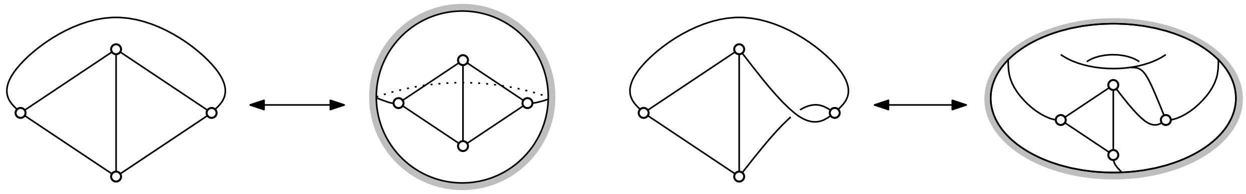

Figure 9.3: Different maps with the same underlying graph. The two maps both have K_4 as their underlying graph, but the left is a sphere map while the right is a torus map. Both maps are represented both abstractly as a graph together with a cyclic permutation of the edges emanating from each vertex (left) and as a graph embedded in a surface (right).

Theorem 9.2.1 ([160]). *Given a connected, locally finite graph G and a collection of cyclic permutations σ_v of the sets $\{e \in E^{\rightarrow} : e^- = v\}$, there exists a unique map M with underlying graph G such that $\sigma(M) = \sigma$.*

In light of Theorem 9.2.1, we identify a map M with the pair (G, σ). Given such a combinatorial specification of a map M as a pair (G, σ), we may form an embedding of the map into a surface $S(M)$ by gluing topological polygons according to the combinatorics of the map (see Figures 9.3 and 9.4). The faces of a map $M = (G, \sigma)$ can be defined abstractly as orbits of the permutation $\sigma^\dagger : E^{\rightarrow} \to E^{\rightarrow}$ defined by $\sigma^\dagger(e) = \sigma^{-1}(-e)$ for each $e \in E^{\rightarrow}$. The dual $e^\dagger$ of a directed edge e is defined to have the orbit of e as its tail and the orbit of $-e$ as its head, so that we again have a bijection between directed edges of M and their duals. Using this bijection, we can consider $\sigma^\dagger$ to acts on dual edges, and the dual map $M^\dagger$ is then constructed abstractly as $M^\dagger = (G^\dagger, \sigma^\dagger)$. We define maps with infinite degree vertices, and duals of maps with infinite degree faces, directly through this abstract formalism.

If e is an oriented edge in a map, we write e^ℓ for the face of M to the left of e and e^r for the face of M to the right of e, so that $e^\ell = (e^\dagger)^-$ and $e^r = (e^\dagger)^+$. Given a map M, we write $f \perp v$ if the face f is incident to the vertex v, that is, if there exists an oriented edge e of M such that $e^- = v$ and $e^\ell = f$. When writing a sum of the form $\sum_{f \perp v}$, we use the convention that a face f is counted with multiplicity according to the number of oriented edges e of M such that $e^- = v$ and $e^\ell = f$. Similarly, when writing a sum of the form $\sum_{u \sim v}$, we count each vertex u with multiplicity according to the number of oriented edges e of M such that $e^- = v$ and $e^+ = u$. The degree of a face is defined to be the number of oriented edges with $e^r = f$, i.e., the degree of the face in the dual.

9.2.2 Unimodularity and the mass transport principle

As noted, a **rooted graph** (G, ρ) is a connected, locally finite (multi)graph $G = (V, E)$ together with a distinguished vertex ρ, called the **root**. A graph isomorphism $\phi : G \to G'$ is an isomorphism of rooted graphs if it maps the root to the root.

The **local topology** (see [50]) is the topology on the set $\mathcal{G}_\bullet$ of isomorphism classes of rooted graphs induced by the metric

$$d_{\mathrm{loc}}\left((G,\rho),(G',\rho')\right) = e^{-R},$$

where

$$R = R\left((G,\rho),(G',\rho')\right) = \sup\left\{R \geq 0 : B_R(G,\rho) \cong B_R(G',\rho')\right\},$$

i.e., the maximal radius such that the balls $B_R(G,\rho)$ and $B_R(G',\rho')$ are isomorphic as rooted graphs.

A **random rooted graph** is a random variable taking values in the space $\mathcal{G}_\bullet$ endowed with the local topology. Similarly, a **doubly-rooted graph** is a graph together with an ordered pair of distinguished (not necessarily distinct) vertices. Denote the space of isomorphism classes of doubly-rooted graphs equipped with this topology by $\mathcal{G}_{\bullet\bullet}$.

The spaces of rooted and doubly-rooted maps are defined similarly and are denoted $\mathcal{M}_\bullet$ and $\mathcal{M}_{\bullet\bullet}$ respectively. For this, an isomorphism ϕ of rooted maps is preserves the roots and the map structure, so that $\phi \circ \sigma = \sigma' \circ \phi$ for vertices inside the balls.

A further generalisation of these spaces will also be useful. A **marked graph** (referred to by Aldous and Lyons [7] as a network) is defined to be a locally finite, connected graph together with a function $m : E \cup V \to \mathbb{X}$ assigning each vertex and edge of G a **mark** in some Polish space $\mathbb{X}$, referred to as the **mark space**. The local topology on the set of isomorphism classes of rooted graphs with marks in $\mathbb{X}$ is the topology induced by the metric

$$d_{\mathrm{loc}}((G,\rho,m),(G',\rho',m')) = e^{-R}$$

where $R((G,\rho,m),(G',\rho',m'))$ is the largest R such that there exists an isomorphism of rooted graphs $\phi : (B_G(\rho,b),m,\rho) \to (B_{G'}(\rho',b),\rho')$ such that $d_{\mathbb{X}}(m'(\phi(x)),m(x)) \leq 1/R$ for every vertex or edge x of $B_G(\rho,n)$, and $d_{\mathbb{X}}$ is a metric compatible with the topology of $\mathbb{X}$. The space of isomorphism classes of rooted marked graphs with marks in $\mathbb{X}$ is denoted $\mathcal{G}_\bullet^{\mathbb{X}}$. The space of rooted marked maps is defined similarly.

It is possible to consider rooted maps and rooted marked maps as rooted marked graphs by encoding the permutations σ_v in the marks – in particular, this means that any statement that holds for all unimodular random rooted marked graphs also holds for all unimodular random rooted marked maps. See [7, Example 9.6].

A **mass transport** is a Borel function $f : \mathcal{G}_{\bullet\bullet} \to [0,\infty]$. A random rooted graph (G,ρ) is said to be **unimodular** if it satisfies the **Mass Transport Principle**: for every mass transport f,

$$\mathbb{E}\left[\sum_{v\in V} f(G,\rho,v)\right] = \mathbb{E}\left[\sum_{u\in V} f(G,u,\rho)\right]. \tag{MTP}$$

That is,

Expected mass out equals expected mass in.

A probability measure $\mathbb{P}$ on $\mathcal{G}_\bullet$ is said to be unimodular if a random rooted map with law $\mathbb{P}$ is unimodular. Unimodular probability measures on rooted maps, marked graphs and marked maps are defined similarly.

The Mass Transport Principle was first introduced by Häggström [110] to study dependent percolation on Cayley graphs. The formulation of the Mass Transport Principle presented here was suggested by Benjamini and Schramm [50] and developed systematically by Aldous and Lyons [7]. The unimodular probability measures on $\mathcal{G}_\bullet$ form a weakly closed, convex subset of the space of probability measures on $\mathcal{G}_\bullet$, so that weak limits of unimodular random graphs are unimodular. In particular, a weak limit of finite graphs with uniformly chosen roots is unimodular: such a limit of finite graphs is referred to as a **Benjamini-Schramm limit**. It is a major open problem to determine whether all unimodular random rooted graphs arise as Benjamini-Schramm limits of finite graphs [7, §10].

We will make frequent use of the fact that unimodularity is stable under most reasonable ways of modifying a graph locally without changing its vertex set. The following lemma is a formalization of this fact.

Lemma 9.2.2. *Let $\mathbb{X}_1$ and $\mathbb{X}_2$ be polish spaces. Let (G,ρ,m) be a unimodular random rooted $\mathbb{X}_1$-marked graph, and suppose that (G',ρ,m') is a random rooted marked graph with the same vertex set as G, such that, for every pair of vertices u,v in G, the conditional distribution of (G',u,v,m') given (G,ρ,m) coincides a.s. with some measurable function of the isomorphism class of (G,u,v,m). Then (G',ρ,m') is unimodular.*

Proof. Let $f:\mathcal{G}_{\bullet\bullet}^{\mathbb{X}_2}\to[0,\infty]$ be a mass transport. For each pair of vertices u,v in G, the conditional expectation

$$F(G,u,v,m)=\mathbb{E}\left[f(G',u,v,m')\mid (G,\rho,m)\right]$$

coincides a.s. with a measurable function of (G,u,v,m), and so is itself a mass transport. Since (G,ρ,m) is unimodular we deduce that

$$\begin{aligned}\mathbb{E}\sum_v f(G',\rho,v,m')&=\mathbb{E}\sum_v F(G,\rho,v,m)\\&=\mathbb{E}\sum_v F(G,v,\rho,m)=\mathbb{E}\sum_u f(G',u,\rho,m'),\end{aligned}\tag{9.2.1}$$

verifying that (G',ρ,m') satisfies the Mass Transport Principle. □

Later in the paper, we will often claim that various random rooted graphs obtained from unimodular random rooted graphs in various ways are unimodular. These claims can always be justified either as a direct consequence of Lemma 9.2.2, or otherwise by applying the Mass-Transport Principle to the conditional expectation as in eq. (9.2.1).

9.2.3 Reversibility

Recall that the **simple random walk** on a locally finite graph $G = (V, E)$ is the Markov process $\langle X_n \rangle_{n\geq 0}$ on the state space V with transition probabilities $p(u, v)$ defined to be the fraction of edges emanating from u that end in v. A probability measure $\mathbb{P}$ on $\mathcal{G}_\bullet$ that is supported on rooted graphs with more than one vertex is said to be **stationary** if, when (G, ρ) is a random rooted graph with law $\mathbb{P}$ and $\langle X_n \rangle_{n\geq 0}$ is a simple random walk on G started at the root, (G, ρ) and (G, X_n) have the same distribution for all n. The measure $\mathbb{P}$ is said to be **reversible** if furthermore

$$(G, \rho, X_n) \stackrel{d}{=} (G, X_n, \rho)$$

for all n. A random rooted graph is said to be reversible if its law is reversible. Reversible random rooted maps, marked graphs and marked maps are defined similarly. Every reversible random rooted graph is clearly stationary, but the converse need not hold in general [36, Examples 3.1 and 3.2].

Let $\mathcal{G}_\leftrightarrow$ be the space of isomorphism classes of connected, locally finite graphs equipped with a distinguished (not necessarily simple) bi-infinite path, which we endow with a natural variant of the local topology. If (G, ρ) is a reversible random rooted graph and $\langle X_n \rangle_{n\geq 0}$ and $\langle X_{-n} \rangle_{n\geq 0}$ are independent simple random walks started from ρ, then the sequence

$$\langle (G, \langle X_{n+k} \rangle_{n\in\mathbb{Z}}) \rangle_{k\in\mathbb{Z}}$$

of $\mathcal{G}_\leftrightarrow$-valued random variables is stationary.

The following correspondence between unimodular and reversible random rooted graphs is implicit in [7] and was proven explicitly in [38]. Similar correspondences also hold between unimodular and reversible random rooted maps, marked graphs and marked maps.

Proposition 9.2.3 ([7, 38])**.** *Let $\mathbb{P}$ be a unimodular probability measure on $\mathcal{G}_\bullet$ that is supported on rooted graphs with more than one vertex and has $\mathbb{P}[\deg(\rho)] < \infty$, and let $\mathbb{P}_{\text{rev}}$ denote the $\deg(\rho)$-biasing of $\mathbb{P}$, defined by*

$$\mathbb{P}_{\text{rev}}((G, \rho) \in \mathscr{A}) := \frac{\mathbb{P}[\deg(\rho)\mathbb{1}((G, \rho) \in \mathscr{A})]}{\mathbb{P}[\deg(\rho)]}$$

for every Borel set $\mathscr{A} \subseteq \mathcal{G}_\bullet$. Then $\mathbb{P}_{\text{rev}}$ is reversible. Conversely, if $\mathbb{P}$ is a reversible probability measure on $\mathcal{G}_\bullet$, then biasing $\mathbb{P}$ by $\deg(\rho)^{-1}$ yields a unimodular probability measure on $\mathcal{G}_\bullet$.

9.2.4 Ergodicity

A probability measure $\mathbb{P}$ on $\mathcal{G}_\bullet$ is said to be **ergodic** if $\mathbb{P}(\mathscr{A}) \in \{0, 1\}$ for every event $\mathscr{A} \subseteq \mathcal{G}_\bullet$ that is invariant to changing the root in the sense that

$$(G, \rho) \in \mathscr{A} \iff (G, v) \in \mathscr{A} \text{ for all } v \in V.$$

A random rooted graph is said to be ergodic if its law is ergodic. Aldous and Lyons [7, §4] proved that the ergodic unimodular probability measures on $\mathcal{G}_\bullet$ are exactly the extreme points of the weakly closed, convex set of unimodular probability measures on. It follows by Choquet theory that every unimodular random rooted graph is a mixture of ergodic unimodular random rooted graphs, meaning that the graph may be sampled by first sampling a random ergodic unimodular probability measure on $\mathcal{G}_\bullet$ from some distribution, and then sampling from this randomly chosen measure. In particular, to prove almost sure statements about unimodular random rooted graphs, it suffices to consider the ergodic case. Analogous statements hold for unimodular random rooted maps, marked graphs, and marked maps.

9.2.5 Duality

Let $\mathbb{P}$ be a unimodular probability measure on $\mathcal{M}_\bullet$ such that $M^\dagger$ is locally finite $\mathbb{P}$-a.s., and let (M,ρ) be a random rooted map with law $\mathbb{P}_{\mathrm{rev}}$. Conditional on (M,ρ), let η be an oriented edge of M sampled uniformly at random from the set $E^{\rightarrow}_\rho$, and let $\rho^\dagger = \eta^r$. We define $\mathbb{P}^\dagger_{\mathrm{rev}}$ to the law of the random rooted map $(M^\dagger,\rho^\dagger)$.

Proposition 9.2.4 (Aldous-Lyons [7, Example 9.6])**.** *Let $\mathbb{P}$ be a unimodular probability measure on $\mathcal{M}_\bullet$ such that $M^\dagger$ is locally finite $\mathbb{P}$-a.s. and $\mathbb{P}[\deg(\rho)] < \infty$. Then $\mathbb{P}^\dagger_{\mathrm{rev}}$ is a reversible probability measure on $\mathcal{M}_\bullet$.*

Intuitively, under the measure $\mathbb{P}_{\mathrm{rev}}$, the edge η is 'uniformly distributed' among the edges of M. Since the map sending each edge to its dual is a bijection, it should follow that $\eta^\dagger$ is uniformly distributed among the edges of $M^\dagger$, making the measure $\mathbb{P}^\dagger_{\mathrm{rev}}$ reversible also.

We define $\mathbb{P}^\dagger$ to be the unimodular measure on $\mathcal{M}_\bullet$ obtained as the $\deg(\rho)^{-1}$-biasing of $\mathbb{P}^\dagger_{\mathrm{rev}}$. We refer to $\mathbb{P}^\dagger$ as the **dual** of $\mathbb{P}$ and say that $\mathbb{P}$ is **self-dual** if $\mathbb{P}^\dagger = \mathbb{P}$. We write $\mathbb{E}_{\mathrm{rev}}$, $\mathbb{E}^\dagger_{\mathrm{rev}}$, and $\mathbb{E}^\dagger$ for the associated expectation operators. We may express $\mathbb{P}^\dagger$ directly in terms of $\mathbb{P}$ by

$$\mathbb{P}^\dagger\big((M^\dagger,\rho^\dagger)\in\mathscr{A}\big) = \mathbb{E}\Big[\sum_{f\perp\rho}\deg(f)^{-1}\Big]^{-1}\mathbb{E}\Big[\sum_{f\perp\rho}\deg(f)^{-1}\mathbb{1}\Big((M^\dagger,f)\in\mathscr{A}\Big)\Big] \tag{9.2.2}$$

for every Borel set $\mathscr{A}\subseteq\mathcal{M}_\bullet$. In particular, we can calculate

$$\begin{aligned}\mathbb{E}^\dagger[\deg(\rho^\dagger)] &= \mathbb{E}\Big[\sum_{f\perp\rho}\deg(f)^{-1}\Big]^{-1}\mathbb{E}\Big[\sum_{f\perp\rho}\deg(f)^{-1}\deg(f)\Big]\\ &= \mathbb{E}\Big[\sum_{f\perp\rho}\deg(f)^{-1}\Big]^{-1}\mathbb{E}[\deg(\rho)] < \infty.\end{aligned} \tag{9.2.3}$$

The factor $\mathbb{E}[\sum_{f\perp\rho}\deg(f)^{-1}]^{-1}$ may be thought of as the ratio of the number of vertices of M to the number of faces of M.

Example 9.2.5. Let M be a finite map and let $\mathbb{P}$ be the law of the unimodular random rooted map (M,ρ) where ρ is a vertex of M chosen uniformly at random. Then $\mathbb{P}^\dagger$ is the law of the unimodular random rooted map $(M^\dagger, \rho^\dagger)$ obtained by rooting the dual $M^\dagger$ of M at a uniformly chosen face $\rho^\dagger$ of M.

9.3 Percolations and invariant amenability

9.3.1 Markings and percolations

Let (G, ρ) be a unimodular random rooted graph with law $\mathbb{P}$. Given a random rooted graph (G, ρ) and a polish space $\mathbb{X}$, an **$\mathbb{X}$-marking** of (G, ρ) is a random assignment $m : E \cup V \to \mathbb{X}$ of marks to the edges and vertices of (G, ρ) (possibly defined on some larger probability space) such that the random rooted marked graph (G, ρ, m) is unimodular. A **percolation** on (G, ρ) is a $\{0, 1\}$-marking of (G, ρ), which we think of as a random subgraph of G consisting of the **open** edges satisfying $\omega(e) = 1$, and open vertices satisfying $\omega(v) = 1$. We assume without loss of generality that if an edge is open then so are both of its endpoints. We call ω **connected** if the subgraph of ω spanned by the vertices $\{v \in V : \omega(v) = 1\}$ is connected. We call ω a **bond percolation** if $\omega(v) = 1$ for every vertex $v \in V$ almost surely. The **cluster** $K_\omega(v)$ of ω at the vertex v is the connected component of v in ω. A percolation is said to be **finitary** if all of its clusters are finite almost surely.

Let us gather here the following simple and useful technical lemmas concerning markings and percolations.

Lemma 9.3.1. *Let (G, ρ, m) be a unimodular random marked graph. Let ω be a percolation on (G, ρ, m) and let $m|_{K_\omega(\rho)}$ denote the restriction of m to $K_\omega(\rho)$. Then the conditional law of $(K_\omega(\rho), \rho, m|_{K_\omega)\rho)})$ given that $\omega(\rho) = 1$ is also unimodular. In particular, if ω is finitary, then ρ is uniformly distributed on its cluster.*

Proof. Let f be a mass transport, and let g be the mass transport

$$g(G, u, v, m, \omega) := \mathbb{1}\big(\omega(u) = \omega(v) = 1,\, v \in K_\omega(u)\big)\, f(K_\omega(u), u, v, m|_{K_\omega(u)}).$$

Then, applying the Mass Transport principle for (G, ρ, ω), we have

$$\begin{aligned}\mathbb{E} \sum_{v \in K_\omega(\rho)} f(K_\omega(\rho), \rho, v, m|_{K_\omega(\rho)}) &= \mathbb{E} \sum_{v \in V(G)} g(G, \rho, v, m, \omega) = \mathbb{E} \sum_{v \in V(G)} g(G, v, \rho, m, \omega) \\ &= \mathbb{E} \sum_{v \in K_\omega(\rho)} f(K_\omega(\rho), v, \rho, m|_{K_\omega(\rho)}).\end{aligned}$$

□

Lemma 9.3.2 (Coupling markings I)**.** *Let (G, ρ) be a unimodular random rooted graph and let $\{m_i : i \in I\}$ be a set of markings of (G, ρ) indexed by a countable set I and with mark spaces $\{\mathbb{X}_i : i \in I\}$. Then there exists a $\prod_{i \in I} \mathbb{X}_i$-marking m of (G, ρ) such that $(G, \rho, \pi_i \circ m)$ and (G, ρ, m_i) have the same distribution for all $i \in I$, where π_i denotes the projection of $\prod_{i \in I} \mathbb{X}_i$ onto $\mathbb{X}_i$ for each $i \in I$.*

Proof. One such marking is given by first sampling (G,ρ) and then sampling the markings m_i independently from their conditional distributions given (G,ρ). The details of this construction are are omitted. $\square$

Lemma 9.3.3 (Coupling markings II)**.** *Let (G,ρ) be a unimodular random rooted graph, let ω be a connected percolation on (G,ρ), and let m be an $\mathbb{X}$-marking of the unimodular random rooted graph given by sampling $(K_\omega(\rho),\rho)$ conditional on $\omega(\rho)=1$. Then there exists a $\mathbb{X}$-marking $\hat{m}$ of (G,ρ) such that the laws of $(K_\omega(\rho),\rho,\hat{m}|_{K_\omega(\rho)})$ and $(K_\omega(\rho),\rho,m)$ coincide.*

Proof. One such marking is given as follows: First sample $(K_\omega(\rho),\rho)$ from its conditional distribution given $\omega(\rho)=1$. Then, independently, sample (G,ρ) and m independently conditional on $(K_\omega(\rho),\rho)$. Extend m to the vertices and edges of G not present in ω by setting m to be some constant $x_0 \in \mathbb{X}$ on those vertices. This yields the law of a random rooted network (G,ρ,ω,m) in which the root always satisfies $\omega(\rho)=1$. Now, for each vertex u of G, let $v(u)$ be chosen uniformly from the set of vertices of ω closest to u, independently from everything else. (In particular, $v(u)=u$ if $\omega(u)=1$.) Let ρ' be chosen uniformly from the set $\{u : v(u)=\rho\}$. Sampling the network (G,ρ',ω,m) biased by $|\{u : v(u)=\rho\}|$ yields the desired coupling. The details of this construction are omitted. $\square$

9.3.2 Amenability

Recall that the **(edge) Cheeger constant** of an infinite graph $G=(V,E)$ is defined to be

$$\mathbf{i}_E(G) = \inf\left\{\frac{|\partial_E W|}{|W|} : W \subset V \text{ finite}\right\}$$

where $\partial_E W$ denotes the set of edges with exactly one end in W. The graph is said to be **amenable** if its Cheeger constant is zero and **nonamenable** if it is positive.

Non-amenability is often too strong a condition to hold for random graphs. For example, it is easily seen that every infinite cluster of a Bernoulli-$(1-\varepsilon)$ percolation on a nonamenable Cayley graph is almost surely amenable, since the cluster will contain 'bad regions' that contain sets of small expansion. In light of this, Aldous and Lyons introduced the following weakened notion of non-amenability for unimodular random graphs. (Another way to overcome this issue is to use *anchored expansion*, which is less relevant in our setting. See [7] for a comparison of the two notions.) The **invariant Cheeger constant** of an ergodic unimodular random rooted graph (G,ρ) is defined to be

$$\mathbf{i}^{\mathrm{inv}}(G,\rho) = \inf\left\{\mathbb{E}\left[\frac{|\partial_E K_\omega(\rho)|}{|K_\omega(\rho)|}\right] : \omega \text{ a finitary percolation on } G\right\}. \tag{9.3.1}$$

(This is a slight abuse of notation: $\mathbf{i}^{\mathrm{inv}}(G,\rho)$ is really a functions of the *law* of (G,ρ).) An ergodic unimodular random rooted graph (G,ρ) is said to be **invariantly amenable** if $\mathbf{i}^{inv}(G,\rho)=0$ and **invariantly nonamenable** otherwise. More generally, we say that a unimodular random rooted graph is invariantly amenable if its ergodic decomposition is supported on invariantly amenable unimodular random rooted graphs, and invariantly nonamenable if its ergodic decomposition is supported on invariantly nonamenable unimodular random rooted graphs. This is a property of the law of (G,ρ) and not of an individual graph. Invariant amenability and non-amenability is defined similarly for ergodic unimodular random rooted maps, marked graphs and marked maps.

Let (G,ρ) be a unimodular random rooted graph and let ω be a finitary percolation on G. Let $\deg_\omega(\rho)$ denote the degree of ρ in ω and let

$$\alpha(G,\rho) = \sup\Big\{\mathbb{E}\big[\deg_\omega(\rho)\big] : \omega \text{ a finitary percolation on } G\Big\}. \tag{9.3.2}$$

An easy application of the mass transport principle [7, Lemma 8.2] shows that

$$\mathbb{E}[\deg_\omega(\rho)] = \mathbb{E}\left[\frac{\sum_{v\in K_\omega(\rho)} \deg_\omega(v)}{|K_\omega(\rho)|}\right]. \tag{9.3.3}$$

It follows that, if $\mathbb{E}[\deg(\rho)] < \infty$,

$$\mathbf{i}^{\mathrm{inv}}(G,\rho) = \mathbb{E}[\deg(\rho)] - \alpha(G,\rho). \tag{9.3.4}$$

Similarly, if (G,ρ) is an ergodic unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] = \infty$, then $\alpha(G,\rho) < \infty$ is a sufficient condition for $\mathbf{i}^{\mathrm{inv}}(G,\rho)$ to be positive.

9.3.3 Hyperfiniteness

A unimodular random rooted graph (G,ρ) is said to be **hyperfinite** if there exists a $\{0,1\}^{\mathbb{N}}$-marking $\langle\omega_i\rangle_{i\geq 1}$ of (G,ρ) such that each of the percolations ω_i is finitary, $\omega_i \subseteq \omega_{i+1}$ almost surely, and $\bigcup_{i\geq 1}\omega_i = G$ almost surely. We call such an $\langle\omega_i\rangle_{i\geq 0}$ a **finitary exhaustion** for (G,ρ). The following is standard in the measured equivalence relations literature and was noted to carry through to the unimodular random rooted graph setting by Aldous and Lyons [7].

Theorem 9.3.4 ([7, Theorem 8.5])**.** *Let (G,ρ) be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then (G,ρ) is invariantly amenable if and only if it hyperfinite.*

Since a complete proof is not provided in [7], we provide a proof below for completeness.

Proof. First suppose that (G,ρ) is hyperfinite, and let $\langle\omega_i\rangle_{i\geq 0}$ be a finitary exhaustion for (G,ρ). By the monotone convergence theorem, $\mathbb{E}[\deg_{\omega_n}(\rho)] \to \mathbb{E}[\deg(\rho)]$, so that

$$\alpha(G,\rho) \geq \limsup_{i\to\infty} \mathbb{E}[\deg_{\omega_i}(\rho)] = \mathbb{E}[\deg(\rho)]$$

and hence (G,ρ) is invariantly amenable. Suppose conversely that (G,ρ) is invariantly amenable. For each $i \geq 1$, there exists a finitary percolation ω_i on G such that $\mathbb{E}[\deg(\rho) - \deg_{\omega_i}(\rho)] \leq 2^{-i}$. By Lemma 9.3.2, there exists $\{0,1\}^{\mathbb{N}}$-marking $\langle\omega_i\rangle_{i\geq 1}$ of (G,ρ) such that $\mathbb{E}[\deg(\rho) - \deg_{\omega_i}(\rho)] \leq 2^{-i}$

for each $i \geq 1$. For each $i \geq 1$, let $\hat{\omega}_i = \bigcap_{j \geq i} \omega_j$. Clearly $\langle \hat{\omega}_i \rangle_{i \geq 1}$ is a $\{0,1\}^{\mathbb{N}}$-marking of (G, ρ), and $\hat{\omega}_i \subseteq \hat{\omega}_{i+1}$ for every $i \geq 1$. Furthermore, by construction,

$$\mathbb{E}[\deg(\rho) - \deg_{\hat{\omega}_i}(\rho)] \leq \sum_{j \geq i} \mathbb{E}[\deg(\rho) - \deg_{\omega_i}(\rho)] \leq 2^{-i+1}$$

and hence, by Borel-Cantelli, ρ is in the interior of its cluster $K_{\omega_i}(\rho)$ for all sufficiently large i almost surely. It follows by unimodularity that $\bigcup \omega_i = G$, so that $\langle \hat{\omega}_i \rangle_{i \geq 0}$ is a finitary exhaustion for (G, ρ). □

Remark 9.3.5. Invariantly amenable unimodular random rooted graphs are always hyperfinite whether or not $\mathbb{E}[\deg(\rho)] < \infty$. However, the graph obtained by replacing each edge of the canopy tree (i.e., the Benjamini-Schramm limit of the balls in a 3-regular tree) at height n by 2^n parallel edges is hyperfinite but nonamenable.

Corollary 9.3.6 (Theorem 9.1.1, (2) implies (7).). *Let (M, ρ) be a hyperfinite ergodic random rooted map. Then (M, ρ) is a Benjamini-Schramm limit of finite planar maps.*

Proof. Let $\langle \omega_i \rangle_{i \geq 1}$ be a finitary exhaustion for (M, ρ). Let M_i be the finite map with underlying graph $K_{\omega_i}(\rho)$ and map structure inherited from M. Then (M_i, ρ_i) is a finite unimodular random rooted map, and converges to (M, ρ) almost surely as $i \to \infty$. □

Theorem 9.3.7 (Theorem 9.1.1: (7) implies (2)). *Let (M, ρ) be an ergodic unimodular random planar map that is a Benjamini-Schramm limit of finite planar maps. Then (M, ρ) is hyperfinite.*

This proposition follows as a corollary to the Lipton-Tarjan planar separator theorem [167]; see [21] for details. Similarly, it is possible to deduce that item (8) of Theorem 9.1.1 implies item (2) as an application of the low-genus separator theorem of Gilbert, Hutchinson, and Tarjan [102].

9.3.4 Ends

Recall that an infinite connected graph $G = (V, E)$ is said to be k-**ended** if k is minimal such that for every finite set of vertices W in G, the graph induced by the complement $V \setminus W$ has at most k distinct infinite connected components. In particular, an infinite tree is one-ended if it does not contain a simple bi-infinite path, and is two-ended if it contains a unique bi-infinite path.

The following proposition, due primarily to Aldous and Lyons [7], connects the number of ends of a unimodular random rooted graph to invariant amenability.

Proposition 9.3.8 (Ends and Amenability [7, 21]). *Let (G, ρ) be an infinite ergodic unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then G has either one, two, or infinitely many ends almost surely. If G has infinitely many ends, then (G, ρ) is invariantly nonamenable. If G has two ends almost surely, then it is almost surely recurrent and (G, ρ) is invariantly amenable.*

Similarly, a connected topological space X is said to be k-ended if over all compact subsets K of X, the complement $X \setminus K$ has a maximum of k connected components that are not precompact.

Lemma 9.3.9. *Let M be a map. Then the underlying graph of M has at least as many ends as the associated surface $S = S(M)$.*

Proof. For each compact subset K of S, let W_K be the set of vertices of M that are adjacent to an edge intersecting K. Since every face of M in S is a topological disc, each non-precompact connected component of $S \setminus K$ contains infinitely many edges of M, and there are no connections in $M \setminus W_K$ between these components, so that $M \setminus W_K$ has at least as many infinite connected components as $S \setminus K$ has non-precompact connected components. □

Note however that the 3-regular tree is infinitely-ended while its associated surface is homeomorphic to the plane and hence one-ended.

Finally, let us note the following simple topological fact.

Lemma 9.3.10. *Let M be a map. Then the underlying graph of M has at least as many ends as M has faces of infinite degree.*

9.3.5 Unimodular couplings and soficity

Let (G_1, ρ_1) and (G_2, ρ_2) be unimodular random rooted graphs. A **unimodular coupling** of (G_1, ρ_1) and (G_2, ρ_2) is a unimodular random rooted $\{0,1\}^2$-marked graph $(G, \rho, \omega_1, \omega_2)$ such that ω_1 and ω_2 are both connected almost surely and the law of the subgraph (ω_i, ρ) conditioned on the event that $\omega_i(\rho) = 1$ is equal to the law of (G_i, ρ_i) for each i. We say that two unimodular random rooted graphs (G_1, ρ_1) and (G_2, ρ_2) are **coupling equivalent** if they admit a unimodular coupling. (It is not difficult to show that this is an equivalence relation, arguing along similar lines to Lemmas 9.3.2 and 9.3.3.) Unimodular couplings and coupling equivalence are defined similarly for unimodular random rooted maps, marked graphs and marked maps. Coupling equivalence is closely related to the notion of two graphings generating the same equivalence relation in the theory of measured equivalence relations (see [7, Example 9.9]).

Example 9.3.11. Let (G, ρ) be a unimodular random rooted graph and let ω be a bond percolation on G that is connected almost surely. Then (G, ρ) and (ω, ρ) are coupling equivalent.

Example 9.3.12. Let (M, ρ) be a unimodular random rooted map with locally finite dual. Then (M, ρ) is coupling equivalent to its unimodular dual: both (M, ρ) and its dual can be represented as percolations of the graph formed by combining M and its dual by adding a vertex wherever a primal edge crosses a dual edge, as well as keeping the original primal and dual edges (see [7, Example 9.6] for a similar construction).

The main use of unimodular coupling will be that hyperfiniteness and strong soficity are both preserved under unimodular coupling. The corresponding statements for measured equivalence relations are due to Elek and Lippner [86].

Proposition 9.3.13. *Let (G_1, ρ_1) and (G_2, ρ_2) be coupling equivalent unimodular random rooted graphs. Then (G_2, ρ_2) is hyperfinite if and only if (G_1, ρ_1) is.*

Proof. It suffices to prove that if (G,ρ) is a unimodular random rooted graph and ω is a percolation on (G,ρ) that is almost surely connected, then (G,ρ) is hyperfinite if and only if (ω,ρ) is hyperfinite. Suppose first that (G,ρ) is hyperfinite and let $\langle\tilde{\omega}_i\rangle_{i\geq 1}$ be a finitary exhaustion for G. By Lemma 9.3.2, we may couple ω and $\langle\omega_i\rangle_{i\geq 1}$ so that $(G,\rho,(\omega,\langle\tilde{\omega}_i\rangle_{i\geq 1}))$ is unimodular. Under such a coupling, $\langle\omega\cap\tilde{\omega}_i\rangle_{i\geq 1}$ is a finitary exhaustion for (ω,ρ), and consequently (ω,ρ) is hyperfinite. Suppose conversely that (ω,ρ) is hyperfinite, and let $\langle\tilde{\omega}_i\rangle_{i\geq 1}$ be a finitary exhaustion for (ω,ρ). By Lemma 9.3.3, we may assume that $(G,\rho,\omega,\langle\tilde{\omega}_i\rangle_{i\geq 1})$ is unimodular. For each vertex v of G, let $w(v)$ be a chosen uniformly from the set of vertices of ω minimising the graph distance to v in G. For each i, define a subgraph $\hat{\omega}_i$ of G by

$$\hat{\omega}_i(e)=1 \iff w(e^-) \text{ and } w(e^+) \text{ are connected in } \tilde{\omega}_i. \tag{9.3.5}$$

Then $\langle\hat{\omega}_i\rangle_{i\geq 1}$ is a finitary exhaustion for (G,ρ). □

Remark 9.3.14. It can be deduced from Proposition 9.3.13 and [7, Theorem 8.9] that a unimodular random rooted graph with finite expected degree is hyperfinite if and only if it is coupling equivalent to $\mathbb{Z}$.

One useful application of Proposition 9.3.13 is the following.

Lemma 9.3.15. *Let (M,ρ) be an ergodic unimodular random rooted map. The number of infinite degree faces is either* 0*,* 1*,* 2*, or* ∞*, and if it is* 1 *or* 2 *then (M,ρ) is hyperfinite.*

Proof. By applying Proposition 9.3.8 and Lemma 9.3.10, it suffices to prove that if M has a nonzero but finite number of infinite degree faces almost surely, then (M,ρ) is hyperfinite. Suppose that M has finitely many infinite degree faces almost surely. Conditional on (M,ρ), choose one such face, f, uniformly at random. If for every edge e of M incident to f there exists a vertex v of M such that e is contained in a finite component of $M\setminus\{v\}$, then inductively there exists a sequence of vertices $\langle v_n\rangle_{n\geq 1}$ of M such that v_n is contained in a finite connected component of $M\setminus\{v_{n+1}\}$ for every $n\geq 1$. We deduce in this case that M is hyperfinite by taking as a finitary exhaustion the sequence of random subgraphs $\langle\tilde{\omega}_m\rangle_{m\geq 1}$ induced by the standard monotone coupling of Bernoulli $1-1/m$ site percolations on M, all of the clusters of which are almost surely finite due to the existence of the cutpoints $\langle v_n\rangle_{n\geq 1}$.

Otherwise, define a percolation ω on (M,ρ) by setting an edge e of M to be open if and only if exactly one side of e is incident to f and there does not exist a vertex v of M such that e is contained in a finite connected component of $M\setminus\{v\}$, and setting a vertex v of M to be open if and only if it is incident to an open edge. Then ω is connected and is isomorphic to the bi-infinite line graph $\mathbb{Z}$. It follows that (G,ρ) is coupling equivalent to $(\mathbb{Z},0)$ and hence hyperfinite by Proposition 9.3.13 □

Remark 9.3.16. An alternate proof in the case that $\mathbb{E}[\deg(\rho)]<\infty$ is as follows. Suppose that M has finitely many infinite degree faces almost surely. Conditional on (M,ρ), choose one such face, f, uniformly at random. We will use the boundary of this face to construct a unimodular coupling between M and $\mathbb{Z}$. let $\langle e_i(f)\rangle_{i\in\mathbb{Z}}$ be the set of oriented edges of M with $e_i^l(f)=f$, enumerated so that $e_i(f)^+ = e_{i+1}^-(f)$ for all $i \in \mathbb{Z}$. (this enumeration is unique up to translations of the indices, and this choice will not matter). Let $G = (V, E)$ be the underlying graph of M, and consider the graph G' with vertex set $V' = V \cup \mathbb{Z}$ and edges given by the edges of G, the edges of $\mathbb{Z}$, and an edge connecting $v \in V$ to $i \in \mathbb{Z}$ if and only if $v = e_i(f)^-$. Let ρ' be chosen uniformly from the set $\{\rho\} \cup \{i : e_i(f)^- = \rho\}$. It is easy to verify that the graph (G', ρ') is unimodular when sampled biased by $1 + |\{i : e_i(f)^- = \rho\}|$, which has finite expectation by assumption. Thus, we have constructed a unimodular coupling between (G, ρ) and $\mathbb{Z}$, so that the claim follows by Proposition 9.3.13.

We call a sofic unimodular random rooted graph (G, ρ) **strongly sofic** if (G, ρ, m) is sofic for every marking m of (G, ρ). The main input to Theorem 9.1.2 is the following, which was proven for Cayley graphs of free groups by [56].

Theorem 9.3.17 (Bowen [56]; Elek [85]; Elek and Lippner [86]; Benjamini, Lyons and Schramm [46])**.** *Every unimodular random rooted tree is strongly sofic.*

The following theorem is an adaptation of a related theorem of Elek and Lippner [86] in the setting of group actions. Even in our setting, it is well-known to experts that treeable unimodular random graphs (i.e., unimodular random graphs admitting a unimodular random spanning tree) are sofic.

Theorem 9.3.18. *Let (G_1, ρ_1) and (G_2, ρ_2) be coupling equivalent unimodular random rooted graphs. Then (G_1, ρ_1) is strongly sofic if and only if (G_2, ρ_2) is strongly sofic.*

The proof can be summarised as follows: Suppose (G_1, ρ_1) is strongly sofic, and let m be a marking of (G_2, ρ_2). We can encode both the structure of G_2 and the marks m as a marking $\hat{m}$ of (G_1, ρ_1). The strong soficity of (G_1, ρ_1) allows us to approximate $(G_1, \rho_1, \hat{m})$ by a sequence of finite graphs. We then use this sequence to define an approximating sequence for (G_2, ρ_2, m). This argument takes some care to make rigorous.

Proof. Since (G_1, ρ_1) and (G_2, ρ_2) can both be considered as percolations on some unimodular random graph (G, ρ), it suffices to prove that if (G, ρ) is a unimodular random rooted graph and ω is an almost surely connected percolation on G, then (G, ρ) is strongly sofic if and only if the unimodular random rooted graph (H, ρ') obtained from (ω, ρ) by conditioning on $\omega(\rho) = 1$ is strongly sofic.

First suppose that (G, ρ) is strongly sofic and let m be a marking of (H, ρ'). By Lemma 9.3.3, there exists a marking m of G such that (G, ρ, ω, m) is unimodular and such that the law of (H, ρ', m) coincides with the law of (ω, ρ, m) conditional on $\omega(\rho) = 1$. Since (G, ρ) is strongly sofic, there exists a sequence of finite unimodular random marked graphs $(G_n, \rho_n, \omega_n, m_n)$ converging to (G, ρ, ω, m) in distribution. Let (H_n, ρ_n', m_n) be the unimodular random rooted graph obtained from (ω_n, ρ_n, m_n) by conditioning on $\omega(\rho) = 1$. The sequence (H_n, ρ_n', m_n) converges to (H, ρ', m) and, since m was arbitrary, (H, ρ') is strongly sofic.

Suppose conversely that (H, ρ') is strongly sofic and let m be an $\mathbb{X}$-marking of (G, ρ). By Lemma 9.3.2, we may assume that (G, ρ, ω, m) is unimodular. For each vertex u of G, let $v(u)$

be chosen uniformly from the set of vertices in ω that minimize the graph distance to u, and for each vertex v of ω let $U_v = \{v(u) = v\}$. Transporting mass 1 from u to $v(u)$ for every vertex u of G shows that $\mathbb{E}|U_\rho| < \infty$, and in particular U_v is finite for every vertex v of ω almost surley. Conditional on (G, ρ, ω, m), let $\{U(v) : v \in V\} \cup \{U(e) : e \in E\}$ be a collection i.i.d. uniform $[0, 1]$ random variables indexed by the vertices and edges of G, and for each vertex v of G such that $\omega(v) = 1$, define

$$\hat{m}(v) = \Big\{ \big(U(u), m(u), \{(U(e), m(e)) : e \text{ is an edge incident to } u \text{ in } G\}\big) : u \in U_v \Big\}.$$

If we denote the space of finite subsets of the metric space X by X^*, then $\hat{m}$ takes values in the metric space

$$\Big([0, 1] \times \mathbb{X} \times \big([0, 1] \times \mathbb{X}\big)^*\Big)^*.$$

By conditioning on $\omega(\rho) = 1$, we obtain a unimodular random rooted marked graph $(H, \rho, \hat{m})$. Since (H, ρ) is strongly sofic, there exists a sequence of finite unimodular random rooted marked graphs $(H_n, \rho_n, \hat{m}_n)$ converging in distribution to $(H, \rho, \hat{m})$, and we may assume that $\mathbb{E}|\hat{m}_n(\rho_n)| < \infty$ for each $n \geq 1$. For each $\varepsilon > 0$ and $R \in \mathbb{N}$, bias $(H_n, \rho_n, \hat{m}_n)$ by $|\hat{m}(\rho)|$ and construct a finite unimodular random rooted marked graph $(G_n^{R,\varepsilon}, \rho_n, m_n^{R,\varepsilon})$ from $(H_n, \rho_n', \hat{m}_n)$ as follows.

1. Let the vertex set of $G_n^{R,\varepsilon}$ be the union $\bigcup_{v \in H_n} \hat{m}(v)$, and let ρ_n be chosen uniformly from $\hat{m}(\rho_n')$. For each vertex u of $G_n^{R,\varepsilon}$, let $u = (u^1, u^2, u^3)$ be the three coordinates of u, and let $m(u) = u^2$.

2. Draw an edge between two vertices u_1 and u_2 of $G_n^{R,\varepsilon}$ if and only if

   (a) $u_1 \in \hat{m}_n(v_1)$ and $u_2 \in \hat{m}_n(v_2)$ for some vertices v_1 and v_2 of H_n that are at distance at most R in H_n, and

   (b) there exists a pair of points

   $$(t, x) \in u^3 \subset [0, 1] \times \mathbb{X} \text{ and } (s, y) \in u^3 \subset [0, 1] \times \mathbb{X}$$

   such that the distance between (t, x) and (s, y) is less than ε in the product metric. If there are multiple such pairs, we draw multiple edges as appropriate.

3. Let $\{Z_n(e)\}$ be a collection of i.i.d. Bernoulli-1/2 random variables indexed by the edges of $G_n^{R,\varepsilon}$. For each edge e of $G_n^{R,\varepsilon}$, let (t, x) and (s, y) be the matching pair of points in $[0, 1] \times \mathbb{X}$ that led us to draw e in step (2), and let $m_n^{R,\varepsilon}(e) = x$ if $Z_n(e) = 0$ and $m_n^{R,\varepsilon}(e) = y$ if $Z_n(e) = 1$.

This construction is continuous for the local topology, and hence for each fixed ε and R, the finite networks $(G_n^{R,\varepsilon}, \rho_n, m_n^{R,\varepsilon})$ converge to the network $(G^{R,\varepsilon}, \rho, m^{R,\varepsilon})$ defined by applying the same procedure to $(H, \rho', \hat{m})$. Taking $\epsilon \to 0$, we obtain the network (G^R, ρ, m^R) which consists of those edges of (G, ρ) whose endpoints are of distance at most R in ω. Finally, taking $R \to \infty$ we recover

(G,ρ,m). Thus, (G,ρ,m) is a weak limit of sofic unimodular random rooted marked graphs, and it follows that (G,ρ,m) is sofic. □

Proof of Theorem 9.1.2. Let (M,ρ) be a simply connected unimodular random rooted map, let (G,ρ) be the underlying graph of M, and let $\mathfrak{F}$ be a sample of FUSF_M. By Theorem 9.5.13, $\mathfrak{F}$ is connected almost surely, and so (G,ρ) is coupling equivalent to the unimodular random rooted tree $(\mathfrak{F},\rho)$. It follows from Theorems 9.3.17 and 9.3.18 that (G,ρ) is strongly sofic. By encoding the map (M,ρ) as a marking of (G,ρ) as in Section 9.2.2 and [7, Example 9.6], we conclude that the unimodular random rooted map (M,ρ) is also sofic. □

9.3.6 Vertex extremal length and recurrence of subgraphs

Let G be an infinite graph. For each vertex v of G, the **vertex extremal length** from v to infinity is defined to be

$$\mathrm{VEL}_G(v,\infty) = \sup_m \frac{\inf_{\gamma: v\to\infty} m(\gamma)^2}{\|m\|^2}, \tag{9.3.6}$$

where the supremum is over measures m on the vertex set of G such that $\|m\|^2 = \sum m(u)^2 < \infty$, and the infimum is over paths γ from v to ∞ in G. A connected graph is said to be **VEL parabolic** if $\mathrm{VEL}(v\to\infty)=\infty$ for some vertex v of G (and hence for every vertex), and **VEL hyperbolic** otherwise. The VEL type is easily seen to be monotone in the sense that subgraphs of VEL parabolic graphs are also VEL parabolic.

Lemma 9.3.19 (He and Schramm [121]: Theorem 9.1.1, (17) implies (5))**.** *Let G be a locally finite, connected graph. If G is VEL hyperbolic, then it is transient. If G has bouneded degrees then the converse also holds, so that G is transient if and only if it is VEL hyperbolic.*

Lemma 9.3.20 (Theorem 9.1.1, (17) implies (6))**.** *Let T be a tree. Then T is transient if and only if it is VEL hyperbolic.*

Proof. Let G be an infinite connected graph. Recall that the **effective resistance** from a v to infinity in a graph G is defined to be

$$\mathscr{R}_{\mathrm{eff}}(v\to\infty) = \sup_m \frac{\inf_{\gamma: v\to\infty} m(\gamma)^2}{\|m\|^2},$$

where the supremum is over measures m on the edge set of G such that $\|m\|^2 = \sum m(u)^2 < \infty$, and the infimum is over paths γ from v to ∞ in G. Recall also that G is transient if and only if the effective resistance from v to infinity is finite for some vertex (and hence every vertex) v of G.

Let T be a VEL parabolic tree, and let v be a vertex of T. For every $M\geq 1$, there exists be a measure m on the vertex set of T such that

$$\frac{\inf_{\gamma: v\to\infty} m(\gamma)^2}{\|m\|^2} \geq M.$$

For each edge e of T, let $u(e)$ be the endpoint of e farthest to v, and define a measure $\hat{m}$ on the edge set of T by setting $\hat{m}(e) = m(u(e))$ for every edge e of T. Then

$$\|\hat{m}\|^2 = \|m\|^2 - m(v)^2 \quad \text{and} \quad \hat{m}(\gamma) = m(\gamma) - m(v)$$

for every path γ from v to ∞. Since $\|m\|^2 \geq m(v)^2$, it follows that

$$\begin{aligned}\mathscr{R}_{\mathrm{eff}}(v \to \infty) &\geq \frac{\inf_{\gamma:v\to\infty} \hat{m}(\gamma)^2}{\|\hat{m}\|^2} \geq \inf_{\gamma:v\to\infty} \frac{\big(m(\gamma)-m(v)\big)^2}{\|m\|^2} \geq \inf_{\gamma:v\to\infty}\left(\frac{m(\gamma)^2}{\|m\|^2} - 2\sqrt{\frac{m(\gamma)^2}{\|m\|^2}}\right)\\ &= \inf_{\gamma:v\to\infty} \frac{m(\gamma)^2}{\|m\|^2} - 2\sqrt{\inf_{\gamma:v\to\infty}\frac{m(\gamma)^2}{\|m\|^2}} \geq M - 2\sqrt{M}.\end{aligned}$$

Since M was arbitrary, we deduce that T is recurrent. □

Vertex extremal length was introduced by He and Schramm [121] and is closely connected to circle packing. Benjamini and Schramm [50] used circle packing to prove the following remarkable theorem. See [21] for an alternative proof.

Theorem 9.3.21 (Benjamini-Schramm [50]: Theorem 9.1.1, (7) implies (17))**.** *Let (M, ρ) be a Benjamini-Schramm limit of finite planar maps. Then M is VEL parabolic almost surely. In particular, if M has bounded degrees, then it is almost surely recurrent for simple random walk.*

The converse is provided by the following theorem.

Theorem 9.3.22 (Benjamini, Lyons and Schramm [45]; Aldous and Lyons [7]: Theorem 9.1.1, (3) implies (2))**.** *Let (G, ρ) be an invariantly nonamenable unimodular random rooted graph. Then there exists a percolation ω on G such that every connected component of ω are nonamenable almost surely, and there exists a constant M such that* $\deg(v) \leq M$ *for every vertex v of G such that $\omega(v) = 1$. Furthermore, the percolation ω can be taken to be a forest.*

This theorem is also very useful for studying random walks on invariantly nonamenable unimodular random rooted graphs; see [21, Section 5.1]. See [21, Section 3.3.1] for a complete proof of the first part of the theorem (in which ω is not taken to be a forest).

9.4 Curvature

In this section, we introduce the average curvature of a unimodular random rooted map, and establish some of its basic properties. We begin by giving a combinatorial definition of the average curvature; in Section 9.4, we show that show that the average curvature is a canonical quantity associated to the random map, in the sense that any unimodular way of embedding the map in a Riemannian manifold (satisfying certain integrability conditions) will result in the same average curvature.

Recall that the internal angles of a regular k-gon are given by $(k-2)\pi/k$. We define the **angle sum** at a vertex v of a map M to be

$$\theta(v) = \theta_M(v) = \sum_{f \perp v} \frac{\deg(f) - 2}{\deg(f)} \pi$$

This definition extends to maps with infinite faces, with the convention that $(\infty - 2)/\infty = 1$. In the case that every face of M has degree at least 3, we interpret $\theta(v)$ as the total angle of the corners at v if we form M by gluing together regular polygons, where we consider the upper half-space $\{x + iy \in \mathbb{C} : y > 0\}$ with edges $\{[n, n+1] : n \in \mathbb{Z}\}$ to be a regular ∞-gon. Of course, M cannot necessarily be drawn in the plane with regular polygons, and the angle sum at a vertex of M need not be 2π. We define the **curvature** of M at the vertex v to be the angle sum deficit

$$\kappa(v) = \kappa_M(v) = 2\pi - \theta(v)$$

and define the **average curvature** of a unimodular random rooted map (M, ρ), denoted $\mathbb{K}(M, \rho)$, to be the expected curvature at the root

$$\mathbb{K}(M, \rho) = \mathbb{E}[\kappa(\rho)] = 2\pi - \mathbb{E}\left[\sum_{f \perp \rho} \frac{\deg(f) - 2}{\deg(f)} \pi\right].$$

Note that if $\mathbb{E}[\deg(\rho)]$ is finite then $\mathbb{K}(M, \rho)$ is also finite.

Example 9.4.1 (Finite Maps)**.** Let M be a finite map and let ρ be a vertex of M chosen uniformly at random. Then

$$\begin{aligned}
\mathbb{K}(M, \rho) = \frac{1}{|V|} \sum_{v \in V} \kappa(v) &= \frac{1}{|V|} \left(2\pi |V| - \sum_{v \in V} \sum_{f \perp v} \frac{\deg(f) - 2}{\deg(f)} \pi \right) \\
&= \frac{1}{|V|} \left(2\pi |V| - \sum_{f \in F} (\deg(f) - 2)\pi \right) \\
&= \frac{1}{|V|} 2\pi \big(|V| - |E| + |F| \big) = \frac{1}{|V|} 2\pi \big(2 - 2\,\mathrm{genus}(M) \big). \qquad (9.4.1)
\end{aligned}$$

where the final equality follows from Euler's formula. If g is a Riemannian metric on $S(M)$ and $\kappa(x)$ is the Gaussian curvature of (S, g) at x, then the Gauss-Bonnet formula implies that

$$\mathbb{K}(M, \rho) = \frac{1}{|V|} \int_S \kappa(x)\, \mathrm{d}x,$$

so that the two notions of average curvature coincide. This is a special case of Theorem 9.4.9 below.

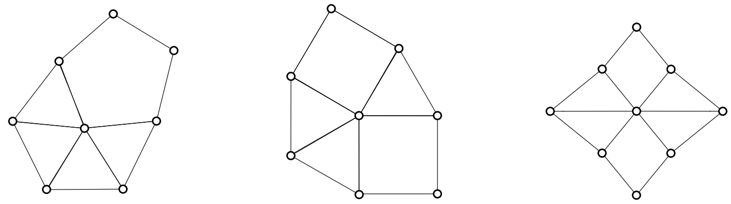

Figure 9.4: Left: A vertex with positive curvature $\pi/15$. Centre: A vertex with zero curvature. Right: A vertex with negative curvature $-\pi/3$.

Example 9.4.2 (k-angulations)**.** If M is a k-angulation (i.e., every face of M has degree k), then

$$\theta(v) = \frac{k-2}{k}\deg(v)\pi$$

for every vertex v of M. In particular, if M is an infinite plane tree, the angle sum at a vertex v of M is simply $\theta(v) = \pi\deg(v)$. Consequently, if (M,ρ) is a unimodular random k-angulation, then

$$\mathbb{K}(M,\rho) = \left(2 - \frac{k-2}{k}\mathbb{E}[\deg(\rho)]\right)\pi.$$

Example 9.4.3 (Curvature of the dual measure)**.** Let $\mathbb{P}$ be a unimodular probability measure on $\mathcal{M}_\bullet$ such that $\mathbb{P}[\deg(\rho)] < \infty$ and M has locally finite dual $\mathbb{P}$-a.s., and let $\mathbb{P}^\dagger$ be the dual measure. Then, by (9.2.2),

$$\begin{aligned}\mathbb{E}^\dagger\big[\kappa(\rho^\dagger)\big] &= 2\pi - \mathbb{E}^\dagger\Big[\sum_{f\perp\rho^\dagger}\frac{\deg(f)-2}{\deg(f)}\Big]\pi\\ &= 2\pi - \mathbb{E}\Big[\sum_{f\perp\rho}\frac{1}{\deg(f)}\Big]^{-1}\mathbb{E}\Big[\sum_{f\perp\rho}\frac{1}{\deg(f)}\sum_{v\perp f}\frac{\deg(v)-2}{\deg(v)}\Big]\pi.\end{aligned}$$

Applying the Mass Transport Principle yields that

$$\mathbb{E}^\dagger\big[\kappa(\rho^\dagger)\big] = 2\pi - \mathbb{E}\Big[\sum_{f\perp\rho}\frac{1}{\deg(f)}\Big]^{-1}\mathbb{E}\Big[\sum_{f\perp\rho}\frac{1}{\deg(f)}\sum_{v\perp f}\frac{\deg(\rho)-2}{\deg(\rho)}\Big]\pi$$

which can be rearranged to give

$$\mathbb{E}^\dagger[\kappa(\rho^\dagger)] = \mathbb{E}\Big[\sum_{f\perp\rho}\frac{1}{\deg(f)}\Big]^{-1}\mathbb{E}[\kappa(\rho)]$$

In particular, the average curvature of (M,ρ) under $\mathbb{P}$ has the same sign as that of $(M^\dagger,\rho^\dagger)$ under $\mathbb{P}^\dagger$.

Example 9.4.4 (Self-dual maps)**.** Let (M,ρ) be a self-dual unimodular random rooted map. Then

(9.2.3) implies that $\mathbb{E}[\sum_{f\perp\rho}\deg(f)^{-1}]=1$ and so

$$\mathbb{K}(M,\rho)=2\pi-\mathbb{E}\left[\sum_{f\perp\rho}\frac{\deg(f)-2}{\deg(f)}\right]\pi$$
$$=2\pi-\mathbb{E}[\deg(\rho)]\pi+2\mathbb{E}\Big[\sum_{f\perp\rho}\deg(f)^{-1}\Big]\pi=4\pi-\mathbb{E}[\deg(\rho)]\pi.$$

9.4.1 Curvature of submaps

Let M be a map with underlying graph G, and let z be a proper embedding of M into an orientable surface S. A **submap** of M is a map represented by a triple (H,S,z), where H is a connected subgraph of G such that the restriction of z to H is a proper embedding of H into S. If M is simply connected, then every connected subgraph of G is also a submap of M. This no longer holds if M is not simply connected. For example, if M is the product $\mathbb{Z}\times K_3$ of the integers $\mathbb{Z}$ with the triangle graph (a.k.a. the complete graph on three vertices) properly embedded into the infinite cylinder, then the subgraph $\mathbb{Z}\times\{v\}$, where v is one of the vertices of K_3, does not correspond to a submap of M.

Proposition 9.4.5 (Curvature of random submaps)**.** *Let (M,ρ) be a unimodular random rooted map with $\mathbb{E}[\deg(\rho)]<\infty$ and let ω be a connected bond percolation on M that is almost surely a submap of M. Then*
$$\mathbb{K}(\omega,\rho)=\mathbb{K}(M,\rho).$$

Proof. First suppose that M and ω both have a locally finite duals a.s. In this case, every face of ω is a union of finitely many faces of M. Define a mass transport as follows. For each vertex v of M and every face f of M incident to v, let $\hat{f}$ be the face of ω containing f, and transport a mass of
$$2\pi\frac{\deg(f)-2}{\deg(f)\deg(\hat{f})}$$
to each vertex u in the boundary of $\hat{f}$, where as usual we count f with multiplicity if it has multiple corners incident to v, and count u with multiplicity if it has multiple corners incident to f. The total mass sent out by each vertex is $\theta_M(v)$. For each face $\hat{f}$ of ω, consider the subgraph $G(\hat{f})$ of M spanned by all the edges of M that are incident to a face of M that are contained in $\hat{f}$. This graph can be drawn in the plane so that its faces correspond to the faces of M contained in $\hat{f}$, along with

an outside face. Applying Euler's formula to this graph, we have that

$$\begin{aligned}\sum_{f\in F(M),f\subseteq\hat f}(\deg(f)-2) &= 2|\{\text{edges of } G(\hat f)\}| - 2|\{\text{faces of } G(\hat f)\}| + 2\\ &\qquad - |\{\text{edges with exactly one side incident to } \hat f\}|\\ &= 2|\{\text{vertices of } G(\hat f)\}| - 2\\ &\qquad - |\{\text{edges with exactly one side incident to } \hat f\}|.\\ &= \deg(\hat f) - 2,\end{aligned}$$

(This is similar to [9, Lemma 18].) It follows that the total mass received by each vertex v is given by

$$\begin{aligned}\sum_{\hat f\in F(\omega),\hat f\perp v}\ \sum_{f\in F(M),f\subseteq\hat f}\ \sum_{u\perp f} 2\pi\frac{\deg(f)-2}{\deg(f)\deg(\hat f)} &= \sum_{\hat f\in F(\omega),\hat f\perp v}\ \sum_{f\in F(M),f\subseteq\hat f} 2\pi\frac{\deg(f)-2}{\deg(\hat f)}\\ &= \sum_{\hat f\in F(\omega),\hat f\perp v} 2\pi\frac{\deg(\hat f)-2}{\deg(\hat f)} = \theta_\omega(v).\end{aligned}$$

The claim now follows by applying the mass-transport principle.

Next, suppose that M has locally finite dual but that ω does not. For each infinite face $\hat f$ of ω, let $\langle e_i(\hat f)\rangle_{i\in\mathbb{Z}}$ be the set of oriented edges of ω that have their left hand side incident to $\hat f$, enumerated in counterclockwise order around the boundary of f. For each $n\geq 1$, define a percolation ω_n by letting $e\in\omega_n$ if and only if either $e\in\omega$ or if e^- and e^+ are both in the boundary of some infinite face $\hat f$ of ω and $e^- = e_i^-(\hat f)$ and $e^+ = e_j^+(\hat f)$ for some $i,j\in\mathbb{Z}$ with $|i-j|\geq n$. Observe that, since M has locally finite dual, ω_n has locally finite dual for every $n\geq 1$. Thus, it follows as above that $\mathbb{K}(\omega_n,\rho)=\mathbb{K}(M,\rho)$ for every $n\geq 1$. The sequence of random variables $|\theta_{\omega_n}(\rho)|$ are bounded by $2\pi\deg(\rho)$, and so it follows from the dominated convergence theorem that $\mathbb{K}(\omega_n,\rho)\to\mathbb{K}(\omega,\rho)$ as $n\to\infty$, completing the proof in this case.

Now suppose that the dual $M^\dagger$ is not locally finite. Let $\mathfrak{F}_\infty$ be the set of infinite degree faces of M. For each $n\geq 1$, let the map M_n be constructed from M as follows. For each infinite face f of (M,ρ), let $\langle e_i(f)\rangle_{i\in\mathbb{Z}}$ be the set of oriented edges of M with $e_i^l(f)=f$, enumerated in counterclockwise order around the boundary of f. Conditional on (M,ρ), let $\{U_i(f) : f\in F_\infty, i\in\mathbb{N}\}$ be a collection of i.i.d. Bernoulli-1/2 random variables. For each infinite face f of G and each $n\geq 1$, let $K_n(f)=\sum_{j=1}^n U_j(f)2^j$. For each $n\geq 1$, we form M_n by drawing an edge between $e(f)^-_{K_m+i2^m}$ and $e(f)^-_{K_m+(i+1)2^m}$ for each $i\in\mathbb{Z}$ and $m\geq n$. (Only M_1 is required for the current proof, we include the definition of M_n here for later use.) See Figure 4 for an illustration. It is clear that M_n is almost surely a locally finite map with a locally finite dual, and it follows from Lemma 9.2.2 that (M_n,ρ) is unimodular.

For each oriented edge e emanating from ρ in M that is incident to an infinite face of M, the

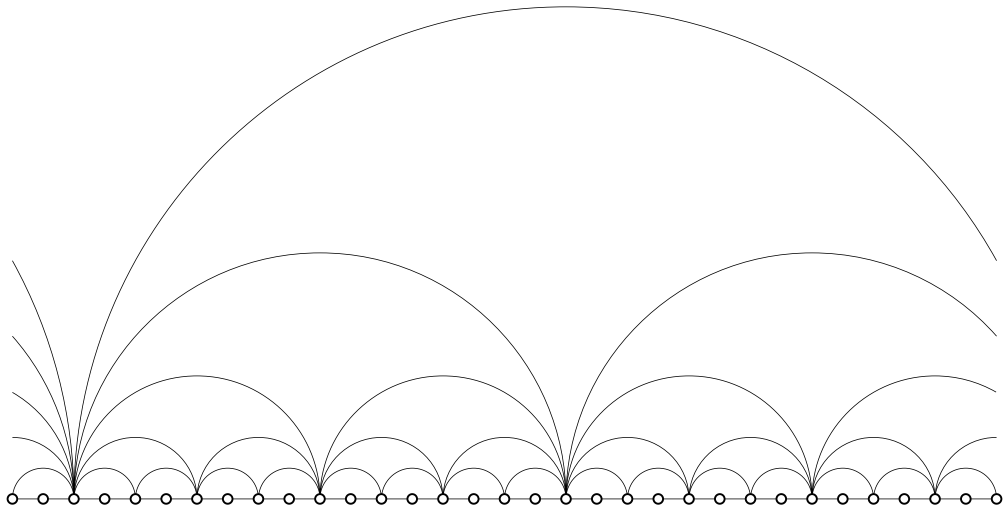

Figure 9.5: The map M_1 is defined by filling in each infinite face of M with a system of arcs. This figure demonstrates this procedure applied to one of the infinite faces of $\mathbb{Z}$.

expected number of additional arcs in M_1 drawn into the corner of this face at e is 2, so that

$$\begin{aligned}\mathbb{E}[\deg_{M_n}(\rho)] \leq \mathbb{E}[\deg_{M_1}(\rho)] &= \mathbb{E}[\deg(\rho)] + 2\mathbb{E}\left|\left\{e \in E^{\rightarrow} : e^- = \rho \text{ and } \deg(e^\ell) = \infty\right\}\right| \\ &\leq 3\mathbb{E}[\deg(\rho)] < \infty.\end{aligned}$$

By Lemma 9.3.2, we can consider both ω and M to be percolations on M_1 that are almost surely submaps, and we deduce from the above argument that $\mathbb{K}(\omega, \rho) = \mathbb{K}(M_1, \rho) = \mathbb{K}(M, \rho)$. $\square$

Remark 9.4.6. A straightforward extension of Proposition 9.4.5 is that, if ω is a connected percolation on (M, ρ) that is almost surely a submap, but might not include every vertex, then, letting (N, ρ') be the unimodular random map obtained from (ω, ρ) by conditioning on the event that ρ is in ω, we have that

$$\mathbb{K}(N, \rho') = \mathbb{P}(\rho \in \omega)^{-1}\mathbb{K}(M, \rho).$$

Proposition 9.4.7 (Upper semicontinuity of the average curvature)**.** *Let (M_n, ρ_n) be a sequence of unimodular random rooted maps with $\mathbb{E}[\deg(\rho_n)] < \infty$ converging weakly to a unimodular random rooted map (M, ρ), and suppose that $\mathbb{E}[\deg(\rho)] < \infty$. Then*

$$\mathbb{K}(M, \rho) \geq \lim_{n\to\infty} \mathbb{K}(M_n, \rho_n).$$

Proof. Let ρ_n be a uniformly chosen root vertex of M_n for each $n \geq 1$. First suppose that none of the maps M_n have any faces of degree 1. In this case, $\theta_{M_n}(v)$ is positive for every $n \geq 1$, and the claim follows from Fatou's lemma. Otherwise, let $(\hat{M}, \rho)$ and $(\hat{M}_n, \rho)$ be the unimodular random rooted maps obtained by removing all self-loops from M and M_n respectively. Clearly $(\hat{M}_n, \rho_n)$ converges weakly to (M, ρ), and by Fatou we have that $\mathbb{K}(\hat{M}, \rho) \geq \lim_{n\to\infty} \mathbb{K}(\hat{M}_n, \rho_n)$ as above. We conclude by applying Proposition 9.4.5 to deduce that $\mathbb{K}(M, \rho) = \mathbb{K}(\hat{M}, \rho)$ and $\mathbb{K}(M_n, \rho_n) = \mathbb{K}(\hat{M}_n, \rho_n)$ for every $n \geq 1$. $\square$

Combining Proposition 9.4.7 with the equation (9.4.1) has the following immediate corollary.

Corollary 9.4.8 (Theorem 9.1.1, (8) implies (1))**.** *Let (M, ρ) be a unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$ that is obtained as a Benjamini-Schramm limit of a sequence of finite maps $\langle M_n \rangle_{n \geq 1}$ such that*

$$\frac{\mathrm{genus}(M_n)}{|V(M_n)|} \xrightarrow[n \to \infty]{} 0.$$

Then $\mathbb{K}(M, \rho) \geq 0$.

We will later prove that $\mathbb{K}(M, \rho) \leq 0$ for any infinite, simply connected, unimodular random rooted planar map with finite expected degree in Section 9.5.

9.4.2 Invariance of the curvature

In this section, we consider *unimodular embeddings* of unimodular random rooted maps. We define a notion of average curvature associated to the embedding and show that, under certain integrability conditions, the average curvature associated to the embedding agrees with the average curvature that we defined combinatorially. This shows that the average curvature is a canonical quantity. Since this section is somewhat tangential to the rest of the paper, we will skip over some of the technical details.

We define a **metric surface embedded map** (MSEM) to be a locally finite map M together with a proper embedding z of M into an oriented metric surface S. A rooted MSEM is a MSEM together with a distinguished root vertex. Two rooted MSEMs are isomorphic if they are isomorphic as rooted maps, and there is an orientation preserving isometry between the two surfaces sending one embedding to the other.

We define the local topology on the set of isomorphism classes of rooted MSEMs by a variation on the local Gromov-Haussdorf topology: Namely, we set the distance between two rooted MSEMs (M_1, ρ_1, S_1, z_1) and (M_2, ρ_2, S_2, z_2) to be e^{-r}, where r is maximal such that there is a rooted graph isomorphism ϕ from the ball of radius r around ρ_1 in M_1 to the (graph distance) ball of radius r about ρ_2 in M_2 such that ϕ preserved the cycling of the edges emanating from each vertex in the interior of the ball, and there exist functions ψ_1 and ψ_2 from the balls of (metric) radius r around $z_1(\rho_1)$ and $z_2(\rho_2)$ in S_1 and S_2 respectively, denoted $B_1(r)$ and $B_2(r)$, into some common metric space X so that the Haussdorf distance between $\psi_1(B_1)$ and $\psi_2(B_2)$ is at most $1/r$, and the Haussdorf distance between $\psi_1(z_1(e) \cap B_1(r))$ and $\psi_2(z_2(\phi(e)) \cap B_2(r))$ is at most $1/r$ for every edge e contained in the graph distance ball of radius r around ρ_1.

We define doubly rooted MSEMs, the local topology on the set of isomorphism classes of doubly rooted MSEMs, and unimodular random rooted MSEMs similarly to the graph case. It is straightforward (but rather tedious) to encode the structure of a MSEM as a marking of the underlying map of the MSEM, so that all of the usual machinery of unimodularity transfers to this setting. If (M, ρ) is a unimodular random rooted map and (M, S, z, ρ) is a unimodular random rooted MSEM with underlying map (M, ρ), we call (S, z) a **unimodular embedding** of M.

In practice, unimodular embeddings often arise as measurable, automorphism equivariant functions of the map, such as the embeddings given by circle packing and the conformal embedding. These are easily seen to be unimodular since every mass-transport on the associated MSEM is induced by a mass-transport on the map.

Let us give an example of a unimodular embedding. Given a map M such that every face of M has degree at least three, a natural way to embed it on a surface is as follows. We associate to each face f of degree d a regular d-gon with sides of length 1. If two faces share an edge, we identify the corresponding edges of the polygons. For infinite faces, the corresponding polygon is a half-plane $\{z \in \mathbb{C} : \mathrm{Im}(z) \geq 0\}$ with edges $\{[n, n+1] : n \in \mathbb{Z}\}$ along the boundary, which we think of as a regular ∞-gon. In this surface, there is no curvature on any of the faces, as they are all flat pieces of $\mathbb{R}^2$. The surface is also smooth along the edges, as two faces are glued along straight segments of unit length. Thus, all of the curvature of the surface is concentrated on the vertices, and in fact the atom of curvature $\kappa_s(v)$ at v (see below) is exactly the value $\kappa(v)$ that we defined combinatorially at the beginning of this section.

One can more generally obtain a unimodular embedding by gluing together (possibly random) shapes that are not necessarily polygons, and, in particular, can also extend the construction to allow for faces of degree 1 and 2.

In what follows, we restrict ourselves to surfaces that have a smooth Riemannian metric, except possibly for cone-like singularities at vertices of the map, and assume that the edges of the map are embedded as smooth curves in the surface. We call a unimodular embedding of a map satisfying these conditions a **smooth embedding**. Given a smooth embedding of a map M, for each oriented edge e of M, we let $\mathrm{ang}(e)$ be the angle between e and the edge following e in the clockwise ordering of the oriented edges emanating from e^-. Every smooth Riemannian surface with cone-like singularities has a Gaussian curvature κ associated with its metric, which is a signed measure on the surface (see e.g. [9]). The curvature measure κ is absolutely continuous with respect to the area measure on the surface except possibly at the cone-like singularities at the vertices. We let $\kappa_s(v)$ be the atom of curvature at a vertex v, and let $\kappa(f)$ be the total curvature of the face f, which is well-defined if f has finite degree. Moreover, every oriented edge e of the map has a well defined total geodesic curvature in the embedding, which we denote $\kappa_g(e)$. In a smooth embedding of a map, the mass of the atom of curvature at a vertex is given by

$$\kappa_s(v) = 2\pi - \sum_{e:\, e^- = v} \mathrm{ang}(e).$$

Given a smooth embedding of a map with a locally finite dual, we define the **total curvature** at a vertex v by

$$\kappa(v) = \kappa_s(v) + \sum_{f \perp v} \frac{\kappa(f)}{\deg(f)}.$$

We are now ready to state our "invariance of the average curvature" theorem.

Theorem 9.4.9. *Let (M, ρ) be a unimodular random rooted map with locally finite dual and let (S, z) be a smooth unimodular embedding of M. Suppose further that either*

1. *$\kappa_g(e) = 0$ for every oriented edge e of M, and $\kappa(f) \leq 0$ for every face f of M, or*

2. *$\sum_{e:e^- = \rho} |\kappa_g(e)|$ and $|\kappa(\rho)|$ both have finite expectation.*

Then

$$\mathbb{E}[\kappa(\rho)] = \mathbb{K}(M, \rho).$$

We shall require the following signed version of the mass-transport principle: If (G, ρ) is a unimodular random rooted graph and ϕ is a measurable function from $\mathcal{G}_{\bullet\bullet}$ to $\mathbb{R}$ such that

$$\mathbb{E}\left[\sum_{v \in V} \left|\phi(G, \rho, v)\right|\right] < \infty, \quad \text{then} \quad \mathbb{E}\left[\sum_{v \in V} \phi(G, \rho, v)\right] = \mathbb{E}\left[\sum_{v \in V} \phi(G, v, \rho)\right].$$

This follows easily from the usual mass-transport principle for ϕ non-negative. A similar signed mass-transport principle holds for unimodular random rooted maps and MSEMs.

Proof. For each oriented edge e of M, let $\operatorname{ang}(e)$ be the angle between e and the edge $\sigma(e)$ following e in the clockwise ordering of the oriented edges emanating from e^-. The Gauss-Bonnet Theorem (see e.g. [9, Chapter VI, Section 7]) implies that for any face f of M,

$$\sum_{e:e^r = f} [\pi - \operatorname{ang}(e)] + \sum_{e:e^r = f} \kappa_g(e) + \kappa(f) = 2\pi.$$

We rewrite this identity as

$$(\deg(f) - 2)\pi = \sum_{e:e^r = f} \left[\operatorname{ang}(e) - \frac{\kappa(f)}{\deg(f)} - \frac{1}{2}\kappa_g(e) - \frac{1}{2}\kappa_g(-\sigma(e))\right]. \tag{9.4.2}$$

We now define a mass transport on M. Given a vertex u, for each face f that u belongs to, u sends to each vertex v in the boundary of f (including u itself) the mass

$$\psi(u, f) = \frac{1}{\deg(f)} \sum_{e:e^r = f,\, e^- = u} \left[\operatorname{ang}(e) - \frac{\kappa(f)}{\deg(f)} - \frac{1}{2}\kappa_g(e) - \frac{1}{2}\kappa_g(-\sigma(e))\right],$$

where, as usual, v receives a multiple of this mass if it is multiply incident to the face. Note that if u and v share more than a single face, there is a term for each face containing the two, and there is also mass sent from u to itself, with a term for each face incident to u.

Let $\phi(u, v)$ be the total mass sent from u to v. If condition (1) holds, then ϕ is positive.

Otherwise, condition (2) holds, and we have that

$$\sum_{v\in V}|\phi(u,v)| \leq \sum_{f\perp u}\deg(f)|\psi(u,f)| \leq |\kappa(u)| + \sum_{e:e^-=u}|\kappa_g(e)|$$

and hence

$$\mathbb{E}\sum_{v\in V}|\phi(\rho,v)| < \infty.$$

We can apply the Mass-Transport Principle to ϕ in either case.

Let us consider the total mass sent out from v. The quantity corresponding to a face $f \perp v$ is sent to each of the $\deg(f)$ corners of f, which cancels the $1/\deg(f)$ factor. The sum over faces adjacent to v of $\mathrm{ang}(e)$ is precisely $2\pi - \kappa_s(v)$. Together with the sum over f of $\kappa(f)/\deg(f)$ this comes to $2\pi - \kappa(v)$. Meanwhile, since $\kappa_g(e) = -\kappa_g(-e)$, the terms involving geodesic curvatures of edges cancel when we sum over all faces incident to v. Thus, we have that the total mass sent out by v is exactly $\kappa(v)$. On the other hand, the total mass received by a vertex v from vertices of a face $f \perp v$ is $(\deg(f) - 2)\pi/\deg(f)$ by (9.4.2), and so the result follows. □

Example 9.4.10 (Integrability conditions are necessary)**.** Consider the map obtained by triangulating each of the two infite faces of $\mathbb{Z}$ as in Figure 9.5. This map admits a unimodular embedding in the hyperbolic plane, defined by drawing the edges of $\mathbb{Z}$ as segments of length one along a bi-infinite geodesic, and drawing the other edges of the map as circular arcs. This embedding is a measurable function of the map, and is therefore unimodular. However, it does not satisfy the integrability requirements of Theorem 9.4.9. Indeed, the conclusion of Theorem 9.4.9 fails for this embedding: the curvature $\kappa_s(v) + \sum_{f\perp v}\kappa(f)/\deg(f)$ is negative for every vertex of the map, but it is easily seen that the map is invariantly amenable and has finite expected degree, so that $\mathbb{K}(M,\rho) \geq 0$ by Corollary 9.4.8 and Corollary 9.3.6.

Example 9.4.11 (Voronoi diagrams and Delaunay triangulations of point processes)**.** We say that a set of points in the plane is in **general position** if no more than three points in the set lie on any given circle or line. Given a such a set of points Z in general position in either the Euclidean plane or the hyperbolic plane, the **Delaunay triangulation** of Z is the simple triangulation that has the points of Z as its vertices, and, for each triple of distinct points u, v and w in Z, contains a geodesic triangle with corners u, v and w if and only if the unique disc containing u, v and w in its boundary does not contain any other points of Z. Suppose that Z is an isometry-invariant, locally finite point process in either the Euclidean plane or the hyperbolic plane, and let $\hat{Z}$ be the Palm version of Z that is conditioned to have a point at the origin. Then the Delaunay triangulation of $\hat{Z}$ is unimodular when rooted at the point at the origin; see [37] for a study of the Poisson case. This embedding satisfies condition (1) of Theorem 9.4.9, and we deduce that, unsurprisingly, Delaunay triangulations of hyperbolic point processes are hyperbolic while Delaunay triangulations of Euclidean point processes are parabolic. (This also follows from the methods of [21].) The dual of the Delaunay triangulation is the Voronoi diagram of the point process, which can also be made unimodular as in Section 9.2.5.

We remark that it is possible to use Theorem 9.4.9 to prove Proposition 9.4.5 by embedding M in the polygonal manifold associated to ω.

9.5 Spanning forests

9.5.1 Uniform spanning forests

Our primary way to relate the average curvature of a map M with the various probabilistic properties listed in Theorem 9.1.1 is a formula relating the average degree of the *free uniform spanning forest* of M to its average curvature. To this aim, we begin by succinctly discussing the topic of uniform spanning forest, referring the reader to [44, 173] for a comprehensive treatment.

For each finite graph G, let UST_G be the uniform measure on spanning trees of G (i.e. connected subgraphs of G containing every vertex and no cycles), which is the law of a percolation on G. There are two natural ways to define infinite volume limits of the uniform spanning tree. Let $G = (V, E)$ be an infinite, locally finite, connected graph. An **exhaustion** of G is an increasing sequence $\langle V_n \rangle_{n \geq 1}$ of finite connected subsets of V such that $\bigcup_{n \geq 1} V_n = V$. Given an exhaustion $\langle V_n \rangle_{n \geq 1}$ of G, we define G_n to be the subgraph of G induced by V_n for each $n \geq 1$. The **free uniform spanning forest** measure of G is defined as the weak limit of the uniform spanning tree measures on the graphs G_n. That is, for each finite set $S \subset E$,

$$\mathsf{FUSF}_G(S \subset \mathfrak{F}) := \lim_{n \to \infty} \mathsf{UST}_{G_n}(S \subset T),$$

where $\mathfrak{F}$ is a sample of the FUSF of G and T is a sample of the UST of G_n. For each $n \geq 1$, we also construct a graph G_n^* from G by identifying every vertex in $V \setminus V_n$ into a single vertex ∂_n, and deleting all of the resulting self loops from ∂_n to itself. We then define the **wired uniform spanning forest** measure of G to be the weak limit of the uniform spanning tree measures on the graphs G_n^*. That is, for each finite set $S \subset E$,

$$\mathsf{WUSF}_G(S \subset \mathfrak{F}) := \lim_{n \to \infty} \mathsf{UST}_{G_n^*}(S \subset T),$$

where $\mathfrak{F}$ is a sample of the WUSF of G and T is a sample of the UST of G_n. The study of uniform spanning forests was pioneered by Pemantle [190], who showed that both limits exist for any graph G and in particular are independent of the choice of exhaustion.

The link between the USFs and amenability is the following.

Theorem 9.5.1 ([7, Proposition 18.14]: Theorem 9.1.1, (2) implies (12))**.** *If (G, ρ) is an invariantly amenable unimodular random rooted graph, then $\mathsf{FUSF}_G = \mathsf{WUSF}_G$ almost surely.*

The USFs enjoy the following properties:

1. **(Free dominates wired.)** The measure FUSF_G stochastically dominates the measure WUSF_G for every graph G.

2. **(Domination and subgraphs)** let H be a connected subgraph of G. Then the FUSF of H stochastically dominates the restriction of the FUSF of G to H.

3. (**Expected degree of the WUSF.**) The expected degree in the WUSF of the root of any unimodular random rooted graph is 2 [7, Proposition 7.3].

Note that, since a connected spanning forest cannot be strictly contained in another spanning forest, the stochastic domination (1) above has the following immediate consequence.

Proposition 9.5.2 (Theorem 9.1.1, (13) implies (12))**.** *Let G be a graph. If the wired uniform spanning forest of G is connected almost surely, then the wired and free spanning forests of G coincide.*

If (G,ρ) is a unimodular random rooted graph and $\mathfrak{F}$ is a sample of either WUSF_G or FUSF_G, then the marked graph $(G,\rho,\mathfrak{F})$ is also unimodular (i.e., $\mathfrak{F}$ is a percolation on (G,ρ)): Since the definitions of FUSF_G and WUSF_G do not depend on the choice of exhaustion, for each mass transport $f : \mathcal{G}_{\bullet\bullet}^{\{0,1\}} \to [0,\infty]$, the expectations

$$f^F(G,u,v) = \mathsf{FUSF}_G\left[f(G,u,v,\mathfrak{F})\right] \qquad \text{and} \qquad f^W(G,u,v) = \mathsf{WUSF}_G\left[f(G,u,v,\mathfrak{F})\right]$$

are also mass transports. Using this observation, we deduce the mass-transport principle for $(G,\rho,\mathfrak{F})$ from that of (G,ρ).

Connections to random walk and potential theory.

Although the uniform spanning tree of each G_n or G_n^* is connected, the limiting random subgraph can be disconnected. Indeed, Pemantle [190] proved that WUSF and FUSF of $\mathbb{Z}^d$ coincide for all $d \geq 1$, and are connected if and only if $d \leq 4$. A complete characterisation of the connectivity of the WUSF was given by Benjamini, Lyons, Peres and Schramm [44]. A connected, locally finite graph is said to have the **intersection property** if the traces of two independent simple random walks started from any two vertices of the graph have infinite intersection almost surely (or, equivalently, if the two traces have non-empty intersection almost surely). A graph is said to have the **non-intersection property** if the traces of two independent simple random walks on the graph have finite intersection almost surely.

Theorem 9.5.3 (Benjamini, Lyons, Peres and Schramm [44]: Theorem 9.1.1, equivalence of (14) and (13))**.** *Let G be an infinite, locally finite, connected graph and let $\mathfrak{F}$ be a sample of WUSF_G. Then $\mathfrak{F}$ is connected almost surely if and only if G has the intersection property. If G has the non-intersection property, then $\mathfrak{F}$ has infinitely many connected components almost surely.*

In general, a graph need not have either of the intersection or non-intersection properties. For example, the graph formed by connecting two disjoint copies of $\mathbb{Z}^3$ by a single edge between their origins does not have either property. However, it is easily seen that this is not the case for reversible random rooted graphs.

Lemma 9.5.4. *Let (G,ρ) be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then G either has the intersection property or the non-intersection property almost surely.*

Proof. By biasing by the degree we may assume that (G,ρ) is reversible. We may assume also that (G,ρ) is ergodic, otherwise taking an ergodic decomposition. Let $\langle X_n\rangle_{n\geq 0}$ and $\langle X_{-n}\rangle_{n\geq 0}$ be independent random walks on G started at ρ. Then the event that the traces of $\langle X_n\rangle_{n\geq 0}$ and $\langle X_{-n}\rangle_{n\geq 0}$ have infinite intersection is an invariant event for the stationary sequence $\langle (G, \langle X_{n+k}\rangle_{n\in\mathbb{Z}}\rangle_{k\in\mathbb{Z}}$ and therefore has probability either zero or one by ergodicity. □

Thus, the WUSF of a unimodular random rooted graph with finite expected degree is either connected or has infinitely many connected components almost surely.

Recall that a function $h : V \to \mathbb{R}$ defined on the vertex set of a graph $G = (V,E)$ is said to be **harmonic** if

$$h(v) = \frac{1}{\deg(v)} \sum_{u\sim v} h(u)$$

for every vertex v of G, or equivalently if $\langle h(X_n)\rangle_{n\geq 0}$ is a martingale when $\langle X_n\rangle_{n\geq 0}$ is a random walk on G. A graph is said to be **Liouville** if it does not admit any non-constant bounded harmonic functions, and **non-Liouville** otherwise. The following proposition, which follows from the martingale convergence theorem, is well-known (see [173, Exercise 14.28]).

Proposition 9.5.5 (Theorem 9.1.1, (14) implies (10))**.** *Let G be a connected graph. If G has the intersection property, then G is Liouville.*

The converse of Proposition 9.5.5 does not hold for general graphs. For example, $\mathbb{Z}^d$ is Liouville for all $d \geq 1$ but has the intersection property only for $d \leq 4$. However, Benjamini, Curien and Georgakopoulos [40] proved that the converse does hold for planar graphs.

Theorem 9.5.6 ([40] Theorem 9.1.1, (10) implies (14))**.** *] Let G be a planar graph. Then G is Liouville if and only if it has the intersection property.*

The **Dirichlet energy** of a function $f : V \to \mathbb{R}$ defined on the vertex set of G is defined to be

$$\mathcal{E}(f) = \frac{1}{2} \sum_{e\in E^{\rightarrow}} \Big(f(e^-) - f(e^+)\Big)^2.$$

The following is classical; see [173, Exercise 9.43].

Proposition 9.5.7 (Theorem 9.1.1, (10) implies (11))**.** *Let G be a connected graph. Then the bounded harmonic functions of finite Dirichlet energy are dense in the space of harmonic functions of finite Dirichlet energy. In particular, if G admits a non-constant harmonic function of finite Dirichlet energy, then G admits a bounded non-constant harmonic function of finite Dirichlet energy.*

Benjamini, Lyons, Peres and Schramm [44] related analytic properties of G to the WUSF and FUSF of G.

Theorem 9.5.8 (Benjamini, Lyons, Peres and Schramm [44]: Theorem 9.1.1, equivalence of items (12) and (11))**.** *Let G be an infinite connected graph. Then the measures FUSF_G and WUSF_G are distinct if and only if G admits harmonic functions of finite Dirichlet energy.*

In general, the FUSF is not nearly as well understood as the WUSF: no criterion for its connectivity is known, nor is it known whether the number of components of the FUSF is non-random in every graph. However, for simply connected maps, the FUSF is relatively well understood thanks to the following duality: Given a map M and a set $W \subset E$, let $W^\dagger := \{e^\dagger \in E^\dagger : e \notin W\}$ be the set of dual edges whose corresponding primal edges are not contained in W. Observe that if t is a spanning tree of a finite planar map M, then the dual $t^\dagger$ is a spanning tree of $M^\dagger$ – it is connected because t has no cycles, and has no cycles because t is connected. This observation leads to the following.

Proposition 9.5.9 (USF Duality [44, Theorem 12.2])**.** *Let M be a simply connected map with locally finite dual $M^\dagger$ and let $\mathfrak{F}$ be a random variable with law FUSF_M. Then $\mathfrak{F}^\dagger$ has the law of $\mathsf{WUSF}_{M^\dagger}$.*

In general, if M is an infinite simply connected map and $\mathfrak{F}$ is an essential spanning forest of M (that is, a spanning forest such that every component is infinite), then $\mathfrak{F}^\dagger$ is an essential spanning forest of $M^\dagger$ – it is a forest because every component of $\mathfrak{F}$ is infinite, and is essential because $\mathfrak{F}$ has no cycles. Moreover, the forest $\mathfrak{F}$ is connected if and only if every component of $\mathfrak{F}^\dagger$ is one-ended. Furthermore, we have the following.

Proposition 9.5.10 (Aldous-Lyons [7, Theorem 8.9]: Theorem 9.1.1, (13) implies (2))**.** *Let (G, ρ) be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$ and let $\mathfrak{F}$ be a sample of WUSF_G. If $\mathfrak{F}$ is connected almost surely, then (G, ρ) is invariantly amenable.*

Proof. Suppose that $\mathfrak{F}$ is connected almost surely. Since $\mathfrak{F}$ has at most two ends almost surely by [7, Theorem 6.2], [7, Theorem 8.9] implies that (G, ρ) is invariantly amenable. $\square$

9.5.2 Results

The first main result of this section relates the average curvature and the expected degree of the FUSF in a simply connected unimodular random rooted map.

Theorem 9.5.11. *Let (M, ρ) be an infinite, simply connected unimodular random map, and suppose that $\mathbb{E}[\deg(\rho)] < \infty$, and let $\mathfrak{F}$ be a sample of FUSF_M. Then*

$$\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] = \frac{1}{\pi}\mathbb{E}[\theta(\rho)] = 2 - \frac{1}{\pi}\mathbb{K}(M, \rho). \tag{9.5.1}$$

As an easy consequence we get the following component of Theorem 9.1.1.

Corollary 9.5.12 (Theorem 9.1.1, non-positivity of average curvature and equivalence of (1) and (12))**.** *Let (M, ρ) be an ergodic, unimodular random graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then the average curvature $\mathbb{P}[\kappa(\rho)]$ is non-positive and is zero if and only if $\mathsf{FUSF}_M = \mathsf{WUSF}_M$ almost surely.*

Proof. Let $\mathfrak{F}$ be a sample of FUSF_M. Since the expected degree of the WUSF in any unimodular random rooted graph is two, and FUSF_M stochastically dominates WUSF_M, we have that $\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] \geq 2$ and that the measures FUSF_M and WUSF_M differ almost surely if and only if this inequality is strict. We conclude by applying Theorem 9.5.11. $\square$

Theorem 9.5.11 follows as an immediate corollary of Proposition 9.4.5 and the following theorem, which is the second main result of this section. In Section 9.5.6, we give an alternative, duality-based proof of Theorem 9.5.11 that does not rely on Theorem 9.5.13 or Proposition 9.4.5, and also applies to the free minimal spanning forest.

Theorem 9.5.13 (Connectivity of the FUSF)**.** *Let (M,ρ) be a simply connected unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$. Then the free uniform spanning forest of M is connected almost surely.*

Since the measure FUSF_G stochastically dominates WUSF_G for every graph G, we deduce the following immediate corollary.

Corollary 9.5.14 (Theorem 9.1.1, equivalence of (12) and (13))**.** *Let (M,ρ) be a simply connected unimodular random map with $\mathbb{E}[\deg(\rho)] < \infty$. Then $\mathsf{FUSF}_M = \mathsf{WUSF}_M$ if and only if the wired uniform spanning forest of M is connected almost surely.*

If M has locally finite dual almost surely, then, by Proposition 9.5.9, Theorem 9.5.13 is equivalent to the statement that every component of the WUSF of $M^\dagger$ is one-ended almost surely. Fortunately, it is known that every component of the WUSF is one-ended almost surely in several large classes of graphs: The following was proven by the second author [127, 128] and followed earlier works by Pemantle [190], Benjamini, Lyons, Peres, and Schramm [44], and Aldous and Lyons [7].

Theorem 9.5.15 ([127, 128])**.** *Let (G,ρ) be a transient unimodular random rooted graph. Then every component of the wired uniform spanning forest of G is one-ended almost surely.*

In the case that the dual $M^\dagger$ is locally finite, Theorem 9.5.13 follows as a corollary to Theorem 9.5.15. The case that dual is not locally finite is contained in Proposition 9.5.18 and is handled by a separate argument. Theorem 9.5.13 is complemented by concurrent work by the second and third authors [130], who prove the corresponding theorem for deterministic simply connected maps with bounded degrees. See [170] for further one-endedness results in the deterministic setting.

It is an open question whether the WUSF is one-ended almost surely in every one-ended unimodular random rooted graph (G,ρ) with $\mathbb{E}[\deg(\rho)] < \infty$, without the assumption of transience. The final result of this section is to prove that this is holds in the planar case.

Theorem 9.5.16. *Let (M,ρ) be a recurrent unimodular random rooted planar map with $\mathbb{E}[\deg(\rho)] < \infty$. Then the WUSF of M has the same number of ends as M almost surely (which is either one or two since M is recurrent).*

When M has locally finite dual, Theorem 9.5.16 follows immediately by duality and the fact that the WUSF of a reversible random rooted graph is either connected or has infinitely many connected components almost surely; our contribution is to handle the case that M has infinite faces.

9.5.3 Proof of Theorem 9.5.13 and Theorem 9.5.16

Proof of Theorem 9.5.13 when the dual is locally finite. Let (M, ρ) be a unimodular random rooted map with $\mathbb{P}[\deg(\rho)] < \infty$, and suppose that the dual $M^\dagger$ is locally finite almost surely. We may assume that (M, ρ) is ergodic, otherwise taking an ergodic decomposition. Let $\mathfrak{F}$ be a sample of FUSF_M and let $\mathfrak{F}^\dagger$ be the dual forest. Since $\mathbb{E}[\deg(\rho)] < \infty$, the law of $(M^\dagger, \rho^\dagger)$ is equivalent to the law of a reversible random rooted map by Proposition 9.2.4. If $M^\dagger$ is almost surely transient, Theorem 9.5.15 implies that every component of $\mathfrak{F}^\dagger$ is one-ended almost surely, and we deduce that $\mathfrak{F}$ is connected almost surely. If $M^\dagger$ is almost surely recurrent, then $\mathfrak{F}^\dagger$ is connected almost surely, and consists of a single tree with at most two ends. It follows that the measures $\mathsf{FUSF}_{M^\dagger}$ and $\mathsf{WUSF}_{M^\dagger}$ coincide and hence the measures FUSF_M coincide WUSF_M by duality. If the single tree of $\mathfrak{F}^\dagger$ were two-ended then $\mathfrak{F}$ would have exactly two components, contradicting Lemma 9.5.4. $\square$

The remainder of this section is dedicated to the proof of Theorems 9.5.13 and 9.5.16 in the presence of infinite faces. We begin by developing a variant of Wilson's algorithm that allows us to sample the dual of the FUSF using random walks when the dual is not locally finite.

Given a graph G and a path γ in G that is either finite or transient, i.e. visits each vertex of G at most finitely many times, the **loop-erasure** $\mathsf{LE}(\gamma)$ is formed by erasing cycles from γ chronologically as they are created. Formally, $\mathsf{LE}(\gamma)_i = \gamma_{t_i}$ where the times t_i are defined recursively by $t_0 = 0$ and $t_i = 1 + \max\{t \geq t_{i-1} : \gamma_t = \gamma_{t_{i-1}}\}$. The loop-erasure of a simple random walk is known as **loop-erased random walk** and was first studied by Lawler [161].

Wilson's algorithm [228] is a method of sampling a uniform spanning tree of a finite graph by joining together loop-erased random walk paths. Benjamini, Lyons, Peres and Schramm [44] later introduced a variant of Wilson's algorithm for sampling the WUSF of an infinite transient graph, known as **Wilson's algorithm rooted at infinity**. Let $G = (V, E)$ be a connected, locally finite graph and let $\langle v_i : i \geq 1\rangle$ be an enumeration of the vertex set V. We sample a sequence of forests $\langle \mathfrak{F}_i \rangle_{i\geq 0}$ in G recursively as follows:

1. If G is finite or recurrent, fix a vertex v_0 of G and let $\mathfrak{F}_0 = \{v\}$. Otherwise G is transient and we set $\mathfrak{F}_0 = \emptyset$.

2. Given $\mathfrak{F}_i$, start a simple random walk from f_i in M independently of everything we have already sampled, stopped if and when it first hits a vertex already included in $\mathfrak{F}_i$.

3. Take the loop-erasure of this random walk path, and let $\mathfrak{F}_{i+1}$ be the union of $\mathfrak{F}_i$ and this loop-erased path.

4. Let $\mathfrak{F} = \bigcup_{i \geq 0} \mathfrak{F}_i$.

This procedure is referred to as **Wilson's algorithm rooted at** v_0 when G is finite or recurrent, and **Wilson's algorithm rooted at infinity** when G is transient: The resulting random forest has law UST_G when G is finite [228] and WUSF_G when G is infinite [44].

Now suppose that M is a simply connected map with at least one infinite degree face. Let $F_{\text{fin}} = \{f \in F(M) : \deg(f) < \infty\}$ be the set of finite degree faces of M and let $F_\infty = F \setminus F_{\text{fin}}$ be the set of infinite degree faces of M. Let $\langle f_i : i \geq 1\rangle$ be an enumeration of F_{fin}. We sample an increasing sequence of forests $\langle \mathfrak{F}_i^\dagger : i \geq 1\rangle$ in $M^\dagger$ recursively as follows:

1. Let $\mathfrak{F}_0^\dagger = F_\infty$.
2. Given $\mathfrak{F}_i^\dagger$, start a simple random walk from f_i in $M^\dagger$ independently of everything we have already sampled, stopped if and when it first hits a vertex already included in $\mathfrak{F}_i^\dagger$.
3. Take the loop-erasure of this random walk path, and let $\mathfrak{F}_{i+1}^\dagger$ be the union of $\mathfrak{F}_i^\dagger$ and this loop-erased path.
4. Let $\mathfrak{F}^\dagger = \bigcup_{i \geq 0} \mathfrak{F}_i^\dagger$.

We call this procedure **Wilson's algorithm rooted at** $\{\infty\} \cup F_\infty$.

Proposition 9.5.17. *Let $M = (V, E, \sigma)$ be a simply connected map with dual $M^\dagger$ and let $\mathfrak{F}^\dagger$ be a random subset of $E^\dagger$ sampled by Wilson's algorithm rooted at $\{\infty\} \cup F_\infty$. Then $\mathfrak{F} = (\mathfrak{F}^\dagger)^\dagger$ is a sample of* FUSF_M.

Proof. Let $\langle V_n \rangle_{n\geq 0}$ be an exhaustion of V such that the submap of M induced by $V \setminus V_n$, denoted M_n, does not have any finite connected components for any n. The dual of M_n may be constructed from $M^\dagger$ by identifying every face f of M that does not have all of its constituent vertices included in M_n into a single vertex ∂_n, and deleting all the self-loops that are created. In particular, all infinite faces of M are identified with ∂_n for every $n \geq 1$. Note also that if we start a simple random walk $\langle X_n \rangle_{n\geq 0}$ on $M^\dagger$, started at some face $f \in F_{\text{fin}}$ and stopped the first time it hits F_∞, then, since every time the walk revisits f it has a constant, positive probability of hitting F_∞ before revisiting f, the stopped random walk path is either finite or transient almost surely. Given these observations, the rest of the proof proceeds similarly to the usual proof of the veracity of Wilson's algorithm rooted at infinity [44, Theorem 5.1].

Let H be a finite set of edges of $M^\dagger$, and let $f_1, \ldots, f_l$ be an enumeration of the set of faces f of M that are endpoints of at least one of the edges in H. Let $\langle \langle X_j^i \rangle_{j\geq 0} : i = 1, \ldots, l\rangle$ be a collection of independent random walks in $M^\dagger$, where the walk $\langle X_j^i \rangle_{j\geq 0}$ is started at f_i and stopped the first time that it hits an infinite face of M. Run Wilson's algorithm in $M_n^\dagger$, rooted at ∂_n and starting with the faces $f_1, \ldots, f_l$ in that order, using the random walks $\langle X_j^i \rangle_{j\geq 0}$: For each $i \in [l]$, let τ_i^n be the first time that the random walk $\langle X_j^i \rangle_{j\geq 0}$ visits the portion of the spanning tree generated up to time $i-1$, so that

$$\mathsf{UST}_{M_n^\dagger}(H \subset T) = \mathbb{P}\Big(H \subseteq \bigcup_{i=1}^{l} \mathsf{LE}\big(\langle X_j^i\rangle_{j=0}^{\tau_i^n}\big)\Big).$$

Now, similarly, run Wilson's algorithm on $M^\dagger$ rooted at $\{\infty\}\cup F_\infty$, starting with the faces $f_1,\ldots,f_l$ in that order and using the random walks $\langle X_j^i\rangle_{j\geq 0}$, and let τ_i be the first time that the random walk $\langle X_j^i\rangle_{j\geq 0}$ visits the portion of the spanning tree generated up to time $i-1$ (which might now be infinite). Since the walks $\langle X_j^i\rangle_{j\geq 0}$ are finite or transient almost surely, we have that

$$\tau_i^n \to \tau_i \text{ and } \mathsf{LE}\big(\langle X_j^i\rangle_{j=0}^{\tau_i^n}\big) \to \mathsf{LE}\big(\langle X_j^i\rangle_{j=0}^{\tau_i}\big)$$

almost surely as $n\to\infty$. It follows that

$$\begin{aligned}\mathsf{FUSF}_M(H\subset \mathfrak{F}^\dagger) = \lim_{n\to\infty}\mathsf{UST}_{M^\dagger}(H\subset T) &= \lim_{n\to\infty}\mathbb{P}\Big(H \subseteq \bigcup_{i=1}^{l} \mathsf{LE}\big(\langle X_j^i\rangle_{j=0}^{\tau_i^n}\big)\Big)\\ &= \mathbb{P}\Big(H \subseteq \bigcup_{i=1}^{l} \mathsf{LE}\big(\langle X_j^i\rangle_{j=0}^{\tau_i}\big)\Big)\end{aligned}$$

completing the proof. □

In the case that M has faces of infinite degree, Theorems 9.5.13 and 9.5.16 follow immediately from the following proposition. (The case of Theorem 9.5.16 in which M is not simply connected is trivial, since in this case M and $\mathfrak{F}$ must both have two ends.)

Proposition 9.5.18. *Let (M,ρ) be a simply connected, unimodular random rooted map with $\mathbb{P}[\deg(\rho)] < \infty$ and suppose that the dual $M^\dagger$ contains a vertex of infinite degree almost surely. Let $\mathfrak{F}$ be a sample of FUSF_M. Then every connected component of $\mathfrak{F}^\dagger \setminus F_\infty$ is finite almost surely, and consequently $\mathfrak{F}$ is connected almost surely.*

The main ingredient of Proposition 9.5.18 is the following simple lemma.

Lemma 9.5.19. *Let (M,η) be a simply connected reversible random edge-rooted map and suppose that the dual $M^\dagger$ contains a vertex of infinite degree almost surely. Let $f\in F_{\mathrm{fin}}$ and let $\langle X_n\rangle_{n\geq 0}$ be a random walk on $M^\dagger$ started at f and stopped at τ_{F_∞}, the first time it visits F_∞. Then $\tau_{F_\infty} < \infty$ almost surely.*

Proof. If there are almost surely no edges $e^\dagger\in E^\dagger$ with both endpoints in F_{fin} then the result is trivial, so suppose not. Form a (possibly disconnected) graph $G^\ddagger$ with edge set $E^\dagger$ by, for each edge $e^\dagger\in E^\dagger$ with an endpoint in F_∞, replacing this endpoint with a vertex of degree one. For each edge $e^\dagger \in E^\dagger$, let $G^\ddagger_{e^\dagger}$ denote the connected component of $G^\dagger$ containing $e^\dagger$. The vertex set of $G^\ddagger$ may be written as $F_{\mathrm{fin}}\cup L$, where L is the set of degree one vertices of $G^\ddagger$ corresponding to edges of M incident to an infinite degree face. It is clear that the random edge-rooted graph $(G^\ddagger_\eta,\eta^\dagger)$ is reversible. Every connected component of $G^\ddagger$ contains a vertex in L and consequently,

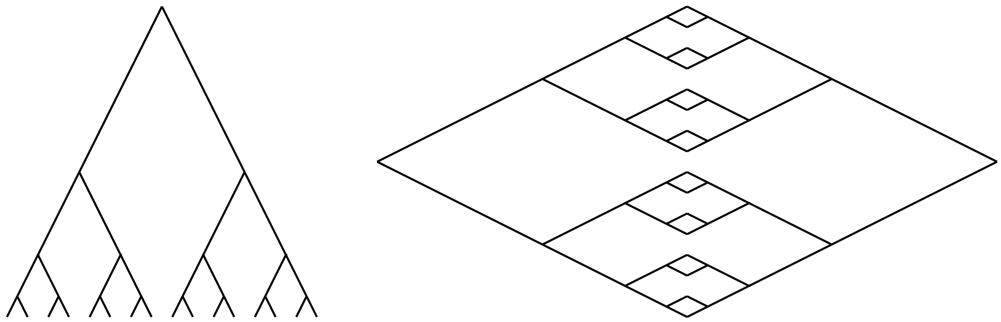

Figure 9.6: The maps T_4 (left) and M_4 (right).

by stationarity, a random walk on $G^\ddagger$ visits L infinitely often almost surely. We conclude by noting that the random walk on $M^\dagger$ started at a vertex $f \in F_{\text{fin}}$ and stopped when it first hits F_∞ can be coupled with the random walk on $G^\ddagger$ started at the same $f \in F_{\text{fin}}$ and stopped when it first hits L so that the two hitting times agree. $\square$

Proof of Proposition 9.5.18. Generate $\mathfrak{F}^\dagger$ with Wilson's algorithm rooted at $\{\infty\} \cup F_\infty$. It follows from Lemma 9.5.19 that for every face f of M, there exists a path in $\mathfrak{F}^\dagger$ connecting f to F_∞, and this path is easily seen to be unique. For each oriented edge e of M with $e^\ell \in F_{\text{fin}}$, transport mass one from e^- to d^-, where $d^\dagger$ is the last edge of the unique path connecting e^ℓ to F_∞ in $\mathfrak{F}^\dagger$. Then each vertex u sends a mass of at most $\deg(u)$, so that the expected mass sent by the root is finite. For each connected component of $\mathfrak{F}^\dagger \setminus F_\infty$, there exists an edge $e^\dagger \in \mathfrak{F}^\dagger$ that connects this infinite connected component to F_∞, and the vertex e^- receives mass equal to at least the number of edges in the component. Thus, the existence of an infinite connected component of $\mathfrak{F}^\dagger \setminus F_\infty$ would contradict the mass transport principle. $\square$

9.5.4 Finite expected degree is needed

Example 9.5.20. Let T_n be a binary tree of height n drawn in the plane. The Benjamini-Schramm limit of T_n as n tends to infinity is known as the canopy tree, and can be thought of as an 'infinite binary tree viewed from a leaf'. Let M_n be the finite map obtained by drawing two copies of T_n so that one is the reflection of the other, and attaching these two copies together at their leaves (see Figure 5). Then the Benjamini-Schramm limit of M_n exists and is formed of two canopy trees attached together at their leaves. Let T'_n be the map formed from T_n by replace each edge at distance k from the leaves of T_n by 3^k parallel edges, and let M'_n be the map obtained by drawing two copies of T'_n so that one is the reflection of the other, and attaching these two copies together at their leaves. These strong drifts ensure that, in the Benjamini-Schramm limit M' of the maps M'_n, the distance between a simple random walk $\langle X_k \rangle_{k\geq 0}$ and the set of former leaves of the canopy tree will increase linearly in k almost surely, and the walk therefore gets stuck in one of the two trees almost surely. Since the walk can get stuck in either of the two trees, we obtain a bounded harmonic function on M defined by letting $h(v)$ be the probability starting from v that the random walk gets stuck in the first tree. Moreover, Wilson's algorithm shows that the WUSF of M' has exactly two components, corresponding to the two trees of M'.

9.5.5 Percolation and minimal spanning forests

While the uniform spanning forests are related to random walks, the *minimal* spanning forests are related to bernoulli bond percolation. Recall that a **Bernoulli-p bond percolation** on G, denoted ω_p, is a random subgraph of G defined by keeping each edge of G independently with probability p and deleting the rest. The Bernoulli bond percolations $\{\omega_p\}_{p\in[0,1]}$ on a graph G may be coupled monotonically by letting $\{U(e)\}_{e\in E}$ be a collection of i.i.d. Uniform($[0,1]$) random variables indexed by the edge set of G and setting $\omega_p(e) = \mathbb{1}(U(e) \leq p)$ for every $e \in E$ and $p \in [0,1]$. The **critical probability** of G is defined by

$$p_c(G) := \inf\big\{p : \mathbb{P}(\omega_p \text{ has an infinite connected component}) = 1\big\}.$$

It is well-known [7, 187] that if (G,ρ) is an ergodic unimodular random rooted graph, then for each $p \in [0,1]$ the number of infinite connected components of ω_p for any is in $\{0,1,\infty\}$ almost surely and is non-random. Moreover, if (G,ρ) is a unimodular random rooted graph and p is such that ω_p has a unique infinite cluster almost surely, then $\omega_{p'}$ has a unique infinite connected component almost surely for every $p' \geq p$ [7, 113, 176]. In light of this, the **uniqueness threshold** of a graph G is defined to be

$$p_u(G) = \inf\{p : \mathbb{P}(\omega_p \text{ has a unique infinite connected component}) = 1\} \geq p_c(G).$$

Note that if (G,ρ) is an ergodic, infinite unimodular random rooted graph, then the quantities $p_c(G)$ and $p_u(G)$ are non-random. It is of interest to determine which graphs have a non-uniqueness phase for Bernoulli bond percolation. For example, $p_c = p_u$ for every amenable transitive graph, while a long standing conjecture of Benjamini and Schramm [51] asserts conversely that every nonamenable transitive graph has $p_c < p_u$.

Lyons, Peres and Schramm [175] related the non-uniqueness phase to minimal spanning forests. Given a finite graph G and an injective function $U : E(G) \to \mathbb{R}$ assigning **weights** to the edges of G, the **minimal spanning tree** of G with respect to U is defined to be the spanning tree T of G minimising the total weight $\sum_{e\in T} U(e)$. Equivalently, an edge e of G is contained in T if and only if there does not exist a simple cycle in G containing e such that e maximises $U(e)$ among the edges in this cycle. We write MST_G for the distribution on spanning trees of G obtained by letting T be the minimal spanning tree of G with respect to weights $\{U(e)\}_{e\in E}$ given by i.i.d. Uniform($[0,1]$) random variables. As for the uniform spanning trees, given an exhaustion $\langle V_n\rangle_{n\geq 0}$ of an infinite graph G, we the **free and wired minimal spanning forests** as the weak limits

$$\mathsf{FMSF}_G(S \subset T) := \lim_{n\to\infty} \mathsf{MST}_{G_n}(S \subset T)$$

and

$$\mathsf{WMSF}_G(S \subset T) := \lim_{n\to\infty} \mathsf{MST}_{G_n^*}(S \subset T).$$

The limits do not depend on the choice of exhaustion and, if (G,ρ) is a unimodular random rooted graph and $\mathfrak{F}$ is a sample of either WMSF_G or FMSF_G, then $\mathfrak{F}$ is a percolation on (G,ρ). Unlike in the uniform case, both of the minimal spanning forests may also be defined directly on the infinite graph G as follows. Let $\{U(e) : e \in E\}$ be a collection of i.i.d. Uniform($[0,1]$) random variables indexed by the edge set of G. An edge e of G is included in free minimal spanning forest of G if and only if it is not the heaviest edge in any simple cycle in G. An edge e of G is included in the wired minimal spanning forest of G if and only if it is not the heaviest edge in any simple cycle in G or in any bi-infinite simple path (or 'cycle through infinity') in G.

Lyons, Peres and Schramm [175] proved that an infinite connected graph G has $\mathsf{FMSF}_G = \mathsf{WMSF}_G$ if and only if for ω_p has a unique infinite cluster for Lebesgue-a.e. $p \in [0,1]$. Combining this with uniqueness monotonicity [7, 113, Theorem 6.7] yields the following.

Theorem 9.5.21 ([7, 113, 175]: Theorem 9.1.1, equivalence of (15) and (16))**.** *Let (G,ρ) be an infinite unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then $p_c(G) < p_u(G)$ if and only if $\mathsf{FMSF}_G \neq \mathsf{WMSF}_G$.*

The minimal spanning forests share several properties with their uniform cousins:

1. **(Free dominates wired.)** The measure FMSF_G stochastically dominates the measure WMSF_G for every graph G.

2. **(Domination and subgraphs)** let H be a connected subgraph of G. Then the FMSF of H stochastically dominates the restriction of the FMSF of G to H.

3. (**Expected degree of the WMSF.**) The expected degree in the WMSF of root of any unimodular random rooted graph is two [7, Proposition 7.3].

4. (**Amenability and boundary conditions.**) If (G,ρ) is an invariantly amenable random rooted graph, then $\mathsf{FMSF}_G = \mathsf{WMSF}_G$ almost surely [7, Proposition 18.14].

5. (**Planar duality**.) If M is a simply connected map with locally finite dual $M^\dagger$ and $\mathfrak{F}$ is a sample of FMSF_M, then $\mathfrak{F}^\dagger$ has law $\mathsf{WMSF}_{M^\dagger}$ [173, §11.5].

Combining items 1 and 2 above, we deduce that if (G,ρ) a unimodular random rooted graph, the measures FMSF_G and WMSF_G coincide almost surely if and only if the expected degree of ρ in the FMSF of G is two.

By using the properties above, the proof of Theorem 9.5.11 also yields a formula relating the expected degree of the FMSF to the average curvature.

Theorem 9.5.22 (Theorem 9.1.1, equivalence of (1) and (15))**.** *Let (M,ρ) be an infinite simply connected unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$ and let $\mathfrak{F}$ be a sample of FMSF_M.*

Then

$$\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] = 2 - \frac{1}{\pi}\mathbb{K}(M,\rho).$$

In particular, $\mathsf{FMSF}_M = \mathsf{WMSF}_M$ *almost surely if and only if* $\mathbb{K}(M,\rho) = 0$.

The equivalence of (16) and (2) in Theorem 9.1.1 can also be proven directly as follows. Let (M,ρ) be a hyperbolic simply connected unimodular random rooted map. If (M,ρ) has locally finite dual, we deduce that $p_c(M) < p_u(M)$ by applying the following two results.

Theorem 9.5.23 (Benjamini, Lyons, Peres and Schramm [36, 43]; Aldous and Lyons [7])**.** *Let (G,ρ) be an invariantly nonamenable unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then ω_{p_c} does not contain any infinite connected components almost surely.*

Theorem 9.5.24 (Benjamini and Schramm [49, Theorem 3.1])**.** *Let (M,ρ) be an invariantly nonamenable, simply connected unimodular random rooted map with locally finite dual $M^\dagger$ and suppose that $\mathbb{E}[\deg(\rho)] < \infty$. Then ω_p has a unique infinite connected component if and only if every component of $\omega_p^\dagger$ is finite. It follows that*

$$p_u(M) = 1 - p_c(M^\dagger)$$

almost surely and that ω_{p_u} contains a unique infinite connected component almost surely.

Benjamini and Schramm proved their theorem for transitive planar graphs, but their proof extends immediately to our setting.

If M does not have locally finite dual, then it must have infinitely many infinite faces by Lemma 9.3.15, so that the underlying graph of M is infinitely ended. In this case, we have that $p_u(M) = 1$ almost surely (see Proposition 9.7.7), while $p_c(M) < 1$ by Theorem 9.5.23.

Thus, we have the following.

Corollary 9.5.25 (Theorem 9.1.1, (2) implies (16))**.** *Let (M,ρ) be a simply connected, invariantly nonamenable, unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$. Then $p_c(M) < p_u(M)$ almost surely.*

9.5.6 Expected degree formula

We prove Theorem 9.5.22. Every property of the minimal spanning forests that we use also holds for the uniform spanning forests, so that we also obtain an alternative proof of Theorem 9.5.11 that does not rely on Theorem 9.5.13.

Proof of Theorem 9.5.11.

Locally finite dual case. Let ω be a percolation on (M,ρ), and let $\omega^\dagger = \{e^\dagger \in E^\dagger : e \notin \omega\}$. As in Section 9.2.5, let η be chosen uniformly at random from the set $E^{\rightarrow}_\rho$ of oriented edges of M emanating from ρ, let $\rho^\dagger = \eta^r$, and let $\mathbb{P}^{\text{rev}}$ be the $\deg(\rho)$-biasing of $\mathbb{P}$ and let $\mathbb{P}^\dagger$ be the $\deg(\rho^\dagger)^{-1}$

biasing of $\mathbb{P}^{\mathrm{rev}}$, so that $(M^\dagger, \rho^\dagger)$ is a unimodular random rooted map under $\mathbb{P}^{\mathrm{rev}}$. We write $\mathbb{E}^\dagger$ for the expectation operator associated to $\mathbb{P}^\dagger$.

Lemma 9.5.26. *The following equation holds.*

$$\mathbb{E}[\deg_\omega(\rho)] = \mathbb{E}[\deg(\rho)] - \mathbb{E}\left[\sum_{f \perp \rho} \deg(f)^{-1}\right]\mathbb{E}^\dagger[\deg_{\omega^\dagger}(\rho^\dagger)]. \tag{9.5.2}$$

Proof. Observe that, since η is uniformly distributed on $E^{\rightarrow}_\rho$ conditional on (M, ρ), we have

$$\mathbb{E}[\deg_\omega(\rho)] = \mathbb{E}[\deg(\rho)\mathbb{1}(\eta \in \omega)] = \mathbb{E}\left[\deg(\rho)\left(1 - \mathbb{1}(\eta^\dagger \in \omega^\dagger)\right)\right],$$

and so

$$\mathbb{E}[\deg_\omega(\rho)] = \mathbb{E}[\deg(\rho)]\left(1 - \mathbb{P}^{\mathrm{rev}}\left(\eta^\dagger \in \omega^\dagger\right)\right).$$

Similarly, since under the measure $\mathbb{P}^\dagger$ and conditional on $(M^\dagger, \rho^\dagger)$, $\eta^\dagger$ is uniformly distributed on $E^{\rightarrow}_{\rho^\dagger}$,

$$\mathbb{E}^\dagger[\deg_{\omega^\dagger}(\rho^\dagger)] = \mathbb{E}^\dagger[\deg(\rho^\dagger)]\mathbb{P}^{\mathrm{rev}}\left(\eta^\dagger \in \omega^\dagger\right).$$

It follows that

$$\mathbb{E}[\deg_\omega(\rho)] = \mathbb{E}[\deg(\rho)]\left(1 - \frac{\mathbb{E}^\dagger[\deg(\rho^\dagger)]}{\mathbb{E}[\deg(\rho)]}\mathbb{E}^\dagger[\deg_{\omega^\dagger}(\rho^\dagger)]\right).$$

Applying the expected degree formula (9.2.3), we deduce that

$$\mathbb{E}[\deg_\omega(\rho)] = \mathbb{E}[\deg(\rho)]\left(1 - \frac{\mathbb{E}\left[\sum_{f\perp\rho}\deg(f)^{-1}\right]}{\mathbb{E}[\deg(\rho)]}\mathbb{E}^\dagger[\deg_{\omega^\dagger}(\rho^\dagger)]\right),$$

which rearranges to give the desired expression. □

Let $\mathfrak{F}$ have law FMSF_M. By Proposition 9.5.9, the dual forest $\mathfrak{F}^\dagger$ is distributed according to $\mathsf{WMSF}_{M^\dagger}$. Since the expected degree at the root of the WMSF in any unimodular random rooted graph is 2, we have $\mathbb{E}^\dagger[\deg_{\mathfrak{F}^\dagger}(\rho^\dagger)] = 2$ and consequently, by Lemma 9.5.26,

$$\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] = \mathbb{E}[\deg(\rho)] - 2\mathbb{E}\left[\sum_{f\perp\rho}\deg(f)^{-1}\right] = \mathbb{E}\left[\sum_{f\perp\rho}\frac{\deg(f)-2}{\deg(f)}\right] = 2 - \frac{1}{\pi}\mathbb{K}(M,\rho).$$

This completes the proof in the case that the dual of M is locally finite.

Non-locally finite dual case. Observe that, if ω is a percolation on a unimodular random map (M, ρ) that is almost surely a spanning forest of M, then every component of $\omega^\dagger$ is infinite, and it follows from [7, Theorem 6.1] that $\mathbb{E}^\dagger[\deg_{\omega^\dagger}(\rho^\dagger)] \geq 2$. Thus, we deduce from Lemma 9.5.26 that

$$\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] \leq 2 - \frac{1}{\pi}\mathbb{K}(M,\rho). \tag{9.5.3}$$

For each $n \geq 1$, let M_n be the locally finite map containing M as a submap that was defined in the proof of Proposition 9.4.5, and let $\mathfrak{F}_n$ be a sample of FMSF_{M_n}. Since (M_n,ρ) has locally finite dual, we have that

$$\mathbb{E}[\deg_{\mathfrak{F}_n}(\rho)] = 2 - \frac{1}{\pi}\mathbb{K}(M_n,\rho) = 2 - \frac{1}{\pi}\mathbb{K}(M,\rho),$$

where the second equality follows from Proposition 9.4.5.

Since the underlying graph of M is a subgraph of the underlying graph of M_n, the forest $\mathfrak{F}$ stochastically dominates the restriction $\mathfrak{F}_n \cap E$ for every $n \geq 1$, and so

$$\begin{aligned}\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] &\geq \mathbb{E}[\deg_{\mathfrak{F}_n}(\rho)] - \mathbb{E}[\deg_{M_n}(\rho) - \deg_M(\rho)] \\ &= 2 - \frac{1}{\pi}\mathbb{K}(M,\rho) - \mathbb{E}[\deg_{M_n}(\rho) - \deg_M(\rho)] \xrightarrow[n\to\infty]{} 2 - \frac{1}{\pi}\mathbb{K}(M,\rho).\end{aligned}$$

(The limit can be deduced either by direct calculation or by invoking the dominated convergence theorem.) To obtain a corresponding upper bound, consider $\mathfrak{F}$ as a subgraph of M_1. Since $\mathfrak{F}$ is a spanning forest of M_1, we can apply (9.5.3) to deduce that

$$\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] \leq 2 - \frac{1}{\pi}\mathbb{K}(M_1,\rho) = 2 - \frac{1}{\pi}\mathbb{K}(M,\rho)$$

completing the proof. □

9.6 The conformal type

Given a map M such that every face of M has degree at least three, we may form a surface $S(M)$ by gluing regular unit polygons together according to the combinatorics of M, the boundaries of these polygons becoming the edges of M embedded in $S(M)$. As before, we consider the upper half-space $\{x+iy \in \mathbb{C} : y > 0\}$ with edges $\{[n,n+1] : n \in \mathbb{Z}\}$ to be a regular ∞-gon. The surface $S(M)$ can be endowed naturally with a conformal structure by defining an atlas as follows.

- For each face f of M, we take as a chart the identity map from the interior of the regular polygon corresponding to f to itself.
- For each edge e of M, we define an open neighbourhood of the interior of e in S by taking the two triangles formed by the endpoints of e and the centres of the two faces adjacent to e (if either face is infinite, we interpret this triangle to be the infinite strip starting at e and perpendicular to the boundary of the face). To define a coordinate chart on this neighbourhood, we simply place the two triangles next to each other in the plane. This chart

is well defined even if both sides of e are incident to the same face.

- For each vertex v of M, we define an open neighbourhood of v in S similarly by intersecting the corners of the faces adjacent to v with open discs of radius $1/2$ centred at v. We define a chart on this neighbourhood by first laying the corners out in a (possibly overlapping) spiral around the origin, and then applying the function $z \mapsto z^{2\pi/\theta(v)}$, suitably interpreted to get an injective map into the plane.

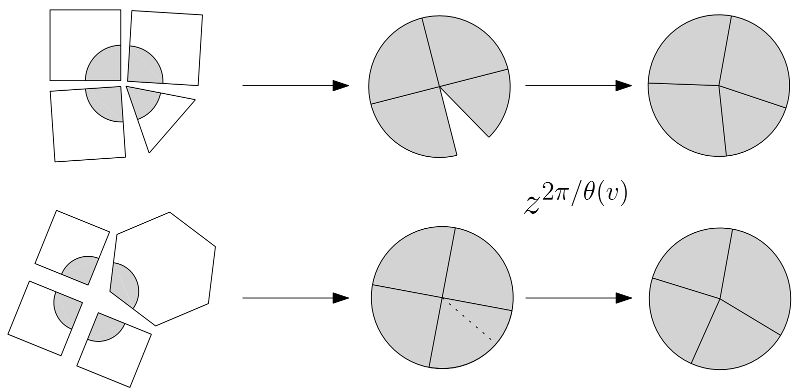

The coordinate changes are easily seen to be analytic, so that this atlas does indeed define a Riemann surface structure on $S(M)$. We denote this Riemann surface by $\mathcal{R}(M)$.

The definition of $\mathcal{R}(M)$ can be extended (somewhat arbitrarily) to maps containing faces of degree 1 and 2 as follows: Given a map M, let $\hat{M}$ be obtained from M by adding a vertex inside each face of M that has degree 1 or 2, and connecting this vertex to each of the corners of the face. Every face of $\hat{M}$ has degree at least three, and we define $\mathcal{R}(M) = \mathcal{R}(\hat{M})$. The map M can be embedded in $\mathcal{R}(M)$ by restricting the natural embedding of $\hat{M}$ into $\mathcal{R}(\hat{M})$.

If M is simply connected, the uniformization theorem implies that $\mathcal{R}(M)$ is conformally equivalent to the sphere, the plane or the disc, and we call M **conformally elliptic, parabolic, or hyperbolic** accordingly. We refer to the embedding of M into the sphere, the plane, or the disc given by uniformizing $\mathcal{R}(M)$, which is unique up to Möbius transformations of the sphere, plane or disc as appropriate, as the **conformal embedding** of M. Conformal embeddings of unimodular random planar maps are conjectured to play a key role in the theory of two-dimensional quantum gravity, see [74] and references therein.

Gill and Rohde [103] proved that every Benjamini-Schramm limit of finite planar maps with uniformly bounded codegrees is conformally parabolic almost surely. The main result of this section generalises and, together with Corollary 9.3.6, provides a converse to their result.

Theorem 9.6.1 (Theorem 9.1.1, equivalence of (9) and (2))**.** *Let (M, ρ) be an infinite, ergodic, simply connected unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$. Then M is conformally parabolic almost surely if and only if (M, ρ) is hyperfinite.*

Proof of Theorem 9.6.1: Conformally parabolic implies hyperfinite. Let (M, ρ) be a unimodular random rooted map such that M is conformally parabolic a.s., and let z be the conformal embedding of M into the plane, which is unique up to translation and scaling. We claim that (M, ρ) is hyperfinite.

We claim that if H and H' are two open half-planes that are disjoint from $z(V)$, then either H is contained in H', H' is contained in H, or H and H' are disjoint. Indeed, if this is not the case, then the set $z(V)$ is contained in some cone. It follows that there exists a unique, not necessarily distinct, pair of angles $\theta_1, \theta_2 \in [0, 2\pi)$ with $|\theta_2 - \theta_1| \bmod 2\pi < \pi$ such that for every $\varepsilon > 0$ and every point $z_0 \in \mathbb{C}$, the cone

$$\{z \in \mathbb{C} : \arg(z - z_0) \in [\theta_1 - \varepsilon, \theta_2 + \varepsilon]\}$$

contains all but finitely many points of $z(V)$, but, if $\theta_1 \neq \theta_2$, then infinitely many points of $z(V)$ are not in the cone

$$\{z \in \mathbb{C} : \arg(z - z_0) \in [\theta_1 + \varepsilon, \theta_2 - \varepsilon]\}.$$

In particular, there exists a finite, nonempty set of vertices of $z(V)$ that have maximally negative inner product with $e^{i\theta_1}$. This set of vertices does not depend on the choice of conformal embedding, and so we obtain a contradiction by transporting a mass of one from each vertex of M to a uniformly chosen element of this set.

If there are a.s. two disjoint open half-planes that are both disjoint from $z(V)$, let H_1 and H_2 be the two disjoint open half-planes disjoint from $z(V)$ that are maximal in the sense that any open-half-plane strictly containing either H_1 or H_2 must intersect $z(V)$. By translating, rotating, and scaling, we may assume that H_1 is the lower half-plane $\{z \in \mathbb{C} : \Im(z) < 0\}$ and H_2 is the half-plane $\{z \in \mathbb{C} : \Im(z) > 1\}$. Let U be a uniform $[0, 1]$ random variable, and consider the randomly shifted real integers $\mathbb{Z} + U \subset \mathbb{C}$. For each $n \in \mathbb{Z}$, let $v(n) \in V$ be chosen uniformly at random from among those $v \in V$ such that $z(v)$ is of minimal distance to $n + U$. Drawing an edge between $v(n)$ and $v(n + 1)$ for every $n \in \mathbb{Z}$ defines a unimodular coupling between (M, ρ) and $\mathbb{Z}$, and it follows from Proposition 9.3.13 that (M, ρ) is hyperfinite in this case.

Now suppose that there is a unique open half-plane H disjoint from $z(V)$ that is maximal in the sense that any open-half-plane strictly containing H must intersect $z(V)$. By rotating and translating, we may assume that H is the right half-plane, $H = \{x + iy \in \mathbb{C} : x > 0\}$. We will use the linear structure of the boundary of H to show that (M, ρ) is coupling equivalent to $\mathbb{Z}$, so that it will follow from Proposition 9.3.13 that (M, ρ) is hyperfinite. Since $z(V)$ is locally finite, there exists a vertex v such that the straight line $\{x + z(v) : x \geq 0\}$ is disjoint from $z(V) \setminus \{z(v)\}$. For any such vertex, again by local finiteness of $z(V)$, there exists a positive angle θ such that the cone $C_\theta(v) = \{z \in \mathbb{C} : \arg(z - z(v)) \in [-\theta, \theta]\}$ is disjoint from $C_\theta(u)$ for every $u \in V$ with $\Re z(u) \geq \Re z(v)$. We say that v is θ-**exposed** if this condition holds. By the above discussion, we may choose $\theta > 0$ such that ρ is θ-exposed with positive probability, in which case it follows from ergodicity and unimodularity that there are infinitely many θ-exposed vertices a.s. Observe that, again using the fact that $z(V)$ is locally finite, for every bounded interval $[a, b]$, there exist at most finitely many θ-exposed vertices v of M that have $\Im z(v) \in [a, b]$. Define a graph whose vertices are the θ-exposed vertices of M, and where we draw an edge between two distinct θ-exposed vertices u and v if and and only if there is no θ-exposed vertex w, distinct from u and v, such that $\Im z(w)$ lies between $\Im z(u)$ and $\Im z(v)$. The above discussion implies that this graph is isomorphic to $\mathbb{Z}$, and it follows that (M, ρ) and $\mathbb{Z}$ are coupling equivalent as claimed.

Now suppose that $z(V)$ intersects every open half-plane. We define D to be the Delaunay tessellation with vertex set given by the set of points $z(V)$. That is, D is the map that has embedded vertex set $z(V)$, and has a straight line between two points $z(u)$ and $z(v)$ if and only if there exists either a disc or a half-plane in the plane such that

1. the disc or half-plane does not have any points of $z(V)$ in its interior,
2. $z(u)$, $z(v)$, and at least one other point of $z(V)$ lie on the boundary of the disc or half-plane, and
3. $z(u)$ and $z(v)$ are adjacent in the cyclic ordering of the set of points of $z(V)$ that intersect the boundary of the disc or half-plane.

Note that the isomorphism class of (D, ρ) is independent of the choice of the conformal embedding.

If the dual of D is not locally finite, observe that the union of the finite degree faces of D is equal to the convex hull of the set $z(V)$, and consequently any infinite face of D must contain a half-space disjoint from $z(V)$. On the other hand, if D is not locally finite, it follows from the definition of D that for every infinite degree vertex v of D, there exists an infinite sequence of vertices u_i and closed discs or half-planes C_i such that C_i contains both $z(v)$ and $z(u_i)$ in its boundary and no points of $z(V)$ in its interior. By taking a subsequential limit, it follows from the fact that $z(V)$ is locally finite that there must exist a half-space containing $z(v)$ in its boundary and no points of $z(V)$ in its interior.

Thus, we may assume that D is finite and has locally finite dual a.s. In this case, it follows from the measurability of conformal embedding that (D, ρ) is a unimodular random rooted map, and that (G, ρ) and (D, ρ) are coupling equivalent. Thus, by Proposition 9.3.13, it suffices to prove that (D, ρ) is hyperfinite. Define a mass transport as follows. For each vertex u and face f of D incident to u, let $\operatorname{ang}(f, u)$ be the angle of the corner of f at u. Transport a mass of $\operatorname{ang}(f, u)/\deg(f)$ from u to each of the vertices incident to f, including u itself. The mass sent out by each vertex u is 2π, while, since each face is a polygon in the plane, the mass received is

$$\sum_{f \perp u} \frac{\text{sum of internal angles of } f}{\deg(f)} = \sum_{f \perp u} \frac{\deg(f) - 2}{\deg(f)} \pi.$$

Applying the mass-transport principle yields that the average curvature of (D, ρ) is zero, and it follows from the equivalence of items (1) and (2) of Theorem 9.1.1 that (D, ρ) is hyperfinite as claimed. $\square$

Before proving the converse, our immediate goal is to prove the following lemma. Suppose M is a conformally hyperbolic map. For each oriented edge e of M, let U_e be the subset of $\mathcal{R}(M)$ defined as follows. If the face e^r to the right of e has finite degree, let U_e be the quadrilateral in the polygon corresponding to f that is formed from the first half of e, the first half of $\sigma(e)$ (i.e., the next edge clockwise from e in the cyclic ordering of the edges emanating from e^-), and the straight lines between the center of f and the midpoints of these two edges. If $\deg(f)$ is infinite, let U_e be the infinite strip in the half-plane corresponding to f whose boundary is given by the first half of e, the first half of $\sigma(e)$, and the two half-infinite straight lines that start at the midpoints of these edges and are perpendicular to the boundary of f. For each vertex v of M, we define

$$U_v = \bigcup_{e^- = v} U_e.$$

Given a subset K of a Riemann surface that is conformally equivalent to the hyperbolic plane, we write $\mathrm{area}_{\mathbb{H}}(K)$ of the hyperbolic area of the image of K under a uniformizing map from the Riemann surface to the hyperbolic plane.

Lemma 9.6.2. *There exists a constant C such that the following holds. Let M be a conformally hyperbolic, simply connected map such that every face of M has degree at least three. Then for every vertex v of M,*

$$\mathrm{area}_{\mathbb{H}}(U_v) \leq C \deg(v).$$

Proof. We shall require the following classical facts about hyperbolic area.

1. **Schwarz-Pick**. If $\mathcal{R}$ is a Riemann surface that is conformally equivalent to the disc, D is a simply connected open subset of $\mathcal{R}$, and K is a Borel subset of D, then the hyperbolic area of K considered as a subset of D is greater than or equal to the hyperbolic area of K considered as a subset of $\mathcal{R}$. (Schwarz-Pick is usually stated in terms of the metrics, but immediately implies this statement.)

2. For every $0 < \varepsilon < 1$, the hyperbolic area of the set $\{z \in \mathbb{C} : |z| < 1 - \varepsilon\}$, considered as a subset of the open unit disc, is given by

$$\frac{4\pi(1-\varepsilon)^2}{1-(1-\varepsilon)^2} \leq \frac{2\pi}{\varepsilon}. \tag{9.6.1}$$

For each oriented edge e, let $U_e^1, \ldots, U_e^4$ be the following subsets of the polygon associated to the face e^r. See Figure 9.7 for an illustration.

1. U_e^1 is the $1/8$ neighbourhood of the midpoint of e.

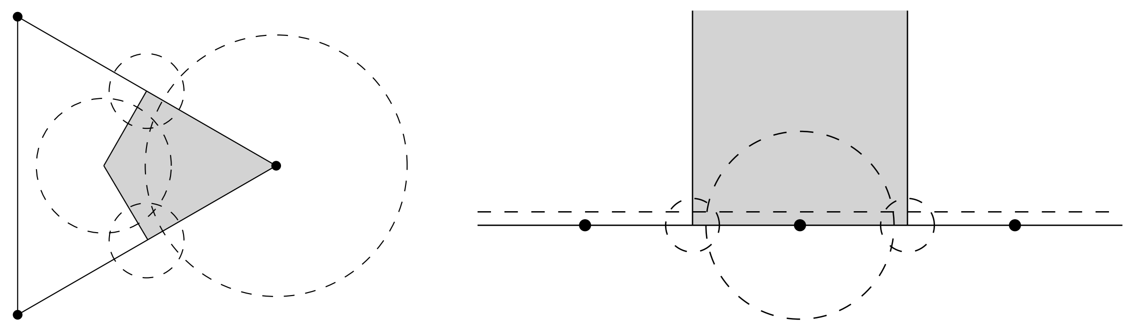

Figure 9.7: The covering of U_e (grey) by discs (dashed boundaries) used in the proof of Lemma 9.6.2. Left: the case that e^r has degree three. Right: the case that e^r has infinite degree.

2. U_e^2 is the 1/8 neighbourhood of the midpoint of the edge following e in the clockwise order of the edges adjacent to e^-.

3. U_e^3 is the 7/16 neighbourhood of e^-.

4. If e^r has finite degree, we define U_e^4 is the intersection of U_e with the disc that is centred at the centre of the polygon corresponding to e^r and that has distance 1/16 from the boundary of the polygon corresponding to e^r. If e^r has infinite degree, we let U_e^4 be the intersection of U_e with the upper half-plane that has distance 1/16 from the boundary of the half-plane corresponding to e^r.

It is easily verified by elementary trigonometry that the U_e is contained in the union $\bigcup_{i=1}^4 U_e^i$.

We first claim that the hyperbolic areas of U_e^1 and U_e^2 are bounded above by a constant. Recall that in the chart at e, we simply place the two triangles formed by the endpoints of e and the centres of the two faces adjacent to e next to each other (where the triangles become infinite strips if a face has infinite degree). It is easily verified that the open ball of radius 1/4 around the midpoint of e is always contained in the domain formed by placing the two triangles together. Since U_e^1 is contained in the ball of radius 1/8 around the midpoint of e, it follows that the hyperbolic area of U_e^1 is bounded above by a constant (namely, $4\pi/3$) by (9.6.1) and Schwarz-Pick. The corresponding claim for U_e^2 follows similarly.

We next claim that the hyperbolic area of U_e^4 is bounded above by a constant. By Schwarz-Pick, it suffices to prove that the hyperbolic area of U_e^4 is bounded above by a constant when considered as a subset of the polygon corresponding to e_r. If e_r has infinite degree, this area is the same as the hyperbolic area of $\{x + iy \in \mathbb{C} : x \in [0,1], y \geq 1/16\}$ considered as a subset of the upper half-plane $\{x + iy \in \mathbb{C} : y > 0\}$, which is finite (in fact, it is equal to 16). On the other hand, if e_r has finite degree, then we can consider the disc D centred at the centre of the polygon corresponding to e_r that has distance 1/16 to the boundary of the polygon. The radius of the disc D' that is centred at the centre of the polygon and just touches the boundary of the polygon is

$$\frac{1 + \cos(2\pi/\deg(e^r))}{4\sin(\pi/\deg(e^r))} \leq \frac{\deg(e^r)}{2\pi}.$$

Thus, it follows that the hyperbolic area of D considered as a subset of D' is at most $16\deg(e^r)$. We deduce that, by symmetry, the area of U_e^4 considered at a subset of D' is at most 16. The claim now follows by applying Schwarz-Pick.

Finally, we claim that for every vertex v, the hyperbolic area of $U_v^3 := \bigcup_{e^-=v} U_e^3$ is at most $C\deg(v)$, where C is a universal constant. Observe that U_v^3 is simply the set of points that are at distance at most $7/16$ from v in one of the polygons corresponding to a face incident to v. Under the chart associated to v, this set gets mapped to a ball of radius $(7/16)^{2\pi/\theta(v)}$, while the set of all points at distance at most $1/2$ from v gets mapped to the ball of radius $(1/2)^{2\pi/\theta(v)}$ with the same centre. It follows that the hyperbolic area of U_v^3 considered as a subset of this larger ball is at most

$$\frac{2\pi}{1-(7/8)^{2\pi/\theta(v)}} \leq C\theta(v) \leq \pi C\deg(v),$$

where C is a constant. Verifying that such a constant exists is a simple calculus exercise. The claim now follows from Schwarz-Pick.

The lemma follows by combining the three estimates given. □

Proof of Theorem 9.6.1: Conformally hyperbolic implies invariantly nonamenable. We may assume that every face of M has degree at least three, since otherwise the map $\hat{M}$ used to define $\mathcal{R}(M)$ can easily be rerooted and biased to be unimodular, and is coupling equivalent to M. Let ϕ be a conformal equivalence between $\mathcal{R}(M)$ and the hyperbolic plane, let Z be a Poisson point process of intensity 1 on the hyperbolic plane, and let D be the Delaunay triangulation associated to Z.

We form a graph G whose vertex set is $V \cup Z$, and has as edges the edges of M, the edges of D, and an edge connecting each $v \in V$ to every point $z \in Z \cap \phi(U_v)$. We also mark the edges of G according to which of these three types they come from. Note that the law of (G,ρ) does not depend on the choice of ϕ. It follows from Lemma 9.6.2 that the expected number of points in $Z \cap \phi(U_\rho)$ is finite. It is easily verified, using the measurability of the conformal embedding, that if we sample (G,ρ) biased by $1+|Z\cap\phi(U_\rho)|$ and then let $\hat\rho$ be uniform on the set $\{\rho\}\cup(Z\cap\phi(U_\rho))$, then the resulting random rooted network $(G,\hat\rho)$ is unimodular. Similarly, if we sample (M,ρ) and Z biased by $|Z\cap\phi(U_\rho)|$, and let ρ' be uniform on the set $Z\cap\phi(U_\rho)$, then the resulting graph (D,ρ') is unimodular. Thus, we have defined a unimodular coupling between (M,ρ) and (D,ρ'). Since D is a Poisson-Delaunay triangulation of the hyperbolic plane, (D,ρ') is invariantly nonamenable (see [37] and [21, Section 2]), and we conclude by applying Proposition 9.3.13. □

We remark that there are many other natural (and inequivalent) ways to associate Riemann surfaces to maps. For example, we could associate to each face of M a disc of circumference k, with boundary split into k arcs of length one corresponding to the edges, and glue adjacent faces according to arc length along their shared edges. The proof of Theorem 9.6.1 extends to many of these alternatively defined conformal embeddings with only minor modifications.

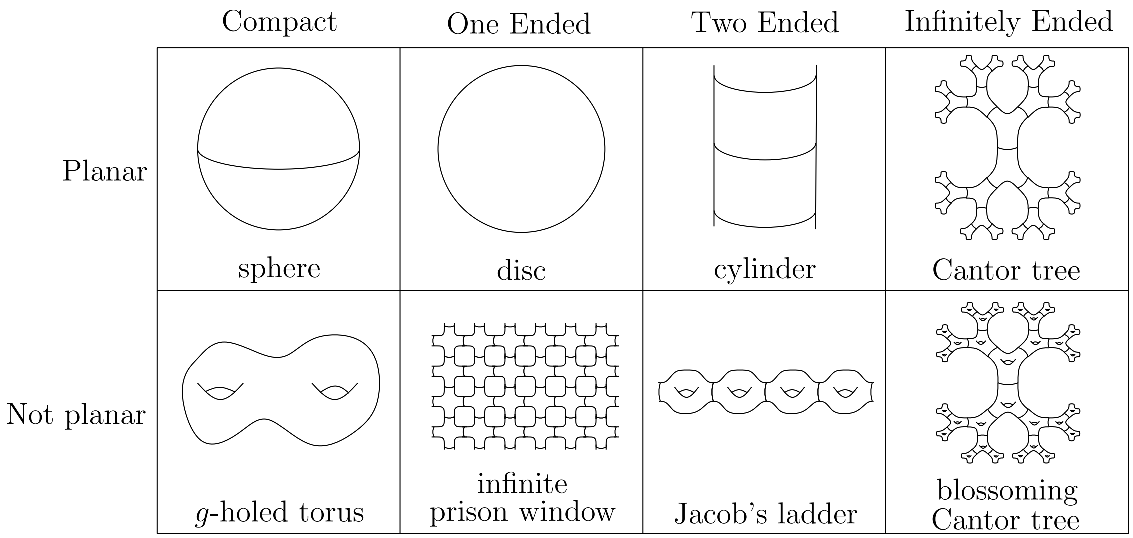

Figure 9.8: Possible topologies of a unimodular random map. The surface $S(M)$ associated to a unimodular random map (M, ρ) is almost surely homeomorphic to one of the above.

9.7 Multiply-connected and non-planar maps

9.7.1 The topology of unimodular random rooted maps.

In this section we study multiply-connected unimodular random rooted maps. We begin by classifying the possible topologies of the surface associated to a unimodular random rooted map (M, ρ). Biringer and Raimbault [54] classified the possible topologies of unimodular random rooted complete, orientable, hyperbolic surfaces. Their methods readily generalise to our setting, yielding the following theorem. In fact, the proof is slightly less technical in our setting, and we provide a quick sketch below.

Theorem 9.7.1 (Topology of unimodular random rooted maps)**.** *Let (M, ρ) be an infinite unimodular random rooted map. Then the surface associated to M is almost surely homeomorphic to one of the following surfaces: the plane, the cylinder, the Cantor tree, the infinite prison window, Jacob's ladder, or the blossoming Cantor tree.*

Here, the **Cantor tree** is a 'tree made of tubes' that is homeomorphic to the complement of the Cantor set in the sphere, the **infinite prison window** is 'the lattice $\mathbb{Z}^2$ made of tubes', **Jacob's ladder** is 'an infinite ladder made of tubes', and the **blossoming Cantor tree** is a Cantor tree with a handle attached near each bifurcation. See Figure 4 for illustrations. Be warned that homeomorphism is an extremely weak notion here. For example, 'the lattices $\mathbb{Z}^d$ made of tubes' are all homeomorphic to each other for all $d \geq 2$, as indeed are any two 'one-ended infinite graphs made of tubes'.

In [54], Biringer and Raimbault must also allow for the surfaces above to be punctured at a locally finite set of points, corresponding to isolated ends of the surface. This does not occur in our setting, as the surfaces corresponding to unimodular random rooted maps do not have isolated ends.

Sketch of proof. An **end** ξ of an infinite graph G is a function that assigns a connected component ξ_K of $G \setminus K$ to each finite set of vertices K of G, and satifies the consistency condition that $\xi_{K'} \subseteq \xi_K$ whenever $K' \supseteq K$. The **space of ends** of G, denoted $\partial\mathscr{E}(G)$ is the topological space with the set of ends of G as its underlying set and with a basis of open sets given by sets of the form $\{\xi \text{ an end of } G : \xi_K = W\}$, where $K \subset V$ is finite and $W \subset V$ is a connected component of $G \setminus K$. Note that the basis sets are also closed, so that the space of ends is always *zero-dimensional*, that is, its topology is induced by a basis of sets that are both open and closed. The space of ends of a surface is defined similarly, replacing instances of the word 'finite' by 'compact' above, and are also zero-dimensional.

It is well-known that every unimodular random rooted graph either has one, two, or infinitely many ends, and, in the last case, the space of ends does not have any isolated points [7, Proposition 6.10]. Since the space of ends of any graph is also compact, it follows in the last case that the space of ends is homeomorphic to the Cantor set (which, by Brouwer's Theorem [61], is the only compact, zero-dimensional Hausdorff space with no isolated points). A similar proof applies to show that if (M, ρ) is a random rooted map with associated surface $S = S(M)$, then S has either one, two, or infinitely many ends and in the last case the space of ends of S is homeomorphic to a Cantor set. Next, standard mass transport arguments show that if S contains handles, then the handles of S accumulate towards every end of S. That is, if S has handles then, for every compact subset K of S, every non-precompact connected component of $S \setminus K$ contains a handle.

We next apply the classification theorem for non-compact surfaces due to Kerékjártó and Richards [199], which states that if two non-compact orientable surfaces S_1 and S_2 have the same number of handles (which in our case will be zero or infinity) and there exists a homeomorphism $\phi : \partial\mathscr{E}(S_1) \to \partial\mathscr{E}(S_2)$ such that the handles of S_1 accumulate to $\xi \in \partial\mathscr{E}(S_1)$ if and only if the handles of S_2 accumulate to $\phi(\xi)$, then S_1 and S_2 are homeomorphic and ϕ extends to a homeomorphism from the ends compactification of S_1 to the ends compactification of S_2. Thus, by the above discussion, the homeomorphism class of $S = S(M)$ is determined almost surely by its number of ends and by the existence or non-existence of handles. This yields the six different possibilities listed in the statement of the theorem (see Figure 6). $\square$

Example 9.7.2. The product $T_3 \times H$ of the 3-regular tree T_3 and the graph H consisting of 3 parallel edges between two vertices is planar but cannot be drawn in the plane without accumulation points. By letting the cyclic ordering of the edges emanating from each vertex of $T_3 \times H$ alternate between T_3-edges and H-edges, and making the cyclic ordering of the edges emanating from $(v, 0)$ and $(v, 1)$ reflections of each other, we obtain a transitive quadrangulation of the Cantor tree.

The main result of this section is that the average curvature of a unimodular random rooted map restricts the possible topologies of the map.

Theorem 9.7.3 (Topology from average curvature)**.** *Let (M, ρ) be an ergodic unimodular random map. Then the almost sure conformal type of M is determined by its average curvature: Either*

1. *The average curvature of (M,ρ) is positive, in which case M is conformally elliptic and $S(M)$ is homeomorphic to the sphere almost surely,*
2. *the average curvature of (M,ρ) is zero, in which case M is conformally parabolic and $S(M)$ is homeomorphic to the plane, the cylinder, or the torus almost surely,*

or else

(3) *the average curvature of (M,ρ) is negative, in which case M is conformally hyperbolic $S(M)$ is homoemorphic to either the plane, the blossoming Cantor tree, Jacob's ladder, the infinite prison window, or a compact surface of genus at least two almost surely.*

The theorem will follow by combining Theorem 9.6.1, Theorem 9.7.1 and the notion of the universal cover of a map.

Proof. Recall that a surjective, holomorphic function $\Pi : S \to S'$ between two Riemann surfaces is a **holomorphic covering** if it is locally a homeomorphism, that is, if for every $x \in S$ there exists an open neighbourhood U of S such that the restriction of Π to U is a homeomorphism between U and its image. Given a Riemann surface S, the **universal cover** of S is a simply connected Riemann surface $\tilde{S}$ together with a covering $\Pi : \tilde{S} \to S$. The universal cover exists for any S, and is unique in the sense that if $\Pi' : \tilde{S}' \to S$ is another simply connected Riemann surface covering S, then there exists a conformal equivalence $\phi : \tilde{S}' \to \tilde{S}$ such that $\Pi' = \Pi \circ \phi$.

The universal cover of a map is defined analogously. Given a pair of maps $M = (G,\sigma)$ and $M' = (G',\sigma')$, we say that a graph homomorphism $\phi : G \to G'$ is a **map homomorphism** if $\sigma' \circ \phi = \phi \circ \sigma$, and say that ϕ is a **covering** if for every vertex v and every face f of M, the restriction of ϕ to each of $\{e \in E^{\rightarrow} : e^- = v\}$ and $\{e \in E^{\rightarrow} : e^r = f\}$ is injective. The **universal cover** of a map M is a simply connected map $\tilde{M}$ together with a covering $\pi : \tilde{M} \to M$. Every map has a universal cover, and the universal cover of a map M is unique in the sense that if $\pi' : \tilde{M}' \to M$ is a covering from a simply connected map $\tilde{M}'$ to M, there exists an isomorphism of maps $\phi : \tilde{M} \to \tilde{M}'$ such that $\pi' \circ \phi = \pi$.

The universal cover $\pi : \tilde{M} \to M$ of M may be constructed by taking every lift of every edge of M in the Riemann surface $\mathcal{R}(M)$ to the universal cover $\Pi : \tilde{\mathcal{R}}(M) \to \mathcal{R}(M)$ of $\mathcal{R}(M)$ (see e.g. [119, p. 60] for the topological notion of path lifting). In particular, if $\mathcal{R}(\tilde{M})$ is the Riemann surface associated to $\tilde{M}$, then there exists a conformal equivalence $\Phi : \mathcal{R}(\tilde{M}) \to \tilde{\mathcal{R}}(M)$ such that $\Pi \circ \Phi \circ \tilde{z} = z \circ \pi$. (See e.g. [215, Section 9.2] for a direct construction.) As for Riemann surfaces, the universal cover of a map is easily seen to be unique in the sense that if $\pi' : \tilde{M}' \to M$ is also a universal cover of M then there exists an isomorphism of maps $f : \tilde{M}' \to \tilde{M}$ such that $\pi' = \pi \circ f$.

The following is proven in [21, Section 4.1].

Lemma 9.7.4. *Let (M,ρ) be a unimodular random rooted map. Let $(\tilde{M},\pi)$ be the universal cover of M and let $\tilde{\rho}$ be an arbitrary element of $\pi^{-1}(\rho)$. Then $(\tilde{M},\tilde{\rho})$ is a unimodular random rooted map.*

Observe that $\kappa_{\tilde{M}}(\tilde{\rho}) = \kappa(\rho)$ for any rooted map (M,ρ). Thus, we conclude by applying Theorem 9.6.1 and the classical theory of Riemann surfaces to obtain that

- $\mathbb{K}(M,\rho) > 0$ if and only if M is finite and simply connected and $\mathcal{R}(M)$ is conformally equivalent to the sphere,
- $\mathbb{K}(M,\rho) = 0$ if and only if $\tilde{\mathcal{R}}(M)$ is conformally equivalent to the plane, if and only if $\mathcal{R}(M)$ is conformally equivalent to one of the plane $\mathbb{C}$, the cylinder $\mathbb{C}/\mathbb{Z}$ or a torus $\mathbb{C}/\Lambda$ for some lattice $\Lambda \subset \mathbb{C}$, and
- $\mathbb{K}(M,\rho) < 0$ if and only if $\tilde{\mathcal{R}}(M)$ is conformally equivalent to the disc.

In the last case there are many possibilities for the conformal equivalence class $\mathcal{R}(M)$; anything other than the sphere, the plane, the cylinder or a torus will do. We conclude by applying the additional topological constraints on $\mathcal{R}(M)$ imposed by Theorem 9.7.1. □

We next connect the topology of $S(M)$ to the number of ends of the underlying graph. Since the number of ends of the underlying graph of a map M is at least the number of ends of $S(M)$, it follows that the underlying graph of M has infinitely many ends if $S(M)$ is homeomorphic to the Cantor tree or the blossoming Cantor tree, and at least two ends if $S(M)$ is homeomorphic to the cylinder or Jacob's ladder.

Lemma 9.7.5. *Let (M,ρ) be a unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$. If $S(M)$ is homeomorphic to the cylinder or to Jacob's ladder almost surely, then the underlying graph of M is two-ended almost surely.*

Proof. We bias by $\deg(\rho)$ and prove the equivalent statement for (M,ρ) reversible. We may also assume that (M,ρ) is ergodic. Suppose for contradiction that M has more than two ends. In this case, M has infinitely many ends almost surely, is invariantly nonamenable, and hence transient almost surely. Since $S(M)$ is two-ended almost surely, there exists some r and D such, with positive probability, the ball $B_r(M,X_n)$ of radius r about X_n in M has degree sum at most D and the complement $S(M)\setminus z(B_r(M,X_n))$ has two non-precompact connected components, each of which is necessarily one-ended. Denote this event by $\mathscr{A}_n$. By stationarity, $\mathscr{A}_n$ occurs for infinitely many n almost surely. Let $n_0 \geq 0$ be the minimal such n, and denote the components of the underlying graph of M that are contained in these two components of $S \setminus z(B_r(M,X_{n_0}))$ by W_1 and W_2. Since M is transient almost surely, the simple random walk $\langle X_n\rangle_{n\geq 0}$ eventually stays in one of the W_i, and so the subgraph of the underlying graph of M induced by this set W_i must be transient. However, taking a subsequence $\langle X_{n_m}\rangle_{m\geq 0}$ of the random walk such that $\mathscr{A}_{n_m}$ occurs for every $m \geq 0$ and the balls $B_r(M,X_{n_m})$ are all disjoint from each other yields an infinite collection of disjoint cutsets of degree sum at most D separating ρ from infinity in the subgraph induced by W_i. Thus, this graph is recurrent by the Nash-Williams criterion [173], a contradiction. □

9.7.2 Theorem 9.1.1 in the multiply-connected planar case

Suppose that (M,ρ) is an infinite, multiply-connected unimodular random planar map with $\mathbb{E}[\deg(\rho)] < \infty$. If $\mathbb{K}(M,\rho) = 0$, then the proof of Theorem 9.7.3 implies that $\mathcal{R}(M)$ is conformally equivalent to the cylinder. Lemma 9.7.5 then implies that the underlying graph of G is two-ended almost surely. It follows that M is recurrent, and we immediately deduce from this that items (2), (10), (11), (14), (17), (5), (6), and (4) of Theorem 9.1.1 hold for (M,ρ). The remaining items of Theorem 9.1.1 hold for (M,ρ) as a consequence of invariant amenability.

Now suppose that $\mathbb{K}(M,\rho) < 0$. In this case, Theorem 9.7.3 implies that $S(M)$ is almost surely homeomorphic to the Cantor tree and consequently that the underlying graph of M is infinitely-ended almost surely by Lemma 9.3.9. The following are well-known to experts.

Proposition 9.7.6. *Let (G,ρ) be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$, and suppose that G is infinitely ended almost surely. Then G admits non-constant harmonic functions of finite Dirichlet energy almost surely.*

Proposition 9.7.7. *Let (G,ρ) be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$, and suppose that G is infinitely ended almost surely. Then $p_u(G) = 1$ almost surely.*

Lemma 9.7.8. *Let (G,ρ) be a unimodular random rooted graph, and suppose that G is infinitely ended almost surely. Let $\langle X_n \rangle_{n\geq 0}$ be a random walk on G. Then for every finite set $K \subset V$ and every infinite connected component W of $G \setminus K$, there is a positive probability that $X_n \in W$ for all sufficiently large n.*

Proof. Suppose for contradiction that there exists a finite set $K \subset V$ and an infinite connected component W of $G \setminus K$ such that X_n does not visit W infinitely often almost surely. It follows that for every vertex v in W, a random walk started at v must hit the set K almost surely, and hence that, for each vertex u of K, we have

$$\inf_{v\in W} P_v(\text{hit } u) \geq \inf_{v\in W} P_v(\text{hit } K) \cdot \inf_{w\in K} P_w(\text{hit } u) > 0.$$

On the other hand, we must have that for every $\varepsilon > 0$, there exist at most finitely many vertices $v \in W$ such that $P_u(\text{hit } v) \geq \varepsilon$, since otherwise the random walk started at u would hit infinitely many vertices of W with probability at least ε by Fatou's Lemma. The identity

$$P_v(\text{hit u}) = P_u(\text{hit } v)\frac{\deg(u)\sum_{n\geq 0} p_n(u,u)}{\deg(v)\sum_{n\geq 0} p_n(v,v)}$$

therefore implies that for every $C < \infty$, there exist at most finitely many vertices v of W with $\deg(v)\sum_{n\geq 0} p_n(u,u) \leq C$.

Choose C sufficiently large that $\deg(\rho)\sum_{n\geq 0} p_n(\rho,\rho) \leq C$ with positive probability, and define a mass transport by, for each vertex v of G, transporting a mass of 1 to the closest vertex to v that has $\deg(w)\sum_{n\geq 0} p_n(w,w) \leq C$. If there are multiple choices of the vertex w, choose one uniformly.

Then every vertex sends a mass of at most one but, in the situation described, there must be some vertices that recieve an infinite amount of mass. This contradicts the Mass-Transport Principle. □

Proof of Proposition 9.7.6. It is well-known that if G is a graph and there exists a finite set K such that $G \setminus K$ has more than one transient connected component, then G admits a non-constant harmonic Dirichlet function: The probability that the walk eventually stays in a particular one of the connected components is such a function. See [173, Exercise 9.23]. □

An end ξ of a graph G is said to be **thin** if there exists a constant C and a sequence $\langle (K_i, W_i)\rangle_{i\geq 1}$ of sets $K_i \subset V$ and connected components W_i of $G \setminus K_i$ such that $|K_i| \leq M$ for all $i \geq 1$, $\operatorname{diam}(K_i) \leq M$ for all $i \geq 1$, and

$$\xi = \bigcap_{i\geq 1} \{\xi' \text{ an end of } G : \xi_{K_i} = W_i\}.$$

Proposition 9.7.9. *Let (G, ρ) be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$, and suppose that G is infinitely ended almost surely. Then the random walk on G converges almost surely in the ends compactification to a thin end of G, and the law of the limiting end has no atoms.*

Proof. We bias by $\deg(\rho)$ to work in the reversible setting, and let $\langle X_n\rangle_{n\in\mathbb{Z}}$ be a bi-infinite random walk started at ρ. Convergence in the space of ends follows immediately from transience. To see that the limiting end is almost surely thin, observe that, by Lemma 9.7.8, there exists $R < \infty$ such that, with positive probability, the random walks $\langle X_n\rangle_{n\geq m}$ and $\langle X_n\rangle_{n\leq m}$ are eventually in different connected components of the complement $G \setminus B_R(G, X_m)$ of the ball $B_R(G, X_m)$. By the ergodic theorem, this event must occur for infinitely many $m \in \mathbb{Z}$. The claim now follows immediately. □

Proof of Proposition 9.7.7. We bias by $\deg(\rho)$ to work in the reversible setting, and let $\langle X_n\rangle_{n\geq 0}$ be a random walk started at ρ. Let $p \in [0, 1]$. If ω_p has a unique infinite cluster almost surely, then Harris's inequality implies that

$$\begin{aligned} &\mathbb{P}(\rho \text{ is connected to } X_n \text{ in } \omega_p \mid (G, \rho, X_n)) \\ &\qquad \geq \mathbb{P}(\rho \text{ is in the unique infinite cluster of } \omega_p \mid (G, \rho)) \\ &\qquad\qquad \cdot \mathbb{P}(X_n \text{ is in the unique infinite cluster of } \omega_p \mid (G, X_n)), \end{aligned}$$

which does not converge to zero by stationarity.

However, Proposition 9.7.9 implies that there almost surely exists an increasing sequence i_n with $i_n \to \infty$ as $n \to \infty$, a constant $C < \infty$, and sequence of disjoint sets K_i with $|K_i| < C$ such that any path connecting ρ to X_m must pass through each of the sets K_i with $i \leq i_n$. If $p < 1$, then the probability that every edge incident to K_i is closed is at least $(1 - p)^C$, and it follows that

$$\mathbb{P}(\rho \text{ is connected to } X_n \text{ in } \omega_p \mid (G, \rho, X_n)) \leq (1 - (1 - p)^C)^{i_n},$$

which converges to zero almost surely if $p < 1$. The claim now follows immediately. □

The negations of the remaining items of Theorem 9.1.1 follow from Proposition 9.7.6 and Proposition 9.7.7 using implications, valid for all unimodular random rooted graphs, that we have already reviewed earlier in the paper; see the green and blue arrows in Figure 1.

9.8 Open problems

We expect that the dichotomy of Theorem 9.1.1 extends to many further properties of planar unimodular random rooted maps. In this section, we discuss several such properties that might be addressed.

9.8.1 Rates of escape of the random walk

Can the type of a unimodular random planar map be determined by the rate of escape of the random walk? The work of Ding, Lee, and Peres [79] (together with the characterization of parabolic unimodular random planar maps as Benjamini-Schramm limits of finite planar maps) implies that the random walk is at most diffusive on any parabolic unimodular random planar map of finite expected degree.

Theorem 9.8.1 ([79])**.** *There exists a universal constant C such that for every parabolic unimodular random rooted map (M, ρ) with $\mathbb{E}[\deg(\rho)] < \infty$, we have*

$$\mathbb{E}\left[\deg(\rho) d(\rho, X_n)^2\right] \leq Cn\, \mathbb{E}\left[\deg(\rho)\right]$$

for all $n \geq 0$.

On the other hand, if (M, ρ) is a hyperbolic unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$ that has at most exponential growth, meaning that

$$\limsup_{n\to\infty} \frac{1}{n} \log |B(\rho, n)| < \infty,$$

then the random walk on M has positive speed, that is,

$$\lim_{n\to\infty} \frac{1}{n} d(\rho, X_n) > 0$$

a.s., where the limit exists a.s. by Kingman's subadditive ergodic theorem. (Note that the exponential growth condition always holds for graphs of bounded degree.) This can be seen in several ways: it is an easy consequence of a theorem of Benjamini, Lyons, and Schramm [45, Theorem 3.2] (see also [7, Theorem 8.13] and [21, Theorem 3.2]) that every invariantly nonamenable unimodular random rooted graph with finite expected degree and at most exponential growth has positive speed.

Meanwhile, it is a result of Benjamini and Curien [38], generalizing the work of Kaimanovich, Vershik, and others [141, 144–148], that every non-Liouville unimodular random rooted graph with finite expected degree and at most exponential growth has positive speed.

In general, however, there do exist invariantly nonamenable, non-Liouville, unimodular random rooted graphs with finite expected degree such that the random walk has zero speed almost surely. An example of such a graph appears in a forthcoming paper by the second author. We do not know of a planar example, which motivates the following question.

Question 9.8.2. *Let (M, ρ) be a hyperbolic unimodular random rooted planar map with $\mathbb{E}[\deg(\rho)] < \infty$. Does the random walk on M have positive speed almost surely?*

See [21] for a related result concerning the positivity of the speed in the hyperbolic metric induced by the circle packing of a hyperbolic unimodular random rooted triangulation with $\mathbb{E}[\deg(\rho)^2] < \infty$.

9.8.2 Positive harmonic functions

Theorem 9.1.1 states that the existence of non-constant bounded harmonic functions and of non-constant harmonic Dirichlet functions are both determined by the type. We conjecture that a similar result holds for positive harmonic functions.

Conjecture 9.8.3. *Let (M, ρ) be a parabolic unimodular random rooted map and suppose that $\mathbb{E}[\deg(\rho)] < \infty$. Then M does not admit any non-constant positive harmonic functions almost surely.*

This conjecture would follow from a positive answer to the following question.

Question 9.8.4. *Let (M, ρ) be a parabolic unimodular random rooted map and suppose that $\mathbb{E}[\deg(\rho)] < \infty$. Let $\langle X_n \rangle_{n \geq 0}$ be a simple random walk on M. Is every component of the complement of the trace of $\langle X_n \rangle_{n \geq 0}$ finite almost surely?*

Note that the answer to Question 9.8.4 is trivially positive if M is recurrent.

Chapter 10

Uniform spanning forests of planar graphs

Summary. We prove that the free uniform spanning forest of any bounded degree, proper plane graph is connected almost surely, answering a question of Benjamini, Lyons, Peres and Schramm [44]. We provide a quantitative form of this result, calculating the critical exponents governing the geometry of the uniform spanning forests of transient proper plane graphs with bounded degrees and codegrees. We find that these exponents are universal in this class of graphs, provided that measurements are made using the hyperbolic geometry of their circle packings rather than their usual combinatorial geometry.

Lastly, we extend the connectivity result to bounded degree planar graphs that admit locally finite drawings in countably connected domains, but show that the result cannot be extended to general bounded degree planar graphs.

10.1 Introduction

The **uniform spanning forests** (USFs) of an infinite, locally finite, connected graph G are defined as weak limits of uniform spanning trees of finite subgraphs of G. These limits can be taken with either **free** or **wired** boundary conditions, yielding the **free uniform spanning forest (FUSF)** and the **wired uniform spanning forest (WUSF)** respectively. Although the USFs are defined as limits of random spanning trees, they need not be connected. Indeed, a principal result of Pemantle [190] is that the WUSF and FUSF of $\mathbb{Z}^d$ coincide, and that they are almost surely (a.s.) a single tree if and only if $d \leq 4$. Benjamini, Lyons, Peres and Schramm [44] (henceforth referred to as BLPS) later gave a complete characterisation of connectivity of the WUSF, proving that the WUSF of a graph is connected a.s. if and only if two independent random walks on the graph intersect a.s. [44, Theorem 9.2].

The FUSF is much less understood. No characterisation of its connectivity is known, nor has it been proven that connectivity is a zero-one event. One class of graphs in which the FUSF is relatively well understood are the *proper plane graphs*. Recall that a **planar graph** is a graph that can be embedded in the plane, while a **plane graph** is a planar graph G together with a specified embedding of G in the plane (or some other topological disc). A plane graph is **proper** if the embedding is proper, meaning that every compact subset of the plane (or whatever topological disc G was embedded in) intersects at most finitely many edges and vertices of the drawing (see

Section 10.2.4 for further details). For example, every tree can be drawn in the plane without accumulation points, while the product of $\mathbb{Z}$ with a finite cycle is planar but cannot be drawn in the plane without accumulation points (and therefore has no proper embedding in the plane).

BLPS proved that the free and wired uniform spanning forests are distinct whenever G is a transient proper plane graph with bounded degrees, and asked [44, Question 15.2] whether the FUSF is a.s. connected in this class of graphs. They proved that this is indeed the case when G is a self-dual plane Cayley graph that is rough-isometric to the hyperbolic plane [44, Theorem 12.7]. These hypotheses were later weakened by Lyons, Morris and Schramm [170, Theorem 7.5], who proved that the FUSF of any bounded degree proper plane graph that is rough-isometric to the hyperbolic plane is a.s. connected.

Our first result provides a complete answer to [44, Question 15.2], obtaining optimal hypotheses under which the FUSF of a proper plane graph is a.s. connected. The techniques we developed to answer this question also allow us to prove quantitative versions of this result, which we describe in the next section. We state our result in the natural generality of proper plane networks. Recall that a **network** (G, c) is a locally finite, connected graph $G = (V, E)$ together with a function $c : E \to (0, \infty)$ assigning a positive **conductance** to each edge of G. The **resistance** of an edge e in a network (G, c) is defined to be $1/c(e)$. Graphs may be considered as networks by setting $c \equiv 1$. A **plane network** is a planar graph G together with specified conductances and a specified drawing of G in the plane.

Theorem 10.1.1. *The free uniform spanning forest is almost surely connected in any bounded degree proper plane network with edge conductances bounded above.*

In light of the duality between the free and wired uniform spanning forests of proper plane graphs (see Section 10.2.4), the FUSF of a proper plane graph G with locally finite dual $G^\dagger$ is connected a.s. if and only if every component of the WUSF of $G^\dagger$ is a.s. *one-ended*. Thus, Theorem 10.1.1 follows easily from the dual statement Theorem 10.1.2 below. (The implication is immediate when the dual graph is locally finite.) Recall that an infinite graph $G = (V, E)$ is said to be **one-ended** if, for every finite set $K \subset V$, the subgraph induced by $V \setminus K$ has exactly one infinite connected component. In particular, an infinite tree is one-ended if and only if it does not contain a simple bi-infinite path. Components of the WUSF are known to be one-ended a.s. in several other classes of graphs [7, 44, 127, 128, 170, 190], and are recurrent in any graph [185]. Recall that a plane graph G is said to have **bounded codegree** if its faces have a bounded number of sides, or, equivalently, if its dual $G^\dagger$ has bounded degrees.

Theorem 10.1.2. *Every component of the wired uniform spanning forest is one-ended almost surely in any bounded codegree proper plane network with edge resistances bounded above.*

The uniform spanning trees of $\mathbb{Z}^2$ and other two dimensional Euclidean lattices are very well understood due to the deep theory of conformally invariant scaling limits. The study of the UST on $\mathbb{Z}^2$ led Schramm, in his seminal paper [208], to introduce the SLE processes, which he conjectured to describe the scaling limits of the loop-erased random walk and UST. This conjecture that was subsequently proven in the celebrated work of Lawler, Werner and Schramm [163]. Overall, Schramm's introduction of SLE has revolutionised the understanding of statistical physics in two dimensions; see e.g. [98, 162, 202] for guides to the extensive literature in this very active field.

Although our own setting is too general to apply this theory, we nevertheless keep conformal invariance in mind throughout this paper. Indeed, the key to our proofs is *circle packing*, a canonical method of drawing planar graphs that is closely related to conformal mapping (see e.g. [122, 123, 200, 202, 215] and references therein). For many purposes, one can pretend that the random walk on the packing is a quasiconformal image of standard planar Brownian motion: Effective resistances, heat kernels, and harmonic measures on the graph can each be estimated in terms of the corresponding Brownian quantities [17, 69].

10.1.1 Universal USF exponents via circle packing

A **circle packing** P is a set of discs in the Riemann sphere $\mathbb{C} \cup \{\infty\}$ that have disjoint interiors (i.e., do not overlap) but can be tangent. The **tangency graph** $G = G(P)$ of a circle packing P is the plane graph with the centres of the circles in P as its vertices and with edges given by straight lines between the centres of tangent circles. The Koebe-Andreev-Thurston Circle Packing Theorem [156, 179, 221] states that every *finite*, simple planar graph arises as the tangency graph of a circle packing, and that if the graph is a triangulation then its circle packing is unique up to Möbius transformations and reflections (see [58] for a combinatorial proof). The Circle Packing Theorem was extended to infinite plane triangulations by He and Schramm [121, 122], who proved that every *infinite, proper*, simple plane triangulation admits a locally finite circle packing in either the Euclidean plane or the hyperbolic plane (identified with the interior of the unit disc), but not both. We call an infinite, simple, proper plane triangulation **CP parabolic** if it admits a circle packing in the plane and **CP hyperbolic** otherwise.

He and Schramm [121] also initiated the use of circle packing to study probabilistic questions on plane graphs. In particular, they showed that a bounded degree, simple, proper plane triangulation is CP parabolic if and only if simple random walk on the triangulation is recurrent (i.e., visits every vertex infinitely often a.s.). Circle packing has since proven instrumental in the study of planar graphs, and random walks on planar graphs in particular. Most relevantly to us, circle packing was used by Benjamini and Schramm [47] to prove that every transient, bounded degree planar graph admits non-constant harmonic Dirichlet functions; BLPS [44] later applied this result to deduce that the free and wired uniform spanning forest of a bounded degree plane graph coincide if and only if the graph is recurrent. We refer the reader to [215] and [202] for background on circle packing, and to [17, 21, 22, 47, 50, 69, 107, 108, 121, 132, 139] for further probabilistic applications.

A guiding principle of the works mentioned above is that circle packing endows a triangulation with a geometry that, for many purposes, is better than the usual graph metric. The results described in this section provide a compelling instance of this principle in action: we find that the critical exponents governing the geometry of the USFs are universal over all transient, bounded degree, proper plane triangulations, provided that measurements are made using the hyperbolic

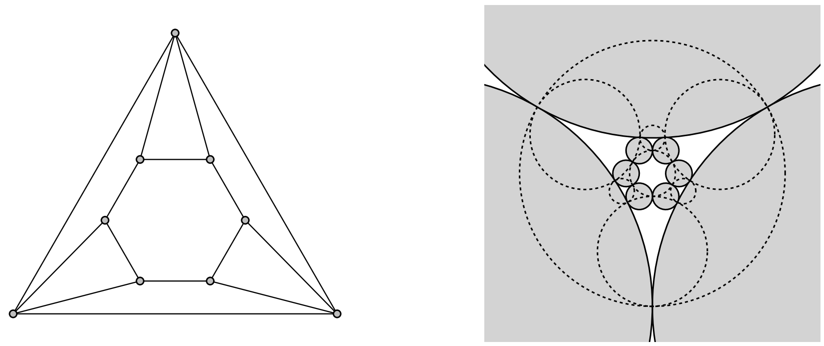

Figure 10.1: A simple, 3-connected, finite plane graph (left) and its double circle packing (right). Primal circles are filled and have solid boundaries, dual circles have dashed boundaries.

geometry of their circle packings rather than the usual combinatorial geometry of the graphs. It is crucial here that we use the circle packing to take our measurements: The exponents in the graph distance are not universal and need not even exist (see Figure 10.2). In Remark 10.6.1 we give an example to show that no such universal exponents hold in the CP parabolic case at this level of generality.

In order to state the quantitative versions of Theorems 10.1.1 and 10.1.2 in their full generality, we first introduce **double circle packing**. Let G be a plane graph with vertex set V and face set F. A double circle packing of G is a pair of circle packings $P = \{P(v) : v \in V\}$ and $P^\dagger = \{P^\dagger(f) : f \in F\}$ satisfying the following conditions (see Figure 10.1):

1. (G **is the tangency graph of** P.) For each pair of vertices u and v of G, the discs $P(u)$ and $P(v)$ are tangent if and only if u and v are adjacent in G.

2. ($G^\dagger$ **is the tangency graph of** $P^\dagger$.) For each pair of faces f and g of G, the discs $P^\dagger(f)$ and $P^\dagger(g)$ are tangent if and only if f and g are adjacent in $G^\dagger$.

3. (**Primal and dual circles are perpendicular**.) For each vertex v and face f of G, the discs $P^\dagger(f)$ and $P(v)$ have non-empty intersection if and only if f is incident to v, and in this case the boundary circles of $P^\dagger(f)$ and $P(v)$ intersect at right angles.

It is easily seen that any finite plane graph admitting a double circle packing must be simple (i.e., not containing any loops or multiple edges) and 3-connected (meaning that the subgraph induced by $V \setminus \{u, v\}$ is connected for each $u, v \in V$). Conversely, Thurston's interpretation of Andreev's Theorem [179, 221] implies that every finite, simple, 3-connected plane graph admits a double circle packing (see also [58]). The corresponding infinite theory was given developed by He [120], who proved that every infinite, simple, 3-connected, proper plane graph G with locally finite dual admits a double circle packing in either the Euclidean plane or the hyperbolic plane (but not both) and that this packing is unique up to Möbius transformations and reflections. As before, we say that G is CP parabolic or CP hyperbolic as appropriate. As in the He-Schramm Theorem [121], CP hyperbolicity is equivalent to transience for graphs with bounded degrees and codegrees [120].

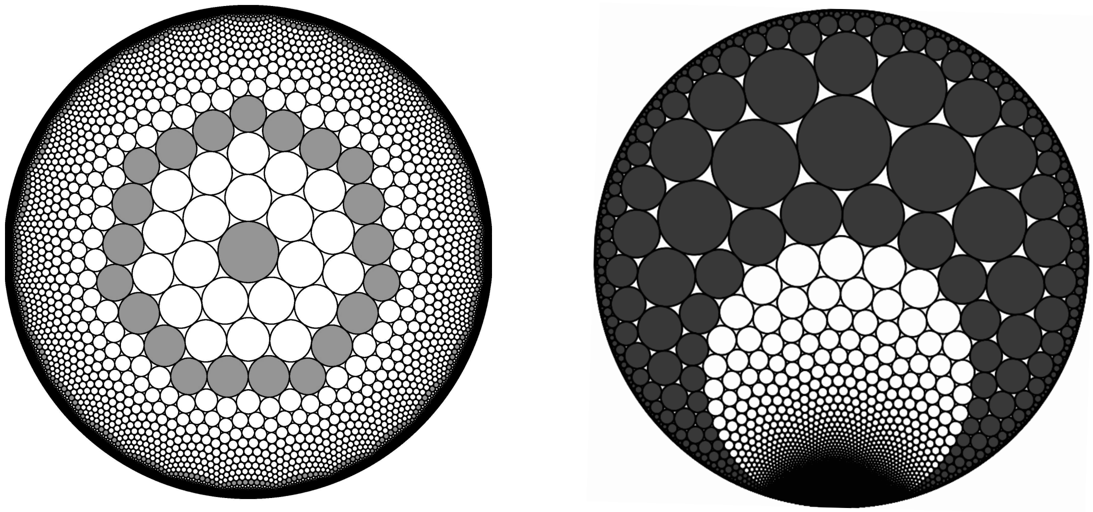

Figure 10.2: Two bounded degree, simple, proper plane triangulations for which the graph distance is not comparable to the hyperbolic distance. Similar examples are given in [215, Figure 17.7]. Left: In this example, rings of degree seven vertices (grey) are separated by growing bands of degree six vertices (white), causing the hyperbolic radii of circles to decay. The bands of degree six vertices can grow surprisingly quickly without the triangulation becoming recurrent [211]. Right: In this example, half-spaces of the 8-regular (grey) and 6-regular (white) triangulations have been glued together along their boundaries; the circles corresponding to the 6-regular half-space are contained inside a horodisc and have decaying hyperbolic radii.

Let G be CP hyperbolic and let $(P, P^\dagger)$ be a double circle packing of G in the hyperbolic plane. We write $r_{\mathbb{H}}(v)$ for the hyperbolic radius of the circle $P(v)$. For each subset $A \subset V(G)$, we define $\operatorname{diam}_{\mathbb{H}}(A)$ to be the hyperbolic diameter of the set of hyperbolic centres of the circles in P corresponding to vertices in A, and define $\operatorname{area}_{\mathbb{H}}(A)$ to be the hyperbolic area of the union of the circles in P corresponding to vertices in A. Since $(P, P^\dagger)$ is unique up to isometries of the hyperbolic plane, $r_{\mathbb{H}}(v)$, $\operatorname{diam}_{\mathbb{H}}(A)$, and $\operatorname{area}_{\mathbb{H}}(A)$ do not depend on the choice of $(P, P^\dagger)$.

We say that a network G has **bounded local geometry** if there exists a constant M such that $\deg(v) \leq M$ for every vertex v of G and $M^{-1} \leq c(e) \leq M$ for every edge e of G. Given a plane network G, we let F be the set of faces of G, and define

$$\mathbf{M} = \mathbf{M}_G = \max\Big\{ \sup_{v \in V} \deg(v),\ \sup_{f \in F} \deg(f),\ \sup_{e \in E} c(e),\ \sup_{e \in E} c(e)^{-1} \Big\}.$$

Given a spanning tree $\mathfrak{F}$ and two vertices x and y in G, we write $\Gamma_{\mathfrak{F}}(x, y)$ for the unique path in $\mathfrak{F}$ connecting x and y.

Theorem 10.1.3 (Free diameter exponent)**.** *Let G be a transient, simple, 3-connected, proper plane network with $\mathbf{M}_G < \infty$, let $\mathfrak{F}$ be the free uniform spanning forest of G, and let $e = (x, y)$ be an edge of G. Then there exist positive constants $k_1 = k_1(\mathbf{M}, r_{\mathbb{H}}(x))$, increasing in $r_{\mathbb{H}}(x)$, and $k_2 = k_2(\mathbf{M})$ such that*

$$k_1 R^{-1} \leq \mathbb{P}\left(\operatorname{diam}_{\mathbb{H}}(\Gamma_{\mathfrak{F}}(x, y)) \geq R\right) \leq k_2 R^{-1}$$

for every $R \geq 1$.

Given a spanning forest $\mathfrak{F}$ of G in which every component is a one-ended tree and an edge e of G, the **past** of e in $\mathfrak{F}$, denoted $\mathrm{past}_{\mathfrak{F}}(e)$, is defined to be the unique finite connected component of $\mathfrak{F} \setminus \{e\}$ if $e \in \mathfrak{F}$ and to be the empty set otherwise. The following theorem is equivalent to Theorem 10.1.3 by duality (see Sections 10.2.4 and 10.5.6).

Theorem 10.1.4 (Wired diameter exponent)**.** *Let G be a transient, simple, 3-connected, proper plane network with $\mathbf{M}_G < \infty$, let $\mathfrak{F}$ be the wired uniform spanning forest of G, and let $e = (x, y)$ be an edge of G. Then there exist positive constants $k_1 = k_1(\mathbf{M}, r_{\mathbb{H}}(x))$, increasing in $r_{\mathbb{H}}(x)$, and $k_2 = k_2(\mathbf{M})$ such that*

$$k_1 R^{-1} \leq \mathbb{P}\left(\mathrm{diam}_{\mathbb{H}}(\mathrm{past}_{\mathfrak{F}}(e)) \geq R\right) \leq k_2 R^{-1}$$

for every $R \geq 1$.

By similar methods, we are also able to obtain a universal exponent of $1/2$ for the tail of the area of $\mathrm{past}_{\mathfrak{F}}(e)$, where $\mathfrak{F}$ is the WUSF of G.

Theorem 10.1.5 (Wired area exponent)**.** *Let G be a transient, simple, 3-connected, proper plane network with $\mathbf{M}_G < \infty$, let $\mathfrak{F}$ be the wired uniform spanning forest of G, and let $e = (x, y)$ be an edge of G. Then there exist positive constants $k_1 = k_1(\mathbf{M}, r_{\mathbb{H}}(x))$, increasing in $r_{\mathbb{H}}(x)$, and $k_2 = k_2(\mathbf{M})$ such that*

$$k_1 R^{-1/2} \leq \mathbb{P}\left(\mathrm{area}_{\mathbb{H}}(\mathrm{past}_{\mathfrak{F}}(e)) \geq R\right) \leq k_2 R^{-1/2}$$

for every $R \geq 1$.

The exponents 1 and 1/2 occurring in Theorems 10.1.4 and 10.1.5 should, respectively, be compared with the analogous exponents for the survival time and total progeny of a critical branching process whose offspring distribution has finite variance (see e.g. [173, § 5,12]).

For *uniformly transient* proper plane graphs (that is, proper plane graphs in which the escape probabilities of the random walk are bounded uniformly away from zero), the hyperbolic radii of circles in $(P, P^\dagger)$ are bounded away from zero uniformly (Proposition 10.5.16). This implies that the hyperbolic metric and the graph metric are rough-isometric (Corollary 10.5.17). Consequently, Theorems 10.1.3–10.1.5 hold with the graph distance and counting measure appropriately (Corollary 10.5.18). This yields the following appealing corollary, which applies, for example, to planar Cayley graphs of co-compact Fuchsian groups.

Corollary 10.1.6 (Free length exponent)**.** *Let G be a uniformly transient, simple, 3-connected, proper plane network with $\mathbf{M}_G < \infty$ and let $\mathfrak{F}$ be the free uniform spanning forest of G. Let $\mathbf{p} > 0$ be a uniform lower bound on the escape probabilities of G. Then there exist positive constants $k_1 = k_1(\mathbf{M}, \mathbf{p})$ and $k_2 = k_2(\mathbf{M}, \mathbf{p})$ such that*

$$k_1 R^{-1/2} \leq \mathbb{P}\left(|\Gamma_{\mathfrak{F}}(x,y)| \geq R\right) \leq k_2 R^{-1/2}$$

for every edge $e = (x,y)$ of G and every $R \geq 1$.

See [31, 52, 150, 170, 180, 210] and the survey [29] for related results on Euclidean lattices.

Organisation

Section 10.2 contains definitions and background on those notions that will be used throughout the paper. Experienced readers are advised that this section also includes proofs of a few simple folklore-type lemmas and propositions. Section 10.3 contains the proofs of Theorems 10.1.1 and 10.1.2, as well as the upper bound of Theorem 10.1.4 in the case that G is a triangulation. Section 10.5 completes the proofs of Theorems 10.1.3–10.1.5 and Corollary 10.1.6; the most substantial component of this section is the proof of the lower bound of Theorem 10.1.4. We conclude with some remarks and open problems in Section 10.6.

10.2 Background and definitions

10.2.1 Notation

We write $E^{\rightarrow}$ for the set of oriented edges of a network $G = (V, E)$. An oriented edge $e \in E^{\rightarrow}$ is oriented from its tail e^- to its head e^+, and has reversal $-e$. For each $r', r > 0$, and $z \in \mathbb{C}$, we define the ball

$$B_z(r) = \{z' \in \mathbb{C} : |z - z'| < r\}$$

and the annulus

$$A_z(r, r') = \{z' \in \mathbb{C} : r \leq |z' - z| \leq r'\}.$$

We write $d_{\mathbb{C}}$ for the Euclidean metric on $\mathbb{C}$ and write $d_{\mathbb{H}}$ for the hyperbolic metric on $\mathbb{D}$.

10.2.2 Uniform Spanning Forests

We begin with a succinct review of some basic facts about USFs, referring the reader to [44] and Chapters 4 and 10 of [173] for a comprehensive overview. For each finite, connected graph G, we define UST_G to be the uniform measure on the set of spanning trees of G (i.e., connected subgraphs of G containing every vertex and no cycles). More generally, for a finite network $G = (G, w)$, we define UST_G to be the probability measure on spanning trees of G for which the measure of each tree is proportional to the product of the conductances of the edges in the tree.

An **exhaustion** of an infinite network G is a sequence $\langle V_n \rangle_{n \geq 0}$ of finite, connected subsets of V such that $V_n \subseteq V_{n+1}$ for all $n \geq 0$ and $\cup_n V_n = V$. Given such an exhaustion, let the network G_n be the subgraph of G induced by V_n together with the conductances inherited from G. The **free uniform spanning forest** measure FUSF_G is defined to be the weak limit of the sequence

$\langle \mathsf{UST}_{G_n} \rangle_{n\geq 1}$, so that

$$\mathsf{FUSF}_G(S \subset \mathfrak{F}) = \lim_{n\to\infty} \mathsf{UST}_{G_n}(S \subset T).$$

for each finite set $S \subset E$, where $\mathfrak{F}$ is a sample of FUSF_G and T is a sample of UST_{G_n}. For each n, we also construct a network G_n^* from G by gluing (wiring) every vertex of $G \setminus G_n$ into a single vertex, denoted ∂_n, and deleting all the self-loops that are created. We identify the set of edges of G_n^* with the set of edges of G that have at least one endpoint in V_n. The **wired uniform spanning forest** measure WUSF_G is defined to be the weak limit of the sequence $\langle \mathsf{UST}_{G_n^*} \rangle_{n\geq 1}$, so that

$$\mathsf{WUSF}_G(S \subset \mathfrak{F}) = \lim_{n\to\infty} \mathsf{UST}_{G_n^*}(S \subset T)$$

for each finite set $S \subset E$, where $\mathfrak{F}$ is a sample of WUSF_G and T is a sample of $\mathsf{UST}_{G_n^*}$. These weak limits were both implicitly proven to exist by Pemantle [190], although the WUSF was not considered explicitly until the work of Häggström [109]. Both measures are easily seen to be concentrated on the set of **essential spanning forests** of G, that is, cycle-free subgraphs of G including every vertex such that every connected component is infinite. The measure FUSF_G stochastically dominates WUSF_G for every infinite network G.

The Spatial Markov Property

Let $G = (V, E)$ be a finite or infinite network and let A and B be subsets of E. We write $(G-B)/A$ for the (possibly disconnected) network formed from G by deleting every edge in B and contracting (i.e., identifying the two endpoints of) every edge in A. We identify the edges of $(G-B)/A$ with $E \setminus B$. Suppose that G is finite, and that

$$\mathsf{UST}_G(A \subseteq T, B \cap T = \emptyset) > 0.$$

Then, given the event that T contains every edge in A and none of the edges in B, the conditional distribution of T is equal to the union of A and the UST of $(G-B)/A$. That is, for every event $\mathscr{A} \subseteq \{0,1\}^E$,

$$\mathsf{UST}_G(T \in \mathscr{A} \mid A \subset T,\ B \cap T = \emptyset) = \mathsf{UST}_{(G-B)/A}(T \cup A \in \mathscr{A}).$$

This is the **spatial Markov property** of the UST. Taking limits over exhaustions, we obtain a corresponding spatial Markov property for the USFs: If $G = (V, E)$ is an infinite network and A and B are subsets of E such that

$$\mathsf{WUSF}_G(A \subseteq \mathfrak{F}, B \cap \mathfrak{F} = \emptyset) > 0$$

and the network $(G-B)/A$ is locally finite, then

$$\mathsf{WUSF}_G(\mathfrak{F} \in \mathscr{A} \mid A \subset \mathfrak{F},\ B \cap \mathfrak{F} = \emptyset) = \mathsf{WUSF}_{(G-B)/A}(\mathfrak{F} \cup A \in \mathscr{A}). \tag{10.2.1}$$

(If $(G-B)/A$ is not connected, we define $\mathsf{WUSF}_{(G-B)/A}$ to be the product of the WUSF measures on the connected components of $(G-B)/A$.) A similar spatial Markov property holds for the FUSF.

The UST and USFs also enjoy a *strong* form of the spatial Markov property. Let G be a finite or infinite network, and suppose that H is a random element of $\{0,1\}^E$, which we think of as a random subgraph of G. We say that a random element K of $\{0,1\}^E$, defined on the same probability space as H, is a **local set for** H if for every set $W \subseteq E$, the event $\{K \subseteq W\}$ is, up to a null set, measurable with respect to the σ-algebra generated by the collection of random variables $\{H(e) : e \in W\}$.

For example, suppose that $\mathfrak{F}$ is a random spanning forest of a network G, and that u and v are vertices of G. Let $K = K(\mathfrak{F})$ be the random set of edges that is equal to the path connecting u and v in $\mathfrak{F}$ if such a path exists, and otherwise equal to all of E. Then K is a local set for $\mathfrak{F}$, since, for every proper subset W of E, the event $\{K \subseteq W\}$ is equal (up to the null event in which $\mathfrak{F}$ is not a forest) to the event that u and v are in the same component of $\mathfrak{F}$ and that the unique path connecting u and v in $\mathfrak{F}$ is contained in W, which is clearly measurable with respect to the restriction of $\mathfrak{F}$ to W.

An example of a random set K that is not a local set for $\mathfrak{F}$ is, for instance, the edges of $\mathfrak{F}$ touching a given vertex v. If v has degree at least 2 and e is some edge touching v, then the event $\{K \subset \{e\}\}$ cannot be determined just by knowing $\mathfrak{F}(e)$. Note also that Proposition 10.2.1 does not hold for this K — because knowing K implies that $\mathfrak{F}(e) = 0$ for all edges e that touch v but $e \notin K$.

Given a random subgraph H of a network G and a local set K for H, we write $\mathcal{F}_K$ to denote the σ-algebra generated by K and the random collection of random variables $\{H(e) : e \in K\}$. We also write $K_o = \{e \in E : e \in K, H(e) = 1\}$ and $K_c = K \setminus K_o$ for the sets of edges in K that are included (open) in H and not included (closed) in H respectively.

Proposition 10.2.1 (Strong Spatial Markov Property for the WUSF)**.** *Let $G = (V, E)$ be an infinite network, let $\mathfrak{F}$ be a sample of the wired uniform spanning forest of G, and let K be a local set for $\mathfrak{F}$ that is either finite or equal to all of E almost surely. Conditional on $\mathcal{F}_K$ and the event that K is finite, let $\hat{\mathfrak{F}}$ be a sample of the wired uniform spanning forest of $(G - K_c)/K_o$. Let $\mathfrak{F}' = K_o \cup \hat{\mathfrak{F}}$ if K_o is finite and $\mathfrak{F}' = \mathfrak{F}$ if $K_o = E$. Then $\mathfrak{F}$ and $\mathfrak{F}'$ have the same distribution.*

Proof. Expand the conditional probability

$$\begin{aligned}&\mathsf{WUSF}_G\left(\mathfrak{F} \in \mathscr{A} \mid \mathcal{F}_K\right) \mathbb{1}(K \neq E)\\ &\qquad\qquad = \sum_{W_1, W_2 \subset E \text{ finite}} \mathsf{WUSF}_G\left(\mathfrak{F} \in \mathscr{A} \mid K_o = W_1, K_c = W_2\right) \mathbb{1}\left(K_o = W_1, K_c = W_2\right).\end{aligned}$$

Since K is a local set for $\mathfrak{F}$, the right-hand side is equal to

$$\sum_{W_1, W_2 \subset E \text{ finite}} \mathsf{WUSF}_G\left(\mathfrak{F} \in \mathscr{A} \mid W_1 \subseteq \mathfrak{F}, W_2 \cap \mathfrak{F} = \emptyset\right) \mathbb{1}\left(K_o = W_1, K_c = W_2\right).$$

Applying the spatial Markov property gives

$$\begin{aligned}\mathsf{WUSF}_G\left(\mathfrak{F}\in\mathscr{A}\mid\mathcal{F}_K\right)\mathbb{1}(K\neq E)&\\ &=\sum_{W_1,W_2\subset E\text{ finite}}\mathsf{WUSF}_{(G-W_2)/W_1}\left(\hat{\mathfrak{F}}\cup W_1\in\mathscr{A}\right)\mathbb{1}\left(K_o=W_1,K_c=W_2\right)\\ &=\mathsf{WUSF}_{(G-K_c)/K_o}\left(\hat{\mathfrak{F}}\cup K_o\in\mathscr{A}\right)\mathbb{1}(K\neq E),\end{aligned}$$

from which the claim follows immediately. □

Similar strong spatial Markov properties hold for the FUSF and UST, and admit very similar proofs.

10.2.3 Random walk, effective resistances

Given a network G and a vertex u of G, we write $\mathbf{P}_u^G$ for the law of the simple random walk on G started at u, and will often write simply $\mathbf{P}_u$ if the choice of network is unambiguous. For each set of vertices A, we let τ_A be the first time the random walk visits A, letting $\tau_A=\infty$ if the walk never visits A. Similarly, τ_A^+ is defined to be the first positive time that the random walk visits A. The conductance $c(u)$ of a vertex u is defined to be the sum of the conductances of the edges emanating from u.

Let A and B be sets of vertices in a finite network G. The **effective conductance** between A and B in G is defined to be

$$\mathscr{C}_{\mathrm{eff}}(A\leftrightarrow B;G)=\sum_{v\in A}c(v)\mathbf{P}_v(\tau_B<\tau_A^+),$$

while the effective resistance $\mathscr{R}_{\mathrm{eff}}(A\leftrightarrow B;G)$ is defined to be the reciprocal of the effective conductance, $\mathscr{R}_{\mathrm{eff}}(A\leftrightarrow B;G)=\mathscr{C}_{\mathrm{eff}}(A\leftrightarrow B;G)^{-1}$. Now suppose that G is an infinite network with exhaustion $\langle V_n\rangle_{n\geq 0}$ and let A and B be finite subsets of V. Let $\langle G_n\rangle_{n\geq0}$ and $\langle G_n^*\rangle_{n\geq0}$ be defined as in Section 10.2.2. The **free effective resistance between** A **and** B is defined to be the limit

$$\mathscr{R}^{\mathrm{F}}_{\mathrm{eff}}(A\leftrightarrow B;\,G)=\lim_{n\to\infty}\mathscr{R}_{\mathrm{eff}}(A\leftrightarrow B;\,G_n),$$

while the **wired effective resistance between** A **and** B is defined to be

$$\mathscr{R}^{\mathrm{W}}_{\mathrm{eff}}(A\leftrightarrow B;\,G)=\lim_{n\to\infty}\mathscr{R}_{\mathrm{eff}}(A\leftrightarrow B;\,G_n^*).$$

Free and wired effective conductances are defined by taking reciprocals. The free effective resistance between two, possibly infinite, sets A and B is defined to be the limit of the free effective resistances between $A\cap V_n$ and $B\cap V_n$, which are decreasing in n. We also define

$$\mathscr{R}_{\mathrm{eff}}\left(A\leftrightarrow B\cup\{\infty\}\right)=\lim_{n\to\infty}\mathscr{R}_{\mathrm{eff}}\left(A\leftrightarrow B\cup\{\partial_n\};G_n^*\right).$$

See e.g. [173] for further background on electrical networks.

We will make frequent use of **Rayleigh' monotonicity principle** [173, Chapter 2.4], which states that the effective conductance between any two sets in a network is an increasing function of the edge conductances. In particular, it follows that the effective conductance between two sets decreases when edges are deleted from the network (which corresponds to taking the conductance of those edges to zero), and increases when edges are contracted (which corresponds to taking the conductance of those edges to infinity).

The proof of Theorem 10.1.2 will require the following simple lemma.

Lemma 10.2.2. *Let A and B be sets of vertices in an infinite network G. Then*

$$\mathscr{R}_{\mathrm{eff}}^{\mathrm{W}}(A \leftrightarrow B;\, G) \leq 3 \max\Big\{\mathscr{R}_{\mathrm{eff}}\left(A \leftrightarrow B \cup \{\infty\};\, G\right),\, \mathscr{R}_{\mathrm{eff}}\left(B \leftrightarrow A \cup \{\infty\};\, G\right)\Big\}.$$

Proof. Let $M = \max\big\{\mathscr{R}_{\mathrm{eff}}(A \leftrightarrow B \cup \{\infty\};\, G),\, \mathscr{R}_{\mathrm{eff}}(B \leftrightarrow A \cup \{\infty\};\, G)\big\}$. Recall that for any three sets A, B and C in G_n^* (or any other finite network) [173, Exercise 2.33],

$$\mathscr{R}_{\mathrm{eff}}(A \leftrightarrow B \cup C; G_n^*)^{-1} \leq \mathscr{R}_{\mathrm{eff}}(A \leftrightarrow B; G_n^*)^{-1} + \mathscr{R}_{\mathrm{eff}}(A \leftrightarrow C; G_n^*)^{-1}.$$

Letting $C = \{\partial_n\}$ and taking the limit as $n \to \infty$, we obtain that

$$M^{-1} \leq \mathscr{R}_{\mathrm{eff}}(A \leftrightarrow B \cup \{\infty\}; G)^{-1} \leq \mathscr{R}_{\mathrm{eff}}^{\mathrm{W}}(A \leftrightarrow B; G)^{-1} + \mathscr{R}_{\mathrm{eff}}(A \leftrightarrow \infty; G)^{-1}.$$

Rearranging, we have that

$$\mathscr{R}_{\mathrm{eff}}(A \leftrightarrow \infty; G) \leq \frac{M\mathscr{R}_{\mathrm{eff}}^{\mathrm{W}}(A \leftrightarrow B; G)}{\mathscr{R}_{\mathrm{eff}}^{\mathrm{W}}(A \leftrightarrow B; G) - M}.$$

By symmetry, the inequality continues to hold when we exchange the roles of A and B. Combining both of these inequalities with the triangle inequality for effective resistances [173, Exercise 9.29] yields that

$$\mathscr{R}_{\mathrm{eff}}^{\mathrm{W}}(A \leftrightarrow B; G) \leq \mathscr{R}_{\mathrm{eff}}(A \leftrightarrow \infty; G) + \mathscr{R}_{\mathrm{eff}}(B \leftrightarrow \infty; G) \leq 2\frac{M\mathscr{R}_{\mathrm{eff}}^{\mathrm{W}}(A \leftrightarrow B; G)}{\mathscr{R}_{\mathrm{eff}}^{\mathrm{W}}(A \leftrightarrow B; G) - M},$$

which rearranges to give the claimed inequality. □

Kirchhoff's Effective Resistance Formula.

The connection between effective resistances and spanning trees was first discovered by Kirchhoff [154] (see also [63]).

Theorem 10.2.3 (Kirchhoff's Effective Resistance Formula)**.** *Let G be a finite network. Then for every $e \in E^{\rightarrow}$*

$$\mathsf{UST}_G(e \in T) = c(e)\mathscr{R}_{\mathrm{eff}}(e^- \leftrightarrow e^+; G).$$

Taking limits over exhaustions, we also have the following extension of Kirchhoff's formula:

$$\mathsf{WUSF}_G(e \in \mathfrak{F}) = c(e)\mathscr{R}^{\mathrm{W}}_{\mathrm{eff}}(e^- \leftrightarrow e^+; G). \tag{10.2.2}$$

A similar equality holds for the FUSF.

The method of random paths.

We say that a path Γ in a network G is **simple** if it does not visit any vertex more than once. Given a probability measure ν on simple paths in a network G, we define the **energy** of ν to be

$$\mathcal{E}(\nu) = \frac{1}{2}\sum_{e \in E^{\rightarrow}} \frac{1}{c(e)}\big(\nu(e \in \Gamma) - \nu(-e \in \Gamma)\big)^2.$$

Effective resistances can be estimated using random paths as follows (see [173, §3] and [**?**]). In particular, if G is an infinite network and A and B are two finite sets of vertices in G, then

$$\mathscr{R}_{\mathrm{eff}}(A \leftrightarrow B \cup \{\infty\}; G) = \min\left\{ \mathcal{E}(\nu) : \begin{array}{l} \nu \text{ a probability measure on simple} \\ \text{paths in } G \text{ starting in } A \text{ that are} \\ \text{either infinite or finite and end in } B \end{array} \right\}.$$

Obtaining resistance bounds by defining flows using random paths in this manner is referred to as the **method of random paths**. It will be convenient to use the following weakening of the method of random paths. Given the law μ of a random subset $W \subset V(G)$, define

$$\mathcal{E}(\mu) = \sum_{v \in V} \mu(v \in W)^2.$$

Lemma 10.2.4 (Method of random sets)**.** *Let A and B be two finite sets of vertices in an infinite network G, and let μ be a measure on subsets $W \subset V(G)$ such that the subgraph of G induced by V almost surely contains a path starting at A that is either infinite or finite and ends at B. Then*

$$\mathscr{R}_{\mathrm{eff}}(A \leftrightarrow B \cup \{\infty\}; G) \leq \sup_{e \in E} c(e)^{-1}\mathcal{E}(\mu). \tag{10.2.3}$$

Proof. Given W, let Γ be a simple path connecting A to B that is contained in W. Then, letting ν be the law of Γ,

$$\mathcal{E}(\nu) \leq \sup_e \frac{1}{c(e)} \sum_{e \in E^{\rightarrow}} \nu(e \in \Gamma)^2.$$

Letting Γ' be an independent random path with the same law as Γ, the sum above is exactly the expected number of oriented edges that are used by both Γ and Γ'. Since these paths are simple, they each contain at most one oriented edge emanating from v for each vertex $v \in V$. It follows that the number of oriented edges included in both paths is at most the number of vertices included in both paths. This yields that

$$\mathcal{E}(\nu) \leq \sup_{e\in E} \frac{1}{c(e)} \sum_{v\in V} \nu(v \in \Gamma)^2 \leq \sup_{e\in E} \frac{1}{c(e)} \sum_{v\in V} \mu(v \in W)^2 = \sup_{e\in E} \frac{1}{c(e)} \mathcal{E}(\mu). \qquad \square$$

10.2.4 Plane Graphs and their USFs

Given a graph $G = (V, E)$, let $\mathbb{G}$ be the metric space defined as follows. For each edge e of G, choose an orientation of e arbitrarily and let $\{I(e) : e \in E\}$ be a set of disjoint isometric copies of the interval $[0, 1]$. The metric space $\mathbb{G}$ is defined as a quotient of the union $\bigcup_e I(e) \cup V$, where we identify the endpoints of $I(e)$ with the vertices e^- and e^+ respectively, and is equipped with the path metric.

Let S be an orientable surface without boundary, which in this paper will always be a domain $D \subseteq \mathbb{C}\cup\{\infty\}$. A **proper embedding** of a graph G into S is a continuous, injective map $z : \mathbb{G} \to S$ satisfying the following conditions:

1. (**Every face is a topological disc.**) Every connected component of the complement $S \setminus z(\mathbb{G})$, called a **face** of (G, z), is homeomorphic to the disc. Moreover, for each connected component U of $S \setminus z(\mathbb{G})$, the set of oriented edges of G that have their left-hand side incident to U forms either a cycle or a bi-infinite path in G.

2. (z **is locally finite.**) Every compact subset of S intersects at most finitely many edges of $z(\mathbb{G})$. Equivalently, the preimage $z^{-1}(K)$ of every compact set $K \subseteq S$ is compact in $\mathbb{G}$.

A locally finite, connected graph is planar if and only if it admits a proper embedding into some domain $D \subseteq \mathbb{C} \cup \{\infty\}$. A **plane graph** $G = (G, z)$ is a planar graph G together with a specified embedding $z : \mathbb{G} \to D \subseteq \mathbb{C} \cup \{\infty\}$. A **plane network** $G = (G, z, c)$ is a planar graph together with a specified embedding and an assignment of positive conductances $c : E \to (0, \infty)$.

Given a pair $G = (G, z)$ of a graph together with a proper embedding z of G into a domain D, the **dual** $G^\dagger$ of G is the graph that has the faces of G as vertices, and has an edge drawn between two faces of G for each edge incident to both of the faces in G. By drawing each vertex of $G^\dagger$ in the interior of the corresponding face of G and each edge of $G^\dagger$ so that it crosses the corresponding edge of G but no others, we obtain an embedding $z^\dagger$ of $G^\dagger$ in D. The edge sets of G and $G^\dagger$ are in natural correspondence, and we write $e^\dagger$ for the edge of $G^\dagger$ corresponding to e. If e is oriented, we let $e^\dagger$ be oriented so that it crosses e from right to left as viewed from the orientation of e. If $G = (G, z, c)$ is a plane network, we assign the conductances $c^\dagger(e^\dagger) = c(e)^{-1}$ to the edges of $G^\dagger$.

USF Duality.

Let G be a plane network with dual $G^\dagger$. For each set $W \subseteq E$, let $W^\dagger := \{e^\dagger : e \notin W\}$. If G is finite and t is a spanning tree of G, then $t^\dagger$ is a spanning tree of $G^\dagger$: the subgraph $t^\dagger$ is connected because t has no cycles, and has no cycles because t is connected. Moreover, the ratio

$$\frac{\prod_{e\in t} c(e)}{\prod_{e^\dagger\in t^\dagger} c^\dagger(e^\dagger)} = \prod_{e\in E} c(e)$$

does not depend on t. It follows that if T is a random spanning tree of G with law UST_G, then $T^\dagger$ is a random spanning tree of $G^\dagger$ with law $\mathsf{UST}_{G^\dagger}$. This duality was extended to infinite proper plane networks by BLPS.

Theorem 10.2.5 ([44, Theorem 12.2 and Proposition 12.5])**.** *Suppose that G is an infinite proper plane network with locally finite dual $G^\dagger$. Then if $\mathfrak{F}$ is a sample of FUSF_G, the subgraph $\mathfrak{F}^\dagger$ is an essential spanning forest of $G^\dagger$ with law $\mathsf{WUSF}_{G^\dagger}$. In particular, $\mathfrak{F}$ is connected almost surely if and only if every component of $\mathfrak{F}^\dagger$ is one-ended almost surely.*

10.2.5 Circle packing

We now give some background on circle packing. The **carrier** of a circle packing P, denoted $\mathrm{carr}(P)$, is the union of all the discs of P and of all the faces of $G(P)$, so that the embedding z of $G(P)$ defined by drawing straight lines between the centres of tangent circles is a proper embedding of $G(P)$ into $\mathrm{carr}(P)$. Similarly, the carrier of a double circle packing $(P, P^\dagger)$ is defined to be the union of all the discs in $P \cup P^\dagger$. Given a domain $D \subset \mathbb{C} \cup \{\infty\}$, a circle packing P (or double circle packing $(P, P^\dagger)$) is said to be **in** D if its carrier is D. In particular, a (double) circle packing P is said to be **in the plane** if its carrier is the plane $\mathbb{C}$ and **in the disc** if its carrier is the open unit disc $\mathbb{D}$. The following theorems, which we stated in the introduction, are the cornerstones of the theory for infinite proper plane graphs.

Theorem 10.2.6. *[He-Schramm Existence and Uniqueness Theorem [120–122, 206]] Let G be an infinite, simple, 3-connected, proper plane graph with locally finite dual. Then G admits a double circle packing either in the plane or the disc, but not both, and this packing is unique up to Möbius transformations and reflections of the plane or the disc as appropriate.*

Theorem 10.2.7. *[He-Schramm Recurrence Theorem [120, 121]] Let G be an infinite, simple, 3-connected, proper plane graph with bounded degrees and codegrees. Then G is CP parabolic if and only if it is recurrent for simple random walk.*

Let us now explain how these theorems follow from the more general statements given in [120]. In that paper, He considers **disc patterns** of simple proper plane triangulations, which are like circle packings except that circles corresponding to adjacent vertices intersect at a specified angle, between 0 and π, rather than being tangent. Given a proper plane network $G = (V, E)$ with with locally finite dual and face set F, we can form a proper plane triangulation T with vertex set $V \cup F$ by adding a vertex inside each face of G and connecting this vertex to each vertex in the boundary of the face. It is easily verified that T is simple if and only if G is simple and 3-connected. Moreover, a double circle packing of T, either in the plane or the disc, can now be obtained using

Theorem 1.3 of [120] by requiring that the angles between circles corresponding to the original edges of G are 0 and $\pi/2$ between circles corresponding to edges of T that are in $V \times F$. It is straightforward to check that conditions (C1) and (C2) of Theorem 1.3 of [120] hold, and that the packing obtained this way is indeed the double circle packing of G. Thus, the existence statement of Theorem 10.2.6 follows. The uniqueness of this packing, as formulated in Theorem 10.2.6, is the content of Theorems 1.1 and 1.2 of [120]. Lastly, Theorem 10.2.7 is a direct consequence of Theorem 1.3 of [120] together with Theorems 2.6 and 8.1 of [121].

Besides their results that we have already mentioned, He and Schramm [122] also proved that every simple triangulation of a countably-connected domain D can be circle packed in a **circle domain**, that is, a domain D' such that every connected component of $\mathbb{C} \cup \{\infty\} \setminus D'$ is a disc or a point.

Given a double circle packing $(P, P^\dagger)$ of a plane graph G in a domain $D \subseteq \mathbb{C}$, we write $\operatorname{diam}_{\mathbb{C}}(A)$ for the Euclidean diameter of the set $\{z(v) : v \in A\}$, and write $d_{\mathbb{C}}(A, B)$ for the Euclidean distance between the sets $\{z(v) : v \in A\}$ and $\{z(v) : v \in B\}$. We write $r(v)$ and $r(f)$ for the Euclidean radii of the circles $P(v)$ and $P^\dagger(f)$. If $(P, P^\dagger)$ has carrier $\mathbb{D}$, we write $\sigma(v)$ for $1 - |z(v)|$, which is the distance between $z(v)$ and the boundary of $\mathbb{D}$, and write $r_{\mathbb{H}}(v)$ for the hyperbolic radius of $P(v)$. (Recall that Euclidean circles in $\mathbb{D}$ are also hyperbolic circles with different centres and radii; see e.g. [13] for further background on hyperbolic geometry.)

We will also use the Ring Lemma of Rodin and Sullivan [200]; see [116] and [2] for quantitative versions. In Section 10.5.1 we formulate and prove a version of the Ring Lemma for double circle packings, Theorem 10.5.1.

Theorem 10.2.8 (The Ring Lemma)**.** *There exists a sequence of positive constants $\langle k_m : m \geq 3\rangle$ such that for every circle packing P of every simple triangulation T, and every pair of adjacent vertices u and v of T, the ratio of radii $r(v)/r(u)$ is at most $k_{\deg(v)}$.*

An immediate corollary of the Ring Lemma is that, whenever P is a circle packing in $\mathbb{D}$ of a CP hyperbolic proper plane triangulation T and v is a vertex of T, the hyperbolic radius of P_v is bounded above by a constant $C = C(\deg(v))$.

10.3 Connectivity of the FUSF

In this section we prove Theorem 10.1.2 and show it easily implies Theorem 10.1.1. We will write $\preceq$, $\succeq$ and $\asymp$ for inequalities that hold up to a positive multiplicative constant depending only on $\sup_{f\in F} \deg(f)$ and $\sup_{e\in E} c(e)^{-1}$.

10.3.1 Preliminaries

We begin with some preliminary estimates that will be used in the proof. We first reduce the statement of Theorem 10.1.2 to the case where the graph is simple and has no *peninsulas*.

Let G be a bounded codegree proper plane network with locally finite dual $G^\dagger$. Recall that a **peninsula** of G is a finite connected component of $G \setminus \{v\}$ for some vertex $v \in V$. In order to apply the theory of circle packing, we first reduce to the case in which G is simple and does not contain a peninsula; this will ensure that the triangulation $T = T(G)$ formed by drawing a star inside each face of G is simple. We will then use the circle packing of T to analyse the WUSF of G.

Lemma 10.3.1. *Let G be a network such that every vertex v of G is contained in a peninsula of G, and let $\mathfrak{F}$ be a sample of WUSF_G. Then $\mathfrak{F}$ is connected and one-ended a.s.*

Proof. Let v_0 be an arbitrary vertex of G, and for each $i \geq 1$ let v_i be a vertex of G such that v_{i-1} is contained in a finite connected component of $G \setminus \{v_i\}$. Let u be another vertex of G and let γ be a path from v_0 to u in G. If u is contained in an infinite connected component of $G \setminus \{v_i\}$, then γ must pass through the vertex v_i. Since γ is finite, we deduce that there exists an integer I such that u is contained in a finite connected component of $G \setminus \{v_i\}$ for all $i \geq I$. Thus, any infinite simple path from u in G must visit v_i for all $i \geq I$. We deduce that every essential spanning forest of G is connected and one-ended. □

Thus, to prove Theorem 10.1.2, it suffices to consider the case that there is some vertex of G that is not contained in a peninsula. In this case, let G' be the simple plane network formed from G by first splitting every edge e of G into a path of length 2 with both edges given the weight $c(e)$, and then deleting every peninsula of the resulting network. The assumption that G has bounded codegrees and edge resistances bounded above ensures that G' does also (indeed, the maximum codegree of G' is at most twice that of G).

Lemma 10.3.2. *Let G be a plane network and let G' be as above. Then every component of the WUSF of G is one-ended a.s. if and only if every component of the WUSF of G' is one-ended a.s.*

Proof. Let $\mathfrak{F}$ be a sample of WUSF_G, and let $\mathfrak{F}'$ be the essential spanning forest of G' defined as follows. For each edge $e \in \mathfrak{F}$ such that neither endpoint of e is contained in a peninsula of G, let both of the edges of the path of length two corresponding to e in G' be included in $\mathfrak{F}'$. For each edge $e \notin \mathfrak{F}$ such that neither endpoint of e is contained in a peninsula of G, choose uniformly and independently exactly one of the two edges of the path of length two corresponding to e in G' to be included in $\mathfrak{F}'$. Then every component of $\mathfrak{F}'$ is one-ended if and only if every component of $\mathfrak{F}$ is one-ended, and it is easily verified that $\mathfrak{F}'$ is distributed according to $\mathsf{WUSF}_{G'}$. □

Thus it suffices to consider the case when G is simple and has no peninsulas. This assumption allows us to circle pack G as follows. Let T be the triangulation obtained by adding a vertex inside each face of G and drawing an edge between this vertex and each vertex of the face it corresponds to. The assumption that G is simple and does not contain a peninsula ensures that T is a simple triangulation. We identify the vertices of T with $V(G) \cup F(G)$, where $F(G)$ is the set of faces of G. Let $P = \{P(v) : v \in V(G)\} \cup \{P(f) : f \in F(G)\}$ be a circle packing of T in either the plane or the unit disc. (Note that this is not the double circle packing of G, which is less useful for us at this stage since we are not assuming that G has bounded degrees.)

We proceed with two geometric lemmas. For each $r' > r > 0$ and z in the carrier of P let $V_z(r, r')$ be the set of vertices v of G such that either the intersection of the circle $P(v)$ with the annulus $A_z(r, r')$ is non-empty or the circle $P(v)$ is contained in the ball $B_z(r)$ and there is a face f incident to v such that the intersection $P(f) \cap A_z(r, r')$ is non-empty.

Lemma 10.3.3 (Existence of a uniformly large number of disjoint annuli)**.** *Suppose that T is CP hyperbolic so that the carrier of P is $\mathbb{D}$. Then the following hold:*

1. *There exists a decreasing sequence $\langle r_n \rangle_{n \geq 0}$ with $r_n \in (0, 1/4)$ such that for every $z \in \mathbb{D}$ with $|z| \geq 1 - r_n$, the sets*
$$V_z(r_i, 2r_i) \qquad 1 \leq i \leq n$$
are disjoint.
2. *If G has bounded degrees, then there exists a constant $C = C(\mathbf{M}_G)$ such that we may take $r_n = C^{-n}$ in (1).*

Proof. We first prove item (1). We construct the sequence recusively, letting $r_0 = 1/8$. Suppose that $\langle r_i \rangle_{i=0}^{n}$ satisfying the conclusion of the lemma have already been chosen, and consider the set of vertices

$$K_n = \left\{ v \in V(G) : r(v) \geq \frac{r_n}{4} \text{ or } r(f) \geq \frac{r_n}{4} \text{ for some face } f \text{ incident to } v \right\}.$$

Since the carrier of P has finite area K_n is a finite set. We define r_{n+1} to be

$$r_{n+1} = \sup \left\{ r \leq r_n/4 \,:\, P(v) \subseteq \overline{B_0(1 - 3r)} \text{ for all } v \in K_n \right\}$$

which is positive since K_n is finite. We claim that $\langle r_i \rangle_{i=0}^{n+1}$ continues to satisfy the conclusion of the lemma. That is, we claim that $V_z(r_{n+1}, 2r_{n+1}) \cap V_z(r_i, 2r_i) = \emptyset$ for every $1 \leq i \leq n$ and every $z \in \mathbb{D}$ with $|z| \geq 1 - r_{n+1}$.

Indeed, let z be such that $|z| \geq 1 - r_{n+1}$ and let $v \in V_z(r_{n+1}, 2r_{n+1})$. By definition of $V_z(r_{n+1}, 2r_{n+1})$, the intersection $P(v) \cap B_z(2r_{n+1})$ is non-empty, and we deduce that $P(v)$ is not contained in the closure of $B_0(1 - 3r_{n+1})$. By definition of r_{n+1}, it follows that $v \notin K_n$ and hence that $r(v) \leq r_n/4$ and $r(f) \leq r_n/4$ for every face f incident to v. Thus, since $r_{n+1} \leq r_n/4$,

$$P(v) \subseteq B_z\left(2r_{n+1} + 2r_n/4\right) \subseteq B_z\left(r_n\right) \text{ and } P(f) \subseteq B_z\left(2r_{n+1} + 2r_n/4\right) \subseteq B_z\left(r_n\right)$$

for every face f incident to v. It follows that $V_z(r_{n+1}, 2r_{n+1}) \cap V_z(r_i, 2r_i) = \emptyset$ for every $1 \leq i \leq n$ as claimed.

To prove (2), observe that, by the Ring Lemma (Theorem 10.2.8), there exists a constant $k = k(\mathbf{M})$ such that for every $C > 1$, every z with $|z| \geq 1 - C^{-m}$, and every $0 \leq n \leq m$, every circle in P that either has centre in $A_z(C^{-n}, 2C^{-n})$ or is tangent to some circle with centre in $A_z(C^{-n}, 2C^{-n})$ has radius at most kC^{-n}. Thus, this set of circles is contained in the ball $B_z((2+4k)C^{-n})$. It follows that taking $C = 1/(4+8k)$ suffices. $\square$

Lastly, we estimate the energy of a random set of vertices that we will frequently use.

Lemma 10.3.4. *Let z be a point in the carrier of P (which may be either $\mathbb{C}$ or $\mathbb{D}$) and U be a uniform random variable on the interval $[1,2]$ and, for each $r > 0$, let μ_r be the law of the random set of vertices $W_z(Ur) := V_z(Ur, Ur)$. Then*

$$\mathcal{E}(\mu_r) \preceq 1$$

uniformly in $r > 0$.

Proof. For a vertex v of G to be included in $W_z(Ur)$, the circle $\{z' \in \mathbb{C} : |z' - z| = Ur\}$ must either intersect the circle $P(v)$ or intersect $P(f)$ for some face f incident to v. We note that the union of $P(v)$ and all the $P(f)$ incident to $P(v)$ is contained in the ball of radius $r(v) + 2\max_{f \sim v} r(f)$ around $z(v)$. Since the codegrees of G are bounded, the Ring Lemma implies that $r(f) \preceq r(v)$ for all incident $v \in V$ and $f \in F$, and so

$$\mu_r(v \in W_z(Ur)) \leq \frac{1}{r} \min\left(2r(v) + 4\max_{f \ni v} r(f),\, r\right) \preceq \frac{1}{r}\min\{r(v), r\}. \tag{10.3.1}$$

We claim that

$$\sum_{v \in V_z(r,2r)} \min\{r(v),\, r\}^2 \leq 16r^2. \tag{10.3.2}$$

To see this, replace each circle of a vertex in $V_z(r, 2r)$ that has radius larger than r with a circle of radius r that is contained in the original circle and intersects $B_z(2r)$: The circles in this new set still have disjoint interiors, are contained in the ball $B_z(4r)$, and have total area $\pi \sum_{v \in V_z(r,2r)} \min(r(v),\, r)^2$, yielding (10.3.2). The claim follows from (10.3.1) and (10.3.2) by definition of the energy $\mathcal{E}(\mu_r)$. $\square$

We are now ready to prove Theorem 10.1.2.

10.3.2 Proof of Theorem 10.1.2

Let G be a simple proper plane network with bounded codegrees and bounded edge resistances. By Lemma 10.3.2 we can assume that G does not contain a peninsula. Let $\mathfrak{F}$ be a sample of WUSF_G and given an edge $e = (x, y)$ let $\mathscr{A}^e$ be the event that x and y are in distinct infinite connected components of $\mathfrak{F} \setminus \{e\}$. It is clear that every component of $\mathfrak{F}$ is one-ended a.s. if and only if

$$\mathsf{WUSF}_G(e \in \mathfrak{F}, \mathscr{A}^e) = 0 \tag{10.3.3}$$

for every edge e of G.

Consider the triangulation T obtained from G and its circle packing $P = \{P(v) : v \in V(G)\} \cup \{P(f) : f \in F(G)\}$, as described in the previous section. By applying a Möbius transformation, we normalise P by setting the centres $z(x)$ and $z(y)$ to be on the negative and positive real axes respectively, setting the circles $P(x)$ and $P(y)$ to have the origin as their tangency point and, in the parabolic case, fixing the scale by setting $z(y) - z(x) = 1$.

Let $\varepsilon > 0$ be arbitrarily small. If T is CP hyperbolic, let V_ε be the set of vertices of G with $|z(v)| \leq 1-\varepsilon$. Otherwise, T is CP parabolic and we define V_ε to be the set of of vertices of G with $|z(v)| \leq \varepsilon^{-1}$. We also denote by E_ε the set of edges that both endpoints are in V_ε.

Let $\mathscr{B}_\varepsilon^e$ be the event that every component of $\mathfrak{F} \setminus \{e\}$ intersects $V \setminus V_\varepsilon$. On the event $\mathscr{B}_\varepsilon^e$, we define η^x to be the rightmost path in $\mathfrak{F} \setminus \{e\}$ from x to $V \setminus V_\varepsilon$ when looking at x from y, and η^y to be the leftmost path in $\mathfrak{F} \setminus \{e\}$ from y to $V \setminus V_\varepsilon$ when looking at y from x. Note that the paths η_x and η_y are not necessarily disjoint. Nonetheless, concatenating the reversal of η^x with e and η^y separates V_ε into two sets of vertices, $\mathcal{L}$ and $\mathcal{R}$, which are to the left and right of e (when viewed from x to y) respectively. See Figure 10.3 for an illustration of the case when η_x and η_y are disjoint (when they are not, $\mathcal{L}$ is a “bubble” separated from $V \setminus V_\varepsilon$).

Let K be the set of edges that touch a vertex in $\mathcal{L}$ or edges that belong to $\eta_x \cup \eta_y$. Note that edges of K do not touch the vertices $\mathcal{R}$. The condition that η^x and η^y are the rightmost and leftmost paths to $V \setminus V_\varepsilon$ from x and y is equivalent to the condition that K does not contain any open path from x to $V \setminus V_\varepsilon$ other than η^x, and does not contain any open path from y to $V \setminus V_\varepsilon$ other than η^y. It follows that, if we define $K = E$ on the complement of the event $\mathscr{B}_\varepsilon^e$, then K is a local set for $\mathfrak{F}$.[10].

Let $\mathscr{A}_\varepsilon^e$ denote the event that $\mathscr{B}_\varepsilon^e$ occurs and that K does not contain an open path from x to y, or equivalently, that η_x and η_y are disjoint. Note that $\mathscr{A}_\varepsilon^e$ is measurable with respect to $\mathcal{F}_K$ (as defined in Section 10.2.2), and that $\mathscr{A}^e = \cap_{\varepsilon>0}\mathscr{A}_\varepsilon^e$. Thus,

$$\mathsf{WUSF}_G(e \in \mathfrak{F}\,, \mathscr{A}^e) \leq \mathsf{WUSF}_G(e \in \mathfrak{F} \mid \mathscr{A}_\varepsilon^e) = \mathbb{E}\big[\mathsf{WUSF}_G(e \in \mathfrak{F} \,|\, \mathcal{F}_K,\, \mathscr{A}_\varepsilon^e)\big]\,.$$

Let $\mathcal{C}$ denote the set of closed edges of $\mathfrak{F}$ in K and let $\mathcal{O}$ denote the set of open edges of $\mathfrak{F}$ in K. By the strong spatial Markov property (Proposition 10.2.1), conditioned on $\mathcal{F}_K$ and the event $\mathscr{A}_\varepsilon^e$, the law of $\mathfrak{F}$ is equal to the union of $\mathcal{O}$ with a sample of the WUSF of the network $(G-\mathcal{C})/\mathcal{O}$ obtained from G by deleting all the revealed closed edges and contracting all the revealed open edges. In particular, by Kirchhoff's Effective Resistance Formula (Theorem 10.2.3),

$$\mathsf{WUSF}_G(e \in \mathfrak{F} \,|\, \mathcal{F}_K,\, \mathscr{A}_\varepsilon^e) \leq c(e)\mathscr{R}_{\mathrm{eff}}^{\mathrm{W}}(\eta^x \leftrightarrow \eta^y;\; G-\mathcal{C}) \tag{10.3.4}$$

Since the edge e was arbitrary to prove (10.3.3) (and hence Theorem 10.1.2) it suffices to prove that there is a upper bound on the effective resistance appearing in (10.3.4) that tends to zero as $\varepsilon \to 0$ uniformly in $\mathcal{F}_K$. We perform this analysis now according to whether T is CP hyperbolic or

[10]Indeed, K can be explored algorithmically, without querying the status of any edge in $E \setminus K$, by performing a right-directed depth-first search of x's component in $\mathfrak{F}$ and a left-directed depth-first search of y's component in $\mathfrak{F}$, stopping each search when it first leaves V_ε.

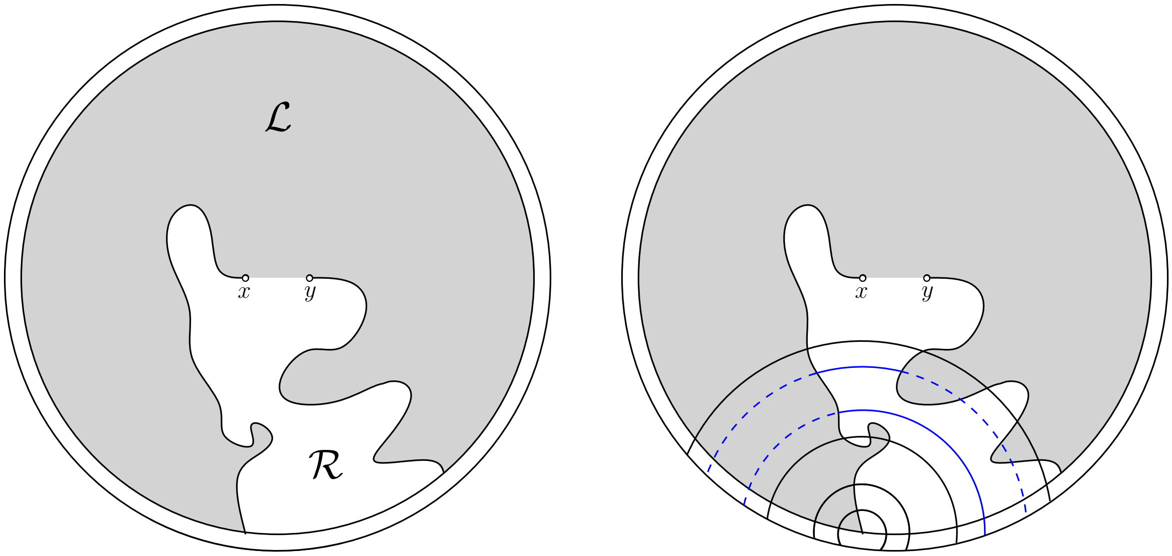

Figure 10.3: Illustration of the proof of Theorem 10.1.2 in the case that T is CP hyperbolic. Left: On the event $\mathscr{B}_\varepsilon^e$, the paths η^x and η^y split V_ε into two pieces, $\mathcal{L}$ and $\mathcal{R}$. Right: We define a random set containing a path (solid blue) from η^x to $\eta^y \cup \{\infty\}$ in $G \setminus \mathcal{C}$ using a random circle (dashed blue). Here we see two examples, one in which the path ends at η^y, and the other in which the path ends at the boundary (i.e., at infinity).

parabolic.

Proof of Theorem 10.1.2, hyperbolic case. Suppose that T is CP hyperbolic and let v^x be the endpoint of the path η^x and put $z_0 = z(v^x)$. Let $\langle r_n \rangle_{n \geq 0}$ be as in Lemma 10.3.3 and let $n(\varepsilon)$ be the maximum n such that $\varepsilon < r_n$. On the event $\mathscr{A}_\varepsilon^e$, for each $1 - |z_0| \leq r \leq 1/4$, we claim that the set $W_{z_0}(r)$, as defined in Lemma 10.3.4, contains a path in G from η^x to $\eta^y \cup \{\infty\}$ that is contained in $\mathcal{R} \cup \eta^x \cup \eta^y$, and is therefore a path in $G \setminus \mathcal{C}$.

Indeed, consider the arc $\mathfrak{A}'(r) = \{z \in \overline{\mathbb{D}} : |z - z_0| = r\}$, parameterised in the clockwise direction. Let $\mathfrak{A}(r)$ be the subarc of $\mathfrak{A}'(r)$ beginning at the last time that $\mathfrak{A}'(r)$ intersects a circle corresponding to a vertex in the trace of η^x, and ending at the first time after this time that $\mathfrak{A}'(r)$ intersects either $\partial\mathbb{D}$ or a circle corresponding to a vertex in the trace of η^y (see Figure 10.3). Thus, on the event $\mathscr{A}_\varepsilon^e$, the set of vertices of T whose corresponding circles in P are intersected by $\mathfrak{A}(r)$ contains a path in T from η^x to $\eta^y \cup \{\infty\}$, for every $1 - |z_0| \leq r \leq 1/4$. (Indeed, for Lebesgue a.e. $1 - |z_0| \leq r \leq 1/4$, the arc $\mathfrak{A}(r)$ is not tangent to any circle in P, and in this case the set is precisely the trace of a simple path in T.) To obtain a path in G rather than T, we divert the path clockwise around each face of G. That is, whenever the path passes from a vertex u of G to a face f of G and then to a vertex v of G, we replace this section of the path with the list of vertices of G incident to f that are between u and v in the counterclockwise order. This construction shows that the subgraph of $G \setminus \mathcal{C}$ induced by the set $W_{z_0}(r)$ contains a path from η^x to $\eta^y \cup \{\infty\}$, as claimed.

By Lemma 10.3.3 the measures μ_{r_i} are supported on sets contained in the disjoint sets $V_z(r_i, 2r_i)$.

Furthermore, by Lemma 10.2.4 and Lemma 10.3.4 we have

$$\mathscr{R}_{\mathrm{eff}}^{\mathrm{W}}(\eta^x \to \eta^y \cup \{\infty\};\, G \setminus \mathcal{C}) \preceq \mathcal{E}\left(\frac{1}{n(\varepsilon)}\sum_{i=1}^{n(\varepsilon)} \mu_{r_i}\right) = \frac{1}{n(\varepsilon)^2}\sum_{i=1}^{n(\varepsilon)} \mathcal{E}(\mu_{r_i}) \preceq \frac{1}{n(\varepsilon)}$$

and hence, by symmetry,

$$\mathscr{R}_{\mathrm{eff}}^{\mathrm{W}}(\eta^y \to \eta^x \cup \{\infty\};\, G \setminus \mathcal{C}) \preceq \frac{1}{n(\varepsilon)}.$$

Applying Lemma 10.2.2 and (10.3.4), we have

$$\mathsf{WUSF}(e \in \mathfrak{F} \,|\, \mathcal{F}_K,\, \mathscr{B}_\varepsilon^e) \preceq \frac{c(e)}{n(\varepsilon)} \tag{10.3.5}$$

which by Lemma 10.3.3 converges to zero as $\varepsilon \to 0$, completing the proof of Theorem 10.1.2 in the case that T is CP hyperbolic. If G has bounded degrees, then combining (10.3.5) with Lemma 10.3.3 implies that there exists a positive constant $C = C(\mathbf{M})$ such that

$$\mathsf{WUSF}(e \in \mathfrak{F} \,|\, \mathcal{F}_K,\, \mathscr{B}_\varepsilon^e) \leq C\frac{c(e)}{\log(1/\varepsilon)} \tag{10.3.6}$$

for all $\varepsilon \leq 1/2$. □

Proof of Theorem 10.1.2, parabolic case. Suppose that T is CP parabolic. Let $r_1 = 1$ and define $\langle r_n \rangle_{n \geq 1}$ recursively by

$$r_n = 1 + \inf\{r \geq r_{n-1} : V_0(r, 2r) \cap V_0(r_{n-1}, 2r_{n-1}) = \emptyset\}.$$

Since $V_0(r_{n-1}, 2r_{n-1})$ is finite, this infimum is finite. By definition, the sets $V_0(r_i, 2r_i)$ are disjoint. A similar analysis to the hyperbolic case shows that, on the event $\mathscr{A}_\varepsilon^e$, for each $r \geq 1$, the set $W_0(r)$ contains a path in G from η^x to η^y that is contained in $\mathcal{R} \cup \eta^x \cup \eta^y$, and is therefore a path in $G \setminus \mathcal{C}$. For each $\varepsilon > 0$, let $n(\varepsilon)$ be the maximal n such that $r_n \leq \varepsilon^{-1}$. Then on the event $\mathscr{A}_\varepsilon^e$, by Lemma 10.2.4 and Lemma 10.3.4,

$$\mathscr{R}_{\mathrm{eff}}^{\mathrm{F}}(\eta^x \leftrightarrow \eta^y;\, G \setminus \mathcal{C}) \preceq \mathcal{E}\left(\frac{1}{n(\varepsilon)}\sum_{i=1}^{n(\varepsilon)} \mu_{r_i}\right) \leq \frac{1}{n(\varepsilon)^2}\sum_{i=1}^{n(\varepsilon)} \mathcal{E}(\mu_{r_i}) \preceq \frac{1}{n(\varepsilon)}.$$

Thus, by (10.3.4),

$$\mathsf{WUSF}_G(e \in \mathfrak{F} \,|\, \mathcal{F}_K,\, \mathscr{B}_\varepsilon^e) \preceq \frac{c(e)}{n(\varepsilon)}. \tag{10.3.7}$$

The right hand side converges to zero as $\varepsilon \to 0$, completing the proof of Theorem 10.1.2. □

Remark 10.3.5. Since the random sets used in the CP parabolic case above are always finite, they can also be used to bound free effective resistances. Therefore, by repeating the proof above with the FUSF in place of the WUSF, we deduce that every component of the FUSF of G is one-ended almost surely if T is CP parabolic. Since the FUSF stochastically dominates the WUSF, and an essential spanning forest with one-ended components does not contain any strict subgraphs that are also essential spanning forests, we deduce that if T is CP parabolic, then the FUSF and WUSF of G coincide. In particular, using [44, Theorem 7.3], we obtain the following.

Theorem 10.3.6. *Let T be a CP parabolic proper plane triangulation. Then T does not admit non-constant harmonic functions of finite Dirichlet energy.*

10.3.3 The FUSF is connected.

Proof of Theorem 10.1.1. First suppose that $G^\dagger$ is locally finite. Since G has bounded degrees and bounded conductances, $G^\dagger$ has bounded codegrees and bounded resistances. Thus, Theorem 10.1.2 implies that every component of the WUSF on $G^\dagger$ is one-ended a.s. and consequently, by Theorem 10.2.5, that the FUSF of G is connected a.s.

Now suppose that G does not have locally finite dual. In this case, we form a plane network (G', c') from G by adding edges to triangulate the infinite faces of G while keeping the degrees bounded. We enumerate these additional edges $\langle e_i \rangle_{i\geq 1}$ and define conductances

$$c'(e_i) = 2^{-i-1}\mathscr{R}_{\text{eff}}^{\text{F}}(e_i^- \leftrightarrow e_i^+; G)^{-1}.$$

Since G' has bounded degrees, bounded conductances and locally finite dual, its FUSF is connected a.s. By Kirchhoff's Effective Resistance Formula (Theorem 10.2.3), Rayleigh monotonicity and the union bound, the probability that the FUSF of G' contains any of the additional edges e_i is at most

$$\sum_{i\geq 0} c'(e_i)\mathscr{R}_{\text{eff}}^{\text{F}}(e_i^- \leftrightarrow e_i^+; G') \leq \sum_{i\geq 0} c'(e_i)\mathscr{R}_{\text{eff}}^{\text{F}}(e_i^- \leftrightarrow e_i^+; G) \leq 1/2.$$

In particular, there is a positive probability that none of the additional edges are contained in the FUSF of G'. The conditional distribution of the FUSF of G' on this event is FUSF_G by the spatial Markov property, and it follows that the FUSF of G is connected a.s. $\square$

10.4 Critical Exponents

10.4.1 The Ring Lemma for double circle packings

In this section we extend the Ring Lemma to double circle packings. While unsurprising, we were unable to find such an extension in the literature.

Theorem 10.4.1 (Ring Lemma)**.** *There exists a family of positive constants $\langle k_{n,m} : n \geq 3, m \geq 3\rangle$ such that if $(P, P^\dagger)$ is a double circle packing in $\mathbb{C} \cup \{\infty\}$ of a simple, 3-connected plane graph G and v is a vertex of G, then for every $f \in F$ incident to v such that $P(v)$ does not contain ∞, then*

$$r(v)/r(f) \leq k_{\deg(v), \max_{g \perp v} \deg(g)}$$

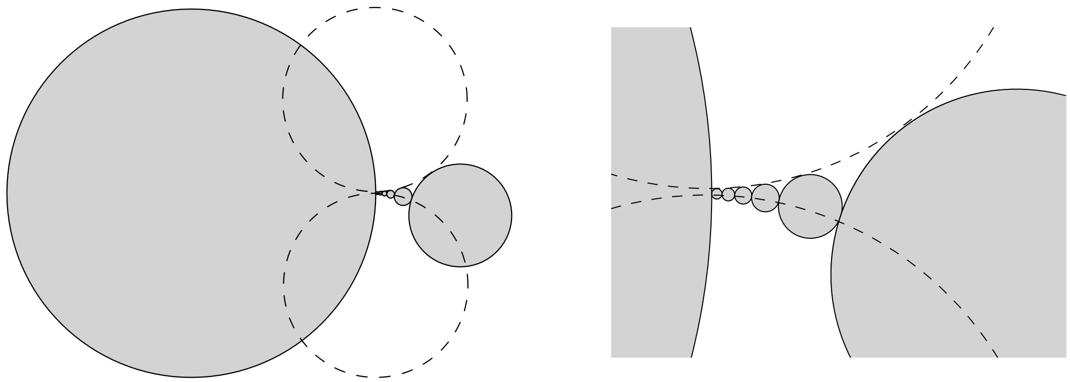

Figure 10.4: Proof of the Double Ring Lemma. If two dual circles are close but do not touch, there must be many primal circles contained in the crevasse between them. This forces the two dual circles to each have large degree. The right-hand figure is a magnification of the left-hand figure.

where $g \perp v$ means that the face g is incident fo v.

As for triangulations, the Ring Lemma immediately implies that whenenever a simple, 3-connected, CP hyperbolic proper plane network G with bounded degrees and codegrees is circle packed in $\mathbb{D}$, the hyperbolic radii of the discs in $P \cup P^\dagger$ are bounded above by a constant $C = C(\mathbf{M}_G, \mathbf{M}_G^\dagger)$.

Proof of Theorem 10.5.1. Let n be the degree of v and let m be the maximum degree of the faces incident to v in G. We may assume that $r(v) = 1$. Note that for each two distinct discs $P^\dagger(f), P^\dagger(f') \in P^\dagger$ that are not tangent, there is at most one disc $P(u) \in P$ that intersects both $P^\dagger(f)$ and $P^\dagger(f')$, while if $P^\dagger(f)$ and $P^\dagger(f')$ are tangent, there exist exactly two discs in P that intersect both $P^\dagger(f)$ and $P^\dagger(f')$. For each two faces f and f' incident to v, the complement $\partial P(v) \setminus \Big(P^\dagger(f_1) \cup P^\dagger(f_2)\Big)$ is either a single arc (if $P^\dagger(f)$ and $P^\dagger(f')$ are tangent), or is equal to the union of two arcs (if $P^\dagger(f)$ and $P^\dagger(f')$ are not tangent).

We claim that there exists an function $\psi_m(\cdot,\cdot) : (0,\infty)^2 \to (0, 2\pi]$, increasing in both coordinates, such that if f and f' are two distinct faces of G incident to v, then each of the (one or two) arcs forming the complement $\partial P(v) \setminus (P^\dagger(f_1) \cup P^\dagger(f_2))$ have length at least $\psi_m(r(f), r(f'))$. Indeed, let $r(f)$ and $r(f')$ be fixed, and suppose that one of the arcs forming the complement $\partial P(v) \setminus P^\dagger(f) \cup P^\dagger(f')$ is extremely small, with length ε. Let the primal circles incident to f be enumerated $v_1, \ldots, v_{\deg(f)}$, where $v_1 = v$, $P(v_2)$ is the primal circle that is tangent to $P(v)$ and intersects $P^\dagger(f)$ on the same side as the small arc, $P(v_3)$ is the next primal circle that is tangent to $P(v_2)$ and intersects $P^\dagger(f)$, and so on. Since ε is small, $P(v_2)$ must also be very small, as it does not intersect $P^\dagger(f')$. Similarly, if ε is sufficiently small, $P(v_3)$ must also be small, since it also does not intersect $P^\dagger(f')$. See Figure 10.5. Applying this argument recursively, we see that, if ε is sufficiently small, then the circles $P(v_2), \ldots, P(v_{\deg(f)})$ are collectively too small to contain $\partial P^\dagger(f) \setminus P(v)$ in their union, a contradiction. We write $\psi_m(r(f), r(f'))$ for the minimal ε that is not ruled out as impossible by this argument.

Let the faces incident to the vertex v be indexed in clockwise order $f_1, \ldots, f_n$, where $n = \deg(v)$

and f_1 has maximal radius among the faces incident to v. For each face f incident to v, the arc $\partial P(v) \cap P^\dagger(f)$ has length $2\tan^{-1}(r(f))$. Since $\partial P(v) = \cup_{i=1}^n (\partial P(v) \cap P^\dagger(f_i))$, we deduce that $r(f_1)$ is bounded below by $\tan(\pi/n)$. By definition of ψ_m, we have that $r(f_k)$ satisfies

$$\psi_m\left(\tan(\pi/n), r(f_k)\right) \leq \psi_m\left(r(f_1), r(f_k)\right) \leq \sum_{i=2}^{k-1} 2\tan^{-1}\left(r(f_i)\right) \tag{10.4.1}$$

for all $3 \leq k \leq n$. For each such k, (10.5.1) yields an implicit upper bound on $r(f_k)$ which converges to zero as $r(f_2)$ converges to zero. This in turn yields a uniform lower bound on $r(f_2)$: if $r(f_2)$ were sufficiently small, the bound (10.5.1) would imply that $\sum_{k=2}^n 2\tan^{-1}(r(f_k))$ would be less than π, a contradiction (since we have trivially that $2\tan^{-1}(r(f_1)) \leq \pi$). We obtain uniform lower bounds on $r(f_k)$ for each $3 \leq k \leq n$ by repeating the above argument inductively. □

10.4.2 Good embeddings of planar graphs

If G is a plane graph and $(P, P^\dagger)$ is a double circle packing of G in a domain $D \subseteq \mathbb{C}$, then drawing straight lines between the centres of the circles in P yields a proper embedding of G in D in which every edge is a straight line. We call such an embedding a **proper embedding of G with straight lines** in D. Following [17], a proper embedding of a graph G with straight lines in a domain $D \subseteq \mathbb{C}$ is said to be **η-good** if the following conditions are satisfed:

1. **(No near-flat, flat, or reflexive angles.)** All internal angles of every face in the drawing are at most $\pi - \eta$. In particular, every face is convex.
2. **(Adjacent edges have comparable lengths.)** For every pair of edges e, e' sharing a common endpoint, the ratio of the lengths of the straight lines corresponding to e and e' in the drawing is at most η^{-1}.

The Ring Lemma (Theorem 10.5.1) has the following immediate corollary.

Corollary 10.4.2. *The straight line embedding given by any double circle packing of a plane graph G with bounded degrees and bounded codegrees in a domain $D \subseteq \mathbb{C}$ is η-good for some positive $\eta = \eta(\mathbf{M}_G)$.*

We remark that, in contrast, the embedding of G obtained by circle packing the triangulation $T(G)$ formed by drawing a star inside each face of G (and then erasing these added vertices) does not necessarily yield a good embedding of G.

For the remainder of this section and in Sections 10.5.3, 10.5.5 and 10.5.6, G will be a fixed transient, simple, 3-connected proper plane network with bounded codegrees and bounded local geometry, and $(P, P^\dagger)$ will be a double circle packing of G in $\mathbb{D}$. We will write $\preceq$, $\succeq$ and $\asymp$ to denote inequalities or equalities that hold up to positive multiplicative constants depending only upon $\mathbf{M}$. We will also fix an edge $e = (x, y)$ of G and, by applying a Möbius transformation if necessary, normalise $(P, P^\dagger)$ by setting the centres $z(x)$ and $z(y)$ to be on the negative real axis and positive real axis respectively and setting the circles $P(x)$ and $P(y)$ to have the origin as their tangency point.

In [17] and [69], several estimates are established that allow one to compare the random walk on a good embedding of a planar graph with Brownian motion. The following estimate, proven for general good embeddings of proper plane graphs in [21, Theorems 1.4 and 1.5], is of central importance to the proofs of the lower bounds in Theorem 10.1.4 and Theorem 10.1.5. Recall that $\sigma(v)$ is defined to be $1-|z(v)|$.

Theorem 10.4.3 (Diffusive Time Estimate)**.** *There exists a constant $C_1 = C_1(\mathbf{M}) \geq 1$ such that the following holds. For each vertex v of G, let $\langle X_n\rangle_{n\geq 0}$ be a random walk on G started at v, let $r(v) \leq r \leq C_1^{-1}\sigma(v)$ and let T_r be the first time n that $|z(X_n) - z(v)| \geq r$. Then*

$$\mathbf{E}_v \sum_{n=0}^{T_r} r(X_n)^2 \asymp r^2.$$

It will be shown in a forthcoming paper of Murugan [personal communication] that the constant C_1 above can in fact be taken to be 1.

Theorem 10.4.4 (Cone Estimate)**.** *Let $\eta = \eta(\mathbf{M}) > 0$ be the constant from Corollary 10.5.2. There exists a positive constant $q_1 = q_1(\mathbf{M})$ such that the following holds. For each vertex v of G, let $\langle X_n\rangle_{n\geq 0}$ be a random walk on G started at v, let $0 \leq r \leq \sigma(v)$ and let T_r be the first time n that $|z(X_n) - z(v)| \geq r$. Then for any interval $I \subset \mathbb{R}/(2\pi\mathbb{Z})$ of length at least $\pi - \eta$ we have*

$$\mathbf{P}_v\Big(\arg\big(z(X_{T_r}) - z(v)\big) \in I\Big) \geq q_1.$$

We also apply the following result of Benjamini and Schramm [47, Lemma 5.3]; their proof was given for simple triangulations but extends immediately to our setting.

Theorem 10.4.5 (Convergence Estimate [47])**.** *For every vertex v of G, we have that*

$$\mathbf{P}_v\left(\big|z(X_n) - z(v)\big| \geq t\sigma(v) \textit{ for some } n \geq 0\right) \preceq \frac{1}{\log t}.$$

We remark that the logarithmic decay in Theorem 10.5.5 is not sharp. Sharp polynomial estimates of the same quantity have been obtained by Chelkak [69, Corollary 7.9].

10.5 Critical exponents

10.5.1 The ring lemma for double circle packings

In this section we extend the Ring Lemma to double circle packings. While unsurprising, we were unable to find such an extension in the literature.

Theorem 10.5.1 (Ring Lemma)**.** *There exists a family of positive constants $\langle k_{n,m} : n \geq 3, m \geq 3\rangle$ such that if $(P, P^\dagger)$ is a double circle packing of a simple, 3-connected plane graph G and v is a vertex of G such that $P(v)$ does not contain ∞, then*

$$r(v)/r(f) \leq k_{\deg(v), \max_{g \perp v} \deg(g)}$$

for all $f \in F$ incident to v.

As for triangulations, the Ring Lemma immediately implies that whenenever a simple, 3-connected, CP hyperbolic proper plane network G with bounded degrees and codegrees is circle packed in $\mathbb{D}$, the hyperbolic radii of the discs in $P \cup P^\dagger$ are bounded above by a constant $C = C(\mathbf{M}_G, \mathbf{M}_G^\dagger)$.

Proof of Theorem 10.5.1. Let n be the degree of v and let m be the maximum degree of the faces incident to v in G. We may assume that $r(v) = 1$. Note that for each two distinct discs $P^\dagger(f), P^\dagger(f') \in P^\dagger$ that are not tangent, there is at most one disc $P(u) \in P$ that intersects both $P^\dagger(f)$ and $P^\dagger(f')$, while if $P^\dagger(f)$ and $P^\dagger(f')$ are tangent, there exist exactly two discs in P that intersect both $P^\dagger(f)$ and $P^\dagger(f')$. For each two faces f and f' incident to v, the complement $\partial P(v) \setminus \Big(P^\dagger(f_1) \cup P^\dagger(f_2)\Big)$ is either a single arc (if $P^\dagger(f)$ and $P^\dagger(f')$ are tangent), or is equal to the union of two arcs (if $P^\dagger(f)$ and $P^\dagger(f')$ are not tangent).

We claim that there exists an function $\psi_m(\cdot,\cdot) : (0,\infty)^2 \to (0, 2\pi]$, increasing in both coordinates, such that if f and f' are two distinct faces of G incident to v, then each of the (one or two) arcs forming the complement $\partial P(v) \setminus (P^\dagger(f_1) \cup P^\dagger(f_2))$ have length at least $\psi_m(r(f), r(f'))$. Indeed, let $r(f)$ and $r(f')$ be fixed, and suppose that one of the arcs forming the complement $\partial P(v) \setminus P^\dagger(f) \cup P^\dagger(f')$ is extremely small, with length ε. Let the primal circles incident to f be enumerated $v_1, \ldots, v_{\deg(f)}$, where $v_1 = v$, $P(v_2)$ is the primal circle that is tangent to $P(v)$ and intersects $P^\dagger(f)$ on the same side as the small arc, $P(v_3)$ is the next primal circle that is tangent to $P(v_2)$ and intersects $P^\dagger(f)$, and so on. Since ε is small, $P(v_2)$ must also be very small, as it does not intersect $P^\dagger(f')$. Similarly, if ε is sufficiently small, $P(v_3)$ must also be small, since it also does not intersect $P^\dagger(f')$. See Figure 10.5. Applying this argument recursively, we see that, if ε is sufficiently small, then the circles $P(v_2), \ldots, P(v_{\deg(f)})$ are collectively too small to contain $\partial P^\dagger(f) \setminus P(v)$ in their union, a contradiction. We write $\psi_m(r(f), r(f'))$ for the minimal ε that is not ruled out as impossible by this argument.

Let the faces incident to the vertex v be indexed in clockwise order $f_1, \ldots, f_n$, where $n = \deg(v)$ and f_1 has maximal radius among the faces incident to v. For each face f incident to v, the arc $\partial P(v) \cap P^\dagger(f)$ has length $2\tan^{-1}(r(f))$. Since $\partial P(v) = \cup_{i=1}^n (\partial P(v) \cap P^\dagger(f_i))$, we deduce that $r(f_1)$ is bounded below by $\tan(\pi/n)$. By definition of ψ_m, we have that $r(f_k)$ satisfies

$$\psi_m\big(\tan(\pi/n), r(f_k)\big) \leq \psi_m\big(r(f_1), r(f_k)\big) \leq \sum_{i=2}^{k-1} 2\tan^{-1}\big(r(f_i)\big) \tag{10.5.1}$$

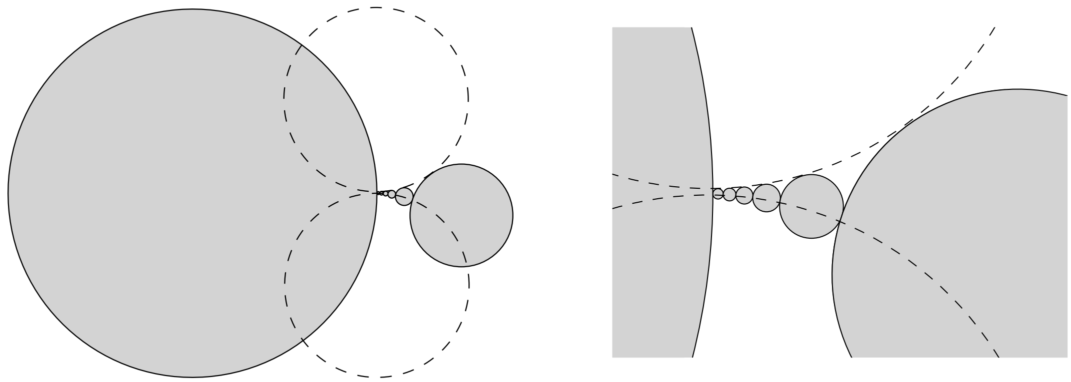

Figure 10.5: Proof of the Double Ring Lemma. If two dual circles are close but do not touch, there must be many primal circles contained in the crevasse between them. This forces the two dual circles to each have large degree. The right-hand figure is a magnification of the left-hand figure.

for all $3 \le k \le n$. For each such k, (10.5.1) yields an implicit upper bound on $r(f_k)$ which converges to zero as $r(f_2)$ converges to zero. This in turn yields a uniform lower bound on $r(f_2)$: if $r(f_2)$ were sufficiently small, the bound (10.5.1) would imply that $\sum_{k=2}^{n} 2\tan^{-1}(r(f_k))$ would be less than π, a contradiction (since we have trivially that $2\tan^{-1}(r(f_1)) \le \pi$). We obtain uniform lower bounds on $r(f_k)$ for each $3 \le k \le n$ by repeating the above argument inductively. □

10.5.2 Good embeddings of planar graphs

If G is a plane graph and $(P, P^\dagger)$ is a double circle packing of G in a domain $D \subseteq \mathbb{C}$, then drawing straight lines between the centres of the circles in P yields a proper embedding of G in D in which every edge is a straight line. We call such an embedding a **proper embedding of G with straight lines** in D. Following [17], a proper embedding of a graph G with straight lines in a domain $D \subseteq \mathbb{C}$ is said to be **η-good** if the following conditions are satisfed:

1. **(No near-flat, flat, or reflexive angles.)** All internal angles of every face in the drawing are at most $\pi - \eta$. In particular, every face is convex.
2. **(Adjacent edges have comparable lengths.)** For every pair of edges e, e' sharing a common endpoint, the ratio of the lengths of the straight lines corresponding to e and e' in the drawing is at most η^{-1}.

The Ring Lemma (Theorem 10.5.1) has the following immediate corollary.

Corollary 10.5.2. *The straight line given by any double circle packing of a plane graph G with bounded degrees and bounded codegrees in a domain $D \subseteq \mathbb{C}$ is η-good for some positive $\eta = \eta(\mathbf{M}_G)$.*

We remark that, in contrast, the embedding of G obtained by circle packing the triangulation $T(G)$ formed by drawing a star inside each face of G (and then erasing these added vertices) does not necessarily yield a good embedding of G.

For the remainder of this section and in Sections 10.5.3, 10.5.5 and 10.5.6, G will be a fixed transient, simple, 3-connected proper plane network with bounded codegrees and bounded local geometry, and $(P, P^\dagger)$ will be a double circle packing of G in $\mathbb{D}$. We will write $\preceq$, $\succeq$ and $\asymp$ to denote inequalities or equalities that hold up to positive multiplicative constants depending only upon $\mathbf{M}$. We will also fix an edge $e = (x, y)$ of G and, by applying a Möbius transformation if necessary, normalise $(P, P^\dagger)$ by setting the centres $z(x)$ and $z(y)$ to be on the negative real axis and positive real axis respectively and setting the circles $P(x)$ and $P(y)$ to have the origin as their tangency point.

In [17] and [69], several estimates are established that allow one to compare the random walk on a good embedding of a planar graph with Brownian motion. The following estimate, proven for general good embeddings of proper plane graphs in [21], is of central importance to the proofs of the lower bounds in Theorem 10.1.4 and Theorem 10.1.5. Recall that $\sigma(v)$ is defined to be $1 - |z(v)|$.

Theorem 10.5.3 (Diffusive Time Estimate)**.** *There exists a constant $C_1 = C_1(\mathbf{M}) \geq 1$ such that the following holds. For each vertex v of G, let $\langle X_n \rangle_{n\geq 0}$ be a random walk on G started at v, let $r(v) \leq r \leq C_1^{-1}\sigma(v)$ and let T_r be the first time n that $|z(X_n) - z(v)| \geq r$. Then*

$$\mathbf{E}_v \sum_{n=0}^{T_r} r(X_n)^2 \asymp r^2.$$

It will be shown in a forthcoming paper of Murugan [personal communication] that the constant C_1 above can in fact be taken to be 1.

Theorem 10.5.4 (Cone Estimate)**.** *Let $\eta = \eta(\mathbf{M}) > 0$ be the constant from Corollary 10.5.2. There exists a positive constant $q_1 = q_1(\mathbf{M})$ such that the following holds. For each vertex v of G, let $\langle X_n \rangle_{n\geq 0}$ be a random walk on G started at v, let $0 \leq r \leq \sigma(v)$ and let T_r be the first time n that $|z(X_n) - z(v)| \geq r$. Then for any interval $I \subset \mathbb{R}/(2\pi\mathbb{Z})$ of length at least $\pi - \eta$ we have*

$$\mathbf{P}_v(\arg(z(X_{T_r}) - z(v)) \in I) \geq q_1.$$

We also apply the following result of Benjamini and Schramm [47, Lemma 5.3]; their proof was given for simple triangulations but extends immediately to our setting.

Theorem 10.5.5 (Convergence Estimate [47])**.** *For every vertex v of G, we have that*

$$\mathbf{P}_v\left(|z(X_n) - z(v)| \geq t\sigma(v) \text{ for some } n \geq 0\right) \preceq \frac{1}{\log t}.$$

We remark that the logarithmic decay in Theorem 10.5.5 is not sharp. Sharp polynomial estimates of the same quantity have been attained by Chelkak [69, Corollary 7.9].

10.5.3 Preliminary estimates

For each vertex v of G, let $a_{\mathbb{H}}(v)$ denote the hyperbolic area of $P(v)$. Note that, since the hyperbolic radii of the discs in P are bounded from above by the Ring Lemma (Theorem 10.5.1),

$$a_{\mathbb{H}}(v) \asymp r_{\mathbb{H}}(v)^2 \asymp \sigma(v)^{-2} r(v)^2 \tag{10.5.2}$$

for every vertex v of G. For the same reason, there exists a constant $s = s(\mathbf{M}) \in (0, 1/2]$ such that for every $\varepsilon > 0$, the set

$$\{v \in V : z(v) \in A_0(1-\varepsilon, 1-s\varepsilon)\}$$

disconnects e from ∞ in G. For each $\varepsilon > 0$, let

$$W_\varepsilon := \left\{v \in V : z(v) \in A_0(1-\varepsilon, 1-s^3\varepsilon)\right\}.$$

In the following section, we will wish to estimate sums of the form

$$\sum_{u \in A} a_{\mathbb{H}}(u) \mathbf{P}_u(\tau_B < \infty) \tag{10.5.3}$$

where A and B are subsets of V. In this section, we prove that, when A is a subset of W_ε for some $\varepsilon > 0$, we can estimate the sum (10.5.3) in terms of geometric quantities.

Lemma 10.5.6 (Cone and half-plane estimate)**.** *Let $r \in (0,1]$, and let $\eta = \eta(\mathbf{M}) > 0$ be the constant from Corollary 10.5.2. There exist positive constants $\delta_1 = \delta_1(\mathbf{M}, r)$ and $q_2 = q_2(\mathbf{M}, r)$, both increasing in r, such that the following holds. For each vertex v of G and $r \in (0,1]$ let $\mathfrak{H}_1$ be the cone*

$$\mathfrak{H}_1 = \mathfrak{H}_1(v) = \{z \in \mathbb{C} : |\arg z - \arg z(v)| \leq \pi/2 - \eta/4\},$$

and let $\mathfrak{H}_2 = \mathfrak{H}_2(v, r)$ be the half-plane containing $z(v)$ whose boundary is the unique straight line with distance $r\sigma(v)$ to $z(v)$ that is orthogonal to the line connecting $z(v)$ to the origin. Then, letting $\tau_{\varepsilon\sigma(v)} = \tau_{W_{\varepsilon\sigma(v)}}$ be the first time the walk visits $W_{\varepsilon\sigma(v)}$, we have that

$$\mathbf{P}_v\left(z(X_n) \in \mathfrak{H}_2 \text{ for all } n \text{ and } z(X_n) \in \mathfrak{H}_1 \text{ for all } n \geq \tau_{\varepsilon\sigma(v)}\right) \geq q_2. \tag{10.5.4}$$

for all $\varepsilon \leq \delta_1$.

Proof. Let $\mathfrak{H}_1'$ and $\mathfrak{H}_1''$ be the slightly thinner cones $\{z \in \mathbb{C} : |\arg(z) - \arg(z(v))| \leq \pi/2 - \eta/3\}$ and $\{z \in \mathbb{C} : |\arg(z) - \arg(z(v))| \leq \pi/2 - \eta/2\}$. By Theorem 10.5.5 there exists some small $\delta' = \delta'(\mathbf{M}) > 0$ such that for each vertex $u \in \mathfrak{H}_1''$ with $\sigma(u) \leq \delta'\sigma(v)$, the random walk started at u has probability at least $1/2$ never to leave $\mathfrak{H}_1'$.

Define numbers $\langle \rho_i \rangle_{i \geq 0}$ and stopping times $\langle T_i \rangle_{i \geq 0}$ by letting $T_0 = 0$ and recursively setting ρ_i to be the distance between $z(X_{T_i})$ and the boundary of $\mathfrak{H}_2 \cap \mathbb{D}$, and T_i to be

$$T_i = \min\left\{n \geq 0 : |z(X_n) - z(X_{T_{i-1}})| \geq \rho_{i-1}\right\}.$$

Let $\mathscr{A}_i$ be the event that $\sigma(X_{T_{i+1}}) \leq \sigma(X_{T_i}) - \rho_i \sin(\eta/2)$, $d_{\mathbb{C}}(X_{T_{i+1}}, \partial\mathfrak{H}_2) \geq d_{\mathbb{C}}(X_{T_i}, \partial\mathfrak{H}_2) + \rho_i \sin(\eta/2)$ and $z(X_{T_{i+1}}) \in \mathfrak{H}_1''$. Applying Theorem 10.5.4 yields that

$$\mathbf{P}_v\left(\mathscr{A}_i \mid z(X_{T_i}) \in \mathfrak{H}_1''\right) \geq q_1 \qquad \text{and hence} \qquad \mathbf{P}_v\Big(\bigcap_{i=0}^{n} \mathscr{A}_i\Big) \geq q_1^n \tag{10.5.5}$$

for every $n \geq 0$. Let $r' = \min\{r, \delta', \sin(\eta/2)\}$ and let $n_0 = \lceil 1/r'^2 \rceil$. If X_{T_i} has distance at least $r'\sigma(v)$ from the boundary of $\mathbb{D} \cap \mathfrak{H}_2$, then $\rho_i \sin(\eta/2) \geq r'^2 \sigma(v)$. It follows that

$$\mathbf{P}_v\Big(z(X_n) \in \mathfrak{H}_2 \text{ for all } n, \lim_{n\to\infty} z(X_n) \in \mathfrak{H}_1'\Big) \geq \frac{1}{2}\mathbf{P}_v\Big(\bigcap_{i=0}^{n_0} \mathscr{A}_i\Big) \geq \frac{1}{2} q_1^{n_0}.$$

Next, applying Theorem 10.5.5, let $\delta_1 = \delta_1(\mathbf{M}, r) > 0$ to be sufficiently small that for each vertex u with $z(u) \notin \mathfrak{H}_1$ and $\sigma(u) \leq \delta_1 \sigma(v)$, the random walk started at u has probability at most $q_1^{n_0}/4$ ever to visit $\mathfrak{H}_1'$. We conclude by observing that the probability appearing in (10.5.4) is at least

$$\mathbf{P}_v\Big(z(X_n) \in \mathfrak{H}_2 \text{ for all } n, \lim_{n\to\infty} z(X_n) \in \mathfrak{H}_1'\Big) - \mathbf{P}_v\Big(z(X_{\tau_{W_{\varepsilon\sigma(v)}}}) \notin \mathfrak{H}_1, \lim_{n\to\infty} z(X_n) \in \mathfrak{H}_1\Big),$$

which is at least $q_1^{n_0}/4$ for all $\varepsilon \leq \delta_1$. We conclude by setting $q_2 = q_1^{n_0}/4$. □

Remark. Although it suffices for our purposes, this argument is rather wasteful. With further effort one can take both q_2 and δ_1 to be polynomials in $1/r$.

Let $p, r \in (0, 1]$. We say that a set $A \subset V$ is (p, r)**-escapable** if for every vertex $v \in A$, the random walk started at v has probability at least p of not returning to A after first leaving the set of vertices whose corresponding discs in P have centres contained in the Euclidean ball of radius $r\sigma(v)$ about $z(v)$. To avoid trivialities, we also declare the empty set to be (p, r)-escapable for all p and r.

Corollary 10.5.7. *Let $r \in (0, 1]$. There exist positive constants $\delta_2 = \delta_2(\mathbf{M}, r)$ and $p_1 = p_1(\mathbf{M}, r)$, both increasing in r, such that the set*

$$\{v \in V : (1 - \delta_2)\varepsilon \leq \sigma(v) \leq \varepsilon\}$$

is (p_1, r)-escapable for every $\varepsilon > 0$.

Proof. Let $\eta = \eta(\mathbf{M})$ be the constant appearing in Corollary 10.5.2, and let

$$\delta_2 = \min\left\{\frac{r}{2}\sin\frac{\eta}{2}, e^{-2C}, 1/4\right\},$$

where $C = C(\mathbf{M})$ is the implicit constant in Theorem 10.5.5. Let $\varepsilon > 0$ be arbitrary, let $v \in \{v \in$

$V : (1-\delta_2)\varepsilon \le \sigma(v) \le \varepsilon\}$, and let T be the stopping time

$$T = \min\left\{n \ge 0 : |z(X_n) - z(v)| \ge r\sigma(v)\right\}.$$

Letting $q_1 = q_1(\mathbf{M})$ be the constant from Theorem 10.5.4, we have that

$$\mathbf{P}_v\left(\sigma(X_T) \le \sigma(v) - r\sin(\eta/2)\sigma(v)\right) \ge q_1. \tag{10.5.6}$$

On the event appearing in the left-hand side of (10.5.6), the half-plane $\mathfrak{H}_2(X_T, \delta_2)$, defined in Lemma 10.5.6, is disjoint from the ball $\{z \in C : 1 - |z| \ge (1-\delta_2)\varepsilon\}$. Thus, the claim follows from (10.5.6) and Lemma 10.5.6 by taking $p = q_1(\mathbf{M})q_2(\mathbf{M}, \delta_2)$, where q_2 is as in Lemma 10.5.6. (Here we are using Lemma 10.5.6 only to lower bound the probability of the first clause defining the event in the left-hand side of (10.5.4).) $\square$

Lemma 10.5.8. *Let A be (p,r)-escapable for some $p \in (0,1)$ and $r \in (0, C_1^{-1})$, where $C_1 = C_1(\mathbf{M})$ is the constant appearing in Theorem 10.5.3. Then for every vertex $u \in V$,*

$$\mathbf{E}_u\left[\sum_{n\ge 0} a_{\mathbb{H}}(X_n)\mathbb{1}(X_n \in A)\right] \preceq \frac{r^2}{p(1-r)^2}. \tag{10.5.7}$$

Proof. Define the sequences of stopping times $\langle T_i^- \rangle_{i\ge 0}$ and $\langle T_i^+ \rangle_{i\ge 0}$ by letting $T_0^- = \tau_A$ and recursively letting

$$T_i^+ = \min\left\{n \ge T_i^- : |z(X_n) - z(X_{T_i^-})| \ge r\sigma(X_{T_i^-})\right\}$$

and $T_i^- = \min\{n \ge T_{i-1}^+ : X_n \in A\}$. Then

$$\mathbf{E}_u\left[\sum_{n\ge 0} a_{\mathbb{H}}(X_n)\mathbb{1}(X_n \in A)\right] \le \sum_{i \ge 0}\mathbf{E}_u\left[\mathbb{1}(T_i^- < \infty)\sum_{n=T_i^-}^{T_i^+ - 1} a_{\mathbb{H}}(X_n)\right].$$

Since $\sigma(X_n) \ge (1-r)\sigma(X_{T_i^-})$ for all $T_i^- \le n < T_i^+$, the Diffusive Time Estimate (Theorem 10.5.3), the strong Markov property, and (10.5.2) imply that

$$\begin{aligned}\mathbf{E}_u&\left[\sum_{n=T_i^-}^{T_i^+ - 1} a_{\mathbb{H}}(X_n) \,\middle|\, T_i^- < \infty,\, X_{T_i^-}\right] \\ &\preceq (1-r)^{-2}\sigma(X_{T_i^-})^{-2}\mathbf{E}_u\left[\sum_{n=T_i^-}^{T_i^+} r(X_n)^2 \,\middle|\, T_i^- < \infty,\, X_{T_i^-}\right] \preceq \frac{r^2}{(1-r)^2}.\end{aligned}$$

Meanwhile, since A is (p,r)-escapable, we have that $\mathbf{P}_u(T_i^- < \infty) \le (1-p)^i$. Combining these estimates yields the desired inequality (10.5.7). $\square$

We say that a set $A \subset V$ is C-**short-lived** if

$$\mathbf{E}_u\left[\sum_{n\geq 0} a_{\mathbb{H}}(X_n)\mathbb{1}(X_n \in A)\right] \leq C$$

for every $u \in V$. Note that if A is a C-short-lived set of vertices, then every subset of A is also C-short-lived. While Lemma 10.5.8 states that escapable sets are short-lived, we also have that unions of boundedly many short-lived sets are short-lived with a larger constant. In particular, letting $s = s(\mathbf{M})$ be as above and letting $\delta = \delta_2(\mathbf{M}, C_1^{-1}(\mathbf{M}))$, we have that

$$W_\varepsilon \subseteq \bigcup_{m=0}^{\lceil 3\log_{1-\delta}(s)\rceil} \{v \in V : (1-\delta)^{m+1}\varepsilon \leq \sigma(v) \leq (1-\delta)^m\varepsilon\}$$

for every $\varepsilon > 0$. Thus, combining Corollary 10.5.7 and Lemma 10.5.8 immediately yields the following.

Corollary 10.5.9. *There exist a constant $C_2 = C_2(\mathbf{M}) > 0$ such that the set W_ε is C_2-short-lived for every $\varepsilon > 0$.*

Lemma 10.5.10. *Let A and B be two sets of vertices in G, and suppose that A is finite and C-short-lived for some $C > 0$. Then*

$$\sum_{v\in A} a_{\mathbb{H}}(v)\mathbf{P}_v(\tau_B < \infty) \preceq C\mathscr{C}^{\mathrm{F}}_{\mathrm{eff}}(A \leftrightarrow B).$$

Proof. Let the stopping times τ_i be defined recursively by setting $\tau_0 = \tau_A$ and $\tau_{i+1} = \min\{t > \tau_i : X_t \in A\}$, so that τ_i is the ith time the random walk $\langle X_n\rangle_{n\geq 0}$ visits A. Then

$$\begin{aligned}\sum_{v\in A} a_{\mathbb{H}}(v)\mathbf{P}_v(\tau_B < \infty) &\leq \sum_{v\in A}\sum_{i\geq 0} a_{\mathbb{H}}(v)\mathbf{P}_v(B \text{ hit between time } \tau_i \text{ and } \tau_{i+1})\\ &= \sum_{v\in A}\sum_{u\in A}\sum_{i\geq 0} a_{\mathbb{H}}(v)\mathbf{P}_v(\tau_i < \infty,\, X_{\tau_i} = u)\mathbf{P}_u(\tau_B < \tau_A^+).\end{aligned}$$

Reversing time then yields that

$$\sum_{v\in A} a_{\mathbb{H}}(v)\mathbf{P}_v(\tau_B < \infty) \preceq \sum_{v\in A}\sum_{u\in A}\sum_{i\geq 0} a_{\mathbb{H}}(v)\mathbf{P}_u(\tau_i < \infty,\, X_{\tau_i} = v)\mathbf{P}_u(\tau_B < \tau_A^+).$$

By exchanging the order of summation, we obtain that

$$\begin{aligned}\sum_{v\in A} a_{\mathbb{H}}(v)\mathbf{P}_v(\tau_B < \infty) &\preceq \sum_{u\in A}\mathbf{E}_u\left[\sum_{n\geq 0} a_{\mathbb{H}}(X_n)\mathbb{1}(X_n \in A)\right]\mathbf{P}_u(\tau_B < \tau_A^+)\\ &\preceq C\sum_{u\in A} c(u)\mathbf{P}_u(\tau_B < \tau_A^+).\end{aligned}$$

To conclude, let $\langle V_j \rangle_{j\geq 1}$ be an exhaustion of V and let G_j be the subgraph of G induced by V_j, with conductances inherited from G, and observe that

$$\begin{aligned}\sum_{u\in A} c(u)\mathbf{P}_u^G(\tau_B < \tau_A^+) &= \lim_{j\to\infty} \sum_{u\in A} c(u)\mathbf{P}_u^G(\tau_B < \min\{\tau_A^+, \tau_{V\setminus V_j}\}) \\ &\leq \lim_{j\to\infty} \sum_{u\in A} c(u)\mathbf{P}_u^{G_j}(\tau_B < \tau_A^+) = \mathscr{C}_{\mathrm{eff}}^{\mathrm{F}}(A \leftrightarrow B). \qquad \square\end{aligned}$$

Recall that $\mathrm{diam}_{\mathbb{C}}(A)$ and $d_{\mathbb{C}}(A,B)$ denote the Euclidean diameter of $\{z(v) : v \in A\}$ and the Euclidean distance between $\{z(v) : v \in A\}$ and $\{z(v) : v \in B\}$ respectively.

Lemma 10.5.11. *Let A and B be disjoint sets of vertices in G. Then*

$$\mathscr{C}_{\mathrm{eff}}^{\mathrm{F}}(A \leftrightarrow B) \preceq \frac{\mathrm{diam}_{\mathbb{C}}(A)^2}{\min\left\{\mathrm{diam}_{\mathbb{C}}(A),\, d_{\mathbb{C}}(A,B)\right\}^2},$$

with the convention that the right-hand side is 1 *if* $\mathrm{diam}_{\mathbb{C}}(A) = 0$.

Proof. Let $D = \mathrm{diam}_{\mathbb{C}}(A)$. If $D = 0$ then A is a single vertex v and $\mathscr{C}_{\mathrm{eff}}^{\mathrm{F}}(A \leftrightarrow B) \leq c(v) \preceq 1$, so assume not. Recall the **extremal length** characterisation of the free effective conductance [173, Exercise 9.42]: For each function $\ell : E \to [0,\infty)$ assigning a non-negative length to every edge e of G, let d_ℓ be the shortest path pseudometric on G induced by ℓ. Then

$$\mathscr{C}_{\mathrm{eff}}^{\mathrm{F}}(A \leftrightarrow B) = \inf\left\{\frac{\sum_{e\in E} c(e)\ell(e)^2}{d_\ell(A,B)^2} : \ell : E \to [0,\infty),\, d_l(A,B) > 0\right\}.$$

Let W be the set of vertices v of G whose corresponding circles intersect the D-neighbourhood of $z(A)$ in $\mathbb{C}$. Define lengths by setting $\ell(e)$ to be

$$\ell(e) = \begin{cases} \min\left\{|z(e^-) - z(e^+)|,\, D\right\} & \text{if } e \text{ has an endpoint in } W \\ 0 & \text{otherwise.}\end{cases}$$

Then

$$d_\ell(A,B) \geq \min\left\{D,\, d_{\mathbb{C}}(A,B))\right\} > 0 \tag{10.5.8}$$

while, since $|z(e^-) - z(e^+)| = r(e^-) + r(e^+)$,

$$\sum_e c(e)\ell(e)^2 \preceq \sum_{v\in W}\sum_{v'\sim v} \min\left\{r(v) + r(v'),\, D\right\}^2 \preceq \sum_{v\in W} \min\left\{r(v),\, D\right\}^2, \tag{10.5.9}$$

where the Ring Lemma (Theorem 10.5.1) is used in the second inequality. As in the proof of Lemma 10.3.4, consider replacing each circle corresponding to a vertex in W that has radius larger than D with a circle of radius D that is contained in the original circle and intersects the D-neighbourhood of $z(A)$. This yields a set of circles contained in the $3D$-neighbourhood of $z(A)$. Comparing the total area of this set of circles with that of the $3D$-neighbourhood of $z(A)$ yields that

$$\sum_{v \in W} \min \left\{ r(v),\, D \right\}^2 \preceq D^2. \tag{10.5.10}$$

We conclude by combining (10.5.8), (10.5.9) and (10.5.10). □

For each $\varepsilon > 0$ and $m \in \mathbb{Z} \setminus \{0\}$, we define

$$U_\varepsilon^m(0) = \left\{ z \in \mathbb{D} : 1 - \varepsilon \leq |z| \leq 1 - s^3\varepsilon \text{ and } \operatorname{sgn}(m)\frac{4^{|m|}}{5^{|m|}}\pi \leq \arg z \leq \operatorname{sgn}(m)\frac{4^{|m|-1}}{5^{|m|-1}}\pi \right\}$$

and define $U_\varepsilon^m(\theta)$ to be the rotated set $e^{i\theta}U_\varepsilon^m$.

Lemma 10.5.12. *There exist universal constants $\delta_3, \delta_4 > 0$ and $k < \infty$ such that*

$$d_{\mathbb{C}}\left(U_\varepsilon^m(\theta), \{re^{i\theta} : r \geq 0\}\right) \geq (2 + \delta_3)\operatorname{diam}_{\mathbb{C}}\left(U_\varepsilon^m(\theta)\right)$$

for all $\theta \in [-\pi, \pi]$, $\varepsilon \leq \delta_4$ and $|m| \leq \log_{5/4}(1/\varepsilon) - k$.

The constants here are more important than usual since we will later need to estimate the difference $\frac{1}{2}d_{\mathbb{C}}\left(U_\varepsilon^m(\theta), \{re^{i\theta} : r \geq 0\}\right) - \operatorname{diam}_{\mathbb{C}}\left(U_\varepsilon^m(\theta)\right)$.

Proof. We begin by calculating, using elementary trigonometry, that

$$\operatorname{diam}_{\mathbb{C}}(U_\varepsilon^m(\theta)) \leq \frac{4^{|m|-1}}{5^{|m|}}\pi + \varepsilon$$

and

$$d_{\mathbb{C}}(U_\varepsilon^m(\theta), \{re^{i\theta} : r \geq 0\}) = \begin{cases} 1 - \varepsilon & |m| \leq 3 \\ (1-\varepsilon)\sin\left(\frac{4^{|m|}}{5^{|m|}}\pi\right) & |m| \geq 4. \end{cases}$$

By concavity of the function $\sin(t)$ on $[0, \pi]$, we have that $\sin(t) \geq 2t/\pi$ for all $t \in [0, \pi/2]$, and in particular

$$\frac{\operatorname{diam}_{\mathbb{C}}(U_\varepsilon^m(\theta))}{d_{\mathbb{C}}(U_\varepsilon^m(\theta), \{re^{i\theta} : r \geq 0\})} \leq \frac{1}{8(1-\varepsilon)}\pi + \frac{5^{|m|}\varepsilon}{2 \cdot 4^{|m|}(1-\varepsilon)} \leq \frac{1}{8(1-\delta_2)}\pi + \frac{4^k\delta_4}{2 \cdot 5^k(1-\delta_4)}$$

for all $\varepsilon \leq \delta_4$ and $|m| \leq \log_{5/4}(1/\varepsilon) - k$. This upper bound is less than $\pi/7 < 1/2$ when δ_4 is sufficiently small and k is sufficiently large. □

10.5.4 Wilson's algorithm

Wilson's algorithm rooted at infinity [44, 228] is a powerful method of sampling the WUSF of an infinite, transient graph by joining together loop-erased random walks. We now give a very brief description of the algorithm (see [173] for a detailed exposition). Let G be a transient network. Let γ be a path in G that visits each vertex of G at most finitely many times. The **loop-erasure** is formed by erasing cycles from γ chronologically as they are created. (The loop-erasure of a random walk path is referred to as **loop-erased random walk** and was first studied by Lawler [161].) Let $\{v_j : j \in \mathbb{N}\}$ be an enumeration of the vertices of G. Let $\mathfrak{F}_0 = \emptyset$ and define a sequence of forests in G as follows:

1. Given $\mathfrak{F}_i$, start an independent random walk from v_{i+1}. Stop this random walk if it hits the set of vertices already included in $\mathfrak{F}_i$, running it forever otherwise.

2. Form the loop-erasure of this random walk path and let $\mathfrak{F}_{i+1}$ be the union of $\mathfrak{F}_i$ with this loop-erased path.

Then the forest $\mathfrak{F} = \bigcup_{i\geq 0} \mathfrak{F}_i$ is a sample of the WUSF of G [44, Theorem 5.1].

10.5.5 Proof of Theorems 10.1.3, 10.1.4 and 10.1.5

Recall that $e = (x, y)$ is a fixed edge. We write $\succeq_e$ to denote a lower bound that holds up to a positive multiplicative constant that depends only on $\mathbf{M}$ and $r_{\mathbb{H}}(x)$, and that is increasing in $r_{\mathbb{H}}(x)$.

Proof of Theorem 10.1.4. Let $\mathfrak{F}$ be the WUSF of G. Let $\mathscr{R}^e_\varepsilon$ be the event that $\mathrm{past}_{\mathfrak{F}}(e)$ contains a vertex whose center is within Euclidean distance ε of the unit circle, and let $\mathscr{D}^e_R$ be the event that $\mathrm{diam}_{\mathbb{H}}(\mathrm{past}_{\mathfrak{F}}(e)) \geq R$. Then, letting $\varepsilon(R) = 1 - \tanh(R/2)$, we have that $\mathscr{R}^e_{\varepsilon(R)/2} \subseteq \mathscr{D}^e_R \subseteq \mathscr{R}^e_{\varepsilon(R)}$. Thus, to prove Theorem 10.1.4, it suffices to prove that

$$\log(1/\varepsilon)^{-1} \preceq_e \mathbb{P}(\mathscr{R}^e_\varepsilon) \preceq \log(1/\varepsilon)^{-1}$$

for all $\varepsilon \leq 1/2$. The proof of Theorem 10.1.2, and in particular of the estimate (10.3.6), adapts immediately to yield the desired upper bound (i.e., if the analysis there is carried out using the double circle packing of G rather than the circle packing of T); the remainder of this proof is devoted to proving the lower bound. This will be done by applying a second moment argument to the random variable

$$Z_\varepsilon = \sum_{v \in W_\varepsilon} a_{\mathbb{H}}(v) \mathbb{1}(v \in \mathrm{past}_{\mathfrak{F}}(e)),$$

which, by definition of W_ε, is positive if and only if $\mathscr{R}^e_\varepsilon$ occurs.

Lemma 10.5.13. *There exists a positive constant $\delta_5 = \delta_5(\mathbf{M}, r_{\mathbb{H}}(x))$, increasing in $r_{\mathbb{H}}(x)$, such that $\mathbb{E}[Z_\varepsilon] \succeq_e 1$ for all $\varepsilon \leq \delta_5$.*

Proof of Lemma 10.5.13. If we generate $\mathfrak{F}$ using Wilson's algorithm, starting with the vertex v, then v is in $\mathrm{past}_{\mathfrak{F}}(e)$ if and only if the loop-erased random walk from v passes through $e = (x, y)$. In particular, we obtain the lower bound

$$\mathbb{P}(v \in \mathrm{past}_{\mathfrak{F}}(e)) \geq \mathbf{P}_v \left(\tau_x < \infty,\, X_{\tau_x+1} = y,\, \langle X_n \rangle_{n \geq \tau_x+1} \text{ disjoint from } \langle X_n \rangle_{n=0}^{\tau_x}\right)$$

and hence, decomposing according to the value of τ_x,

$$\mathbb{E}[Z_\varepsilon] \geq \sum_{v\in W_\varepsilon}\sum_{m\geq 1} a_{\mathbb{H}}(v)\mathbf{P}_v\left(\tau_x = m,\, X_{m+1} = y,\, \langle X_n\rangle_{n\geq m} \text{ disjoint from } \langle X_n\rangle_{n=0}^{m-1}\right).$$

Letting $\langle Y_n\rangle_{n\geq 0}$ be a random walk started at x independent of $\langle X_n\rangle_{n\geq 0}$ and reversing time yields that

$$\mathbb{E}[Z_\varepsilon] \succeq \sum_{v\in W_\varepsilon}\sum_{m\geq 1} a_{\mathbb{H}}(v)\mathbf{P}_x\left(X_m = v,\, \mathscr{C}_m\right) = \mathbf{E}_x\left[\sum_{m\geq 0} a_{\mathbb{H}}(X_m)\mathbb{1}\left(X_m \in W_\varepsilon,\, \mathscr{C}_m\right)\right],$$

where $\mathscr{C}_m$ is the event

$$\mathscr{C}_m = \left\{Y_1 = y, \text{ and } \langle X_n\rangle_{n=1}^m \text{ disjoint from } \langle Y_n\rangle_{n\geq 0}\right\}.$$

(Note that $\langle X_n\rangle_{n=0}^m$ does not return to x after time 0 on the event $\mathscr{C}_m$.) Let τ_1 be the first time that the random walk $\langle X_m\rangle_{m\geq 0}$ visits $\{v\in V : s^2\varepsilon \leq \sigma(v) \leq s\varepsilon\}$, which is finite a.s. by definition of s, and let τ_2 be the first time m after τ_1 that $|z(X_m) - z(X_{\tau_1})| \geq C_1^{-1}s^2\varepsilon$, where $C_1 = C_1(\mathbf{M}) \geq 1$ is the constant from Theorem 10.5.3. Then $X_m \in W_\varepsilon$ for all $\tau_1 \leq m \leq \tau_2$, and so

$$\begin{aligned}\mathbb{E}[Z_\varepsilon] &\succeq \mathbf{E}_x\left[\sum_{m=\tau_1}^{\tau_2} a_{\mathbb{H}}(X_m)\mathbb{1}\left(\mathscr{C}_m\right)\right]\\ &\geq \mathbf{E}_x\left[\sum_{m=\tau_1}^{\tau_2} a_{\mathbb{H}}(X_m)\mathbb{1}\left(d_{\mathbb{C}}(X_{\tau_1}, \{Y_n : n\geq 0\}) > \varepsilon/s,\, \mathscr{C}_{\tau_2}\right)\right].\end{aligned}$$

The events $\mathscr{C}_{\tau_1}\cap\{d_{\mathbb{C}}(X_{\tau_1}, \{Y_n : n\geq 0\}) > \varepsilon/s\}$ and $\mathscr{C}_{\tau_2}\cap\{d_{\mathbb{C}}(X_{\tau_1}, \{Y_n : n\geq 0\}) > \varepsilon/s\}$ are equal, and so

$$\begin{aligned}\mathbb{E}[Z_\varepsilon] &\succeq \mathbf{E}_x\left[\sum_{m=\tau_1}^{\tau_2} a_{\mathbb{H}}(X_m)\mathbb{1}\left(d_{\mathbb{C}}(X_{\tau_1}, \{Y_n : n\geq 0\}) > \varepsilon/s, \mathscr{C}_{\tau_1}\right)\right]\\ &= \mathbf{E}_x\left[\mathbf{E}_x\left[\sum_{m=\tau_1}^{\tau_2} a_{\mathbb{H}}(X_m)\,\middle|\, X_{\tau_1}\right]\mathbb{1}\left(d_{\mathbb{C}}(X_{\tau_1}, \{Y_n : n\geq 0\}) > \varepsilon/s, \mathscr{C}_{\tau_1}\right)\right].\end{aligned}$$

Applying (10.5.2) and the Diffusive Time Estimate (Theorem 10.5.3), we obtain that

$$\mathbf{E}_x\left[\sum_{m=\tau_1}^{\tau_2} a_{\mathbb{H}}(X_m)\,\middle|\, X_{\tau_1}\right] \asymp \varepsilon^{-2}\mathbf{E}_x\left[\sum_{m=\tau_1}^{\tau_2} r(X_m)^2\,\middle|\, X_{\tau_1}\right] \asymp 1,$$

and hence that

$$\mathbb{E}[Z_\varepsilon] \succeq \mathbf{P}_x(d_{\mathbb{C}}(X_{\tau_1}, \{Y_n : n\geq 0\}) > \varepsilon/s, \mathscr{C}_{\tau_1}). \tag{10.5.11}$$

By the Ring Lemma (Theorem 10.5.1), there exists a positive constant $k = k(\mathbf{M})$ such that $\min\{r(x), r(y)\} \geq r_{\mathbb{H}}(x)/k$. Let $\delta_1(\mathbf{M}, r_{\mathbb{H}}(x)/k)$ be the constant appearing in Lemma 10.5.6, and let $\delta_5 = \delta_5(\mathbf{M}, r_{\mathbb{H}}(x)) \leq \delta_1(\mathbf{M}, r_{\mathbb{H}}(x)/k)$ be sufficiently small that every vertex u that has $\sigma(u) \leq \delta_5$ and is contained in the cone $\mathfrak{H}_1(x)$ has distance at least δ_5/s from the half-plane $\mathfrak{H}_2(y, r_{\mathbb{H}}(x)/k)$ (as defined in Lemma 10.5.6). Applying Lemma 10.5.6 to both random walks X and Y conditioned on $Y_1 = y$ yields that

$$\mathbf{P}_x(d_{\mathbb{C}}(X_{\tau_1}, \{Y_n : n \geq 0\}) > \varepsilon/s, \mathscr{C}_{\tau_1}) \succeq q_2(\mathbf{M}, r_{\mathbb{H}}(x)/k)^2 \succeq_e 1$$

for all $\varepsilon \leq \delta_5$, concluding the proof. □

Lemma 10.5.14. $\mathbb{E}[Z_\varepsilon^2] \preceq \mathbb{E}[Z_\varepsilon]\log(1/\varepsilon)$ *for all* $0 < \varepsilon \leq \delta_4$*, where* δ_4 *is the constant from Lemma 10.5.12.*

Proof. For each two vertices u and v in G, let

$$H(u,v) := \left\{w \in V : |z(w) - z(u)| \leq |z(w) - z(v)|\right\}$$

be the set of vertices closer to u than to v with respect to the Euclidean metric on the circle packing. More generally, for a vertex u and a set $A \subset V$, let

$$H(u,A) := \left\{w \in V : |z(w) - z(u)| \leq |z(w) - z(v)| \text{ for some } u \in A\right\} = \bigcup_{v \in A} H(u,v).$$

Note that for every vertex u and set A,

$$d_{\mathbb{C}}(A, H(u,A)) \geq \inf_{v_1, v_2 \in A} d_{\mathbb{C}}(v_2, H(u, v_2)) - d_{\mathbb{C}}(v_1, v_2) \geq \frac{1}{2} d_{\mathbb{C}}(u, A) - \operatorname{diam}_{\mathbb{C}}(A). \tag{10.5.12}$$

Expand $\mathbb{E}[Z_\varepsilon^2]$ as the sum

$$\mathbb{E}[Z_\varepsilon^2] = \sum_{u,v \in W_\varepsilon} a_{\mathbb{H}}(u) a_{\mathbb{H}}(v) \mathbb{P}\left(u, v \in \operatorname{past}_{\mathfrak{F}}(e)\right).$$

If u and v are both in the past of e in $\mathfrak{F}$, let $w(u,v)$ be the first vertex at which the unique simple paths in $\mathfrak{F}$ from u to e and from v to e meet. Then

$$\mathbb{E}[Z_\varepsilon^2] \leq 2 \sum_{u,v \in W_\varepsilon} a_{\mathbb{H}}(u) a_{\mathbb{H}}(v) \mathbb{P}\left(u, v \in \operatorname{past}_{\mathfrak{F}}(e) \text{ and } w(u,v) \in H(u,v)\right).$$

Consider generating $\mathfrak{F}$ using Wilson's algorithm rooted at infinity, starting first with u and then v. In order for u and v both to be in the past of e and for $w(u,v)$ to be in $H(u,v)$, we must have that the simple random walk from v hits $H(u,v)$, so that

$$\mathbb{E}[Z_\varepsilon^2] \le 2 \sum_{u\in W_\varepsilon} a_{\mathbb{H}}(u)\mathbb{P}(u \in \mathrm{past}_{\mathfrak{F}}(e)) \sum_{v\in W_\varepsilon} a_{\mathbb{H}}(v)\mathbf{P}_v(\tau_{H(u,v)} < \infty). \tag{10.5.13}$$

Thus, it suffices to prove that

$$\sum_{v\in W_\varepsilon} a_{\mathbb{H}}(v)\mathbf{P}_v(\tau_{H(u,v)} < \infty) \preceq \log(1/\varepsilon) \tag{10.5.14}$$

for all $\varepsilon \le \delta_4$ and $u \in W_\varepsilon$.

Recall the definition of U_ε^m from Lemma 10.5.12. Let k be the universal constant from Lemma 10.5.12, let $\ell(\varepsilon) = \lceil \log_{5/4}(\varepsilon) - k\rceil$, and let $\theta = \arg(u)$. For each $m \in \mathbb{Z}$ with $1 \le |m| \le \ell(\varepsilon)$, let $S_\varepsilon^m(\theta)$ be the set of vertices whose centres are contained in $U_\varepsilon^m(\theta)$, and let

$$S_\varepsilon^0 = S_\varepsilon^0(u) := \left\{ v \in W_\varepsilon : \left|\arg\frac{z(v)}{z(u)}\right| \le \frac{5^k\pi\varepsilon}{4^k}\right\}.$$

Then

$$\begin{aligned}\sum_{v\in W_\varepsilon} a_{\mathbb{H}}(v)\mathbf{P}_v(\tau_{H(u,v)} < \infty) &= \sum_{m=-\ell(\varepsilon)}^{\ell(\varepsilon)} \sum_{v\in S_\varepsilon^m} a_{\mathbb{H}}(v)\mathbf{P}_v(\tau_{H(u,v)} < \infty)\\ &\le \sum_{m=-\ell(\varepsilon)}^{\ell(\varepsilon)} \sum_{v\in S_\varepsilon^m} a_{\mathbb{H}}(v)\mathbf{P}_v(\tau_{H(u,S_\varepsilon^m)} < \infty).\end{aligned} \tag{10.5.15}$$

Since S_ε^m is contained in W_ε, it is C-short-lived for some $C = C(\mathbf{M})$ by Corollary 10.5.9. Thus, applying Lemmas 10.5.10–10.5.12 together with (10.5.13) yields that

$$\begin{aligned}\sum_{v\in S_\varepsilon^m} a_{\mathbb{H}}(v)\mathbf{P}_v(\tau_{H(u,S_\varepsilon^m)} < \infty) &\preceq \frac{\mathrm{diam}_{\mathbb{C}}(S_\varepsilon^m)^2}{\left(\frac{1}{2}d_{\mathbb{C}}\left(S_\varepsilon^m, u\right) - \mathrm{diam}_{\mathbb{C}}\left(S_\varepsilon^m\right)\right)^2}\\ &\preceq \frac{\mathrm{diam}_{\mathbb{C}}(U_\varepsilon^m)^2}{\left(\frac{1}{2}d_{\mathbb{C}}\left(U_\varepsilon^m(\theta), \{re^{i\theta} : r \ge 0\}\right) - \mathrm{diam}_{\mathbb{C}}\left(U_\varepsilon^m(\theta)\right)\right)^2} \preceq 1\end{aligned} \tag{10.5.16}$$

for $\varepsilon \le \delta_4$ and $1 \le |m| \le \ell(\varepsilon)$. This in turn yields (10.5.14) when combined with (10.5.15) and the fact that $\sum_{v\in S_\varepsilon^0} a_{\mathbb{H}}(v) \preceq 1$. □

Let δ_4 be the constant from Lemma 10.5.12 and let $\delta_5 = \delta_5(\mathbf{M}, r_{\mathbb{H}}(x))$ be the constant from Lemma 10.5.13. The lower bound of Theorem 10.1.4 now follows from Lemmas 10.5.13 and 10.5.14 together with the Cauchy-Schwartz inequality, which imply that

$$\mathbb{P}(Z_\varepsilon > 0) \ge \frac{\mathbb{E}[Z_\varepsilon]^2}{\mathbb{E}[Z_\varepsilon^2]} \succeq_e \frac{1}{\log(1/\varepsilon)}$$

for all $\varepsilon \le \min\{\delta_4, \delta_5\}$. Since $\mathbb{P}(Z_\varepsilon > 0)$ is an increasing function of ε and $\min\{\delta_4,\delta_5\}$ is an increasing function of $r_{\mathbb{H}}(x)$, it follows that

$$\mathbb{P}(Z_\varepsilon > 0) \ge \mathbb{P}(Z_{\min\{\delta_4,\delta_5\}\varepsilon} > 0) \succeq_e \frac{1}{\log(1/\varepsilon)}$$

for all $\varepsilon \le 1/2$. □

Lemma 10.5.15. $\mathbb{E}[Z_\varepsilon] \preceq 1$ *for all* $\varepsilon > 0$.

Proof. By Wilson's algorithm, $\mathbb{P}(v \in \mathrm{past}_{\mathfrak{F}}(e)) \le \mathbf{P}_v(\tau_x < \infty)$. Thus, the claim follows immediately from Corollary 10.5.9 and Lemma 10.5.10, taking $A = W_\varepsilon \setminus \{x\}$ and $B = \{x\}$, since $\mathscr{C}^{\mathrm{F}}_{\mathrm{eff}}(x \leftrightarrow V \setminus \{x\}) \le c(x) \preceq 1$. The claim can also be proven more directly by a time reversal argument similar to that carried out in the proof of Lemma 10.5.13. □

Proof of Theorem 10.1.5. We continue to use the notation from the proof of Theorem 10.1.4. We first prove the upper bound. Let $R \ge 1$. Applying Markov's inequality and Lemma 10.5.15 we obtain that

$$\mathbb{P}\left(\sum_{k=0}^{R^{1/2}} Z_{s^{-3k}} \ge R\right) \le \frac{1}{R}\mathbb{E}\left[\sum_{k=0}^{R^{1/2}} Z_{s^{-3k}}\right] \preceq R^{-1/2}, \tag{10.5.17}$$

and hence that

$$\mathbb{P}\Big(\mathrm{area}_{\mathbb{H}}(\mathrm{past}_{\mathfrak{F}}(e)) \ge R\Big) \le \mathbb{P}\left(\sum_{k=0}^{R^{1/2}} Z_{s^{-3k}} \ge R\right) + \mathbb{P}\left(\sum_{k=R^{1/2}}^{\infty} Z_{s^{-3k}} > 0\right) \preceq R^{-1/2}, \tag{10.5.18}$$

where the second inequality follows from Theorem 10.1.4 and (10.5.17).

We now obtain a matching lower bound. Let δ_4 be the constant from Lemma 10.5.12 and let $\delta_5 = \delta_5(\mathbf{M}, r_{\mathbb{H}}(x))$ be the constant from Lemma 10.5.13, and let $R_0 = R_0(\mathbf{M}, r_{\mathbb{H}}(x)) = (4/9)\log^2_{1/s}(1/\min\{\delta_4,\delta_5\})$, so that that $s^{-3R^{1/2}/2}$ is less than both δ_4 and δ_5 for all $R \ge R_0$. Let $R \ge R_0$ and let

$$W = \bigcup_{k=\frac{1}{2}R^{1/2}}^{R^{1/2}} W_{s^{-3k}} \qquad \text{and let} \qquad Z = \sum_{k=\frac{1}{2}R^{1/2}}^{R^{1/2}} Z_{s^{-3k}}.$$

The argument used to derive (10.5.13) in the proof of Lemma 10.5.14 also yields that

$$\mathbb{E}[Z^2] \le 2\sum_{u\in W} a_{\mathbb{H}}(u)\mathbb{P}(u \in \mathrm{past}_{\mathfrak{F}}(e)) \sum_{v \in W} a_{\mathbb{H}}(v)\mathbf{P}_v(\tau_{H(u,v)} < \infty). \tag{10.5.19}$$

Let $u \in W$ and let $\theta = \arg(z(u))$. Let $\ell(s^{-3k})$ and the sets $S^m_{s^{-3k}}(\theta)$ be defined as in the proof of

Lemma 10.5.14. Then,

$$\sum_{v \in W} a_{\mathbb{H}}(v) \mathbf{P}_v(\tau_{H(u,v)} < \infty) \leq \sum_{k=\frac{1}{2}R^{1/2}}^{R^{1/2}} \sum_{m=-\ell(s^{-3k})}^{\ell(s^{-3k})} \left[\sum_{v \in S_\varepsilon^m} a_{\mathbb{H}}(v) \mathbf{P}(\tau_{H(u,S_\varepsilon^m)} < \infty) \right]. \qquad (10.5.20)$$

The arguments yielding the estimate (10.5.16) in the proof of Lemma 10.5.14 also yield that the sums appearing in the square brackets on the right-hand side of (10.5.20) are bounded above by a constant depending only on $\mathbf{M}$. It follows that

$$\sum_{v \in W} a_{\mathbb{H}}(v) \mathbf{P}_v(\tau_{H(u,v)} < \infty) \preceq R$$

for all $u \in W$, and hence that $\mathbb{E}[Z^2] \preceq R\mathbb{E}[Z]$.

Next, Lemma 10.5.15 implies that $\mathbb{E}[Z] \succeq_e R^{1/2}$, while Theorem 10.1.4 implies that $\mathbb{P}(Z > 0) \preceq R^{-1/2}$. Thus, there exists a positive constant $C = C(\mathbf{M})$ such that $\mathbb{E}[Z \mid Z > 0] \geq CR$ for all sufficiently large R. Applying the second moment estimate above, the Paley-Zigmund Inequality implies that

$$\mathbb{P}\left(Z \geq \frac{C}{2} R \,\middle|\, Z > 0 \right) \geq \frac{\mathbb{E}[Z \mid Z > 0]^2}{4\mathbb{E}[Z^2 \mid Z > 0]} = \frac{\mathbb{E}[Z]^2}{4\mathbb{E}[Z^2]\mathbb{P}(Z > 0)} \succeq_e 1. \qquad (10.5.21)$$

Combining (10.5.21) with the lower bound of Theorem 10.1.4 yields that

$$\mathbb{P}\left(\mathrm{area}_{\mathbb{H}}(\mathrm{past}_{\mathfrak{F}}(e)) \geq \frac{C}{2} R \right) \geq \mathbb{P}\left(Z \geq \frac{C}{2} R \right) \succeq_e R^{-1/2}$$

for all $R \geq R_0$. Since the probability on the left hand side is decreasing in R, we conclude that

$$\mathbb{P}(\mathrm{area}_{\mathbb{H}}(\mathrm{past}_{\mathfrak{F}}(e)) \geq R) \succeq_e R^{-1/2}$$

as claimed. $\square$

Proof of Theorem 10.1.3. Let $\mathfrak{F}$ be a spanning forest of G and let $e^\dagger$ be the edge of $G^\dagger$ dual to $e = (x, y)$. The past of $e^\dagger$ in the dual forest $\mathfrak{F}^\dagger$ is contained in the region of the plane bounded by e and $\Gamma_{\mathfrak{F}}(x, y)$, so that

$$\mathrm{diam}_{\mathbb{H}}(\Gamma_{\mathfrak{F}}(x, y)) \geq \mathrm{diam}_{\mathbb{H}}(\mathrm{past}_{\mathfrak{F}^\dagger}(e^\dagger)).$$

Meanwhile, if $\mathrm{past}_{\mathfrak{F}^\dagger}(e^\dagger)$ is non-empty, then every edge in the path $\Gamma_{\mathfrak{F}}(x, y)$ is incident to a face of G that is in $\mathrm{past}_{\mathfrak{F}^\dagger}(e^\dagger)$. By the Ring Lemma (Theorem 10.5.1), the hyperbolic radii of circles in $P \cup P^\dagger$ are bounded above. We deduce that there exists a constant $C = C(\mathbf{M})$ such that

$$\mathrm{diam}_{\mathbb{H}}(\Gamma_{\mathfrak{F}}(x, y)) \leq \mathrm{diam}_{\mathbb{H}}(\mathrm{past}_{\mathfrak{F}^\dagger}(e^\dagger)) + C.$$

We deduce Theorem 10.1.3 from Theorem 10.1.4 by applying these estimates when $\mathfrak{F}$ is the FUSF of G and $\mathfrak{F}^\dagger$ is the WUSF of $G^\dagger$. $\square$

10.5.6 The uniformly transient case

Recall that a graph is said to be **uniformly transient** if $\mathbf{p} = \inf_{v\in V} \mathbf{P}_v(\tau_v^+ = \infty)$ is positive. If the graph has bounded degrees, this is equivalent to the property that $\mathscr{C}_{\mathrm{eff}}(v \to \infty)$ is bounded away from zero uniformly in v.

Proposition 10.5.16. *Then there exists a constant $C = C(\mathbf{M})$ such that*

$$r_{\mathbb{H}}(v) \geq \frac{1}{C} \exp\left(-C\mathscr{R}_{\mathrm{eff}}(v \to \infty)\right).$$

Proof. By applying a Möbius transformation if necessary, we may assume that the circle $P(v)$ is centered at the origin. By the Ring Lemma (Theorem 10.5.1), $r_{\mathbb{H}}(v)$ is bounded above by a constant depending only on the maximum degree and codegree of G. Applying [107, Corollary 3.3] together with Theorem 10.5.1 yields that $\mathscr{R}_{\mathrm{eff}}(v \to \infty) \geq c\log(1/r(v))$ for some constant $c = c(\mathbf{M})$. Since the hyperbolic radii are bounded, the Euclidean radius $r(v)$ is comparable to $r_{\mathbb{H}}(v)$. $\square$

We are now ready to prove Corollary 10.1.6. In the rest of this subsection, we will use $\preceq, \succeq$ and $\asymp$ to denote inequalities or equalities that hold up to positive multiplicative constants depending only on $\mathbf{M}$ and $\mathbf{p}$.

Proof of Corollary 10.1.6. Let $\mathfrak{F}$ be a spanning forest of G and let $e^\dagger$ be the edge of $G^\dagger$ dual to $e = (x, y)$. The past of $e^\dagger$ in the dual forest $\mathfrak{F}^\dagger$ is contained in the region of the plane bounded by e and $\Gamma_{\mathfrak{F}}(e)$. The non-amenability of the hyperbolic plane implies that the perimeter of any set is at least a constant multiple of its area, and so

$$\sum_{v\in\Gamma_{\mathfrak{F}}(x,y)} r_{\mathbb{H}}(v) \succeq \mathrm{area}_{\mathbb{H}}\left(\mathrm{past}_{\mathfrak{F}^\dagger}(e^\dagger)\right). \tag{10.5.22}$$

On the other hand, if $\mathrm{past}_{\mathfrak{F}^\dagger}(e^\dagger)$ is non-empty, then every edge in the path $\Gamma_{\mathfrak{F}}(x, y)$ is incident to a face of G that is in $\mathrm{past}_{\mathfrak{F}^\dagger}(e^\dagger)$. We deduce that if $\mathrm{past}_{\mathfrak{F}^\dagger}(e^\dagger)$ is non-empty then

$$\mathrm{area}_{\mathbb{H}}\left(\mathrm{past}_{\mathfrak{F}^\dagger}(e^\dagger)\right) \succeq \sum_{v\in\Gamma_{\mathfrak{F}}(x,y)} a_{\mathbb{H}}(v). \tag{10.5.23}$$

Note that neither estimate (10.5.22) or (10.5.23) required uniform transience. Proposition 10.5.16 and the Ring Lemma (Theorem 10.5.1) imply that

$$|\Gamma_{\mathfrak{F}}(x, y)| \asymp \sum_{v\in\Gamma_{\mathfrak{F}}(x,y)} r_{\mathbb{H}}(v) \asymp \sum_{v\in\Gamma_{\mathfrak{F}}(x,y)} a_{\mathbb{H}}(v).$$

Combining the above estimates in the case that $\mathfrak{F}$ is the FUSF of G and $\mathfrak{F}^\dagger$ is the WUSF of $G^\dagger$ allows us to deduce Corollary 10.1.6 from Theorem 10.1.5. $\square$

Let (X_1, d_1) and (X_2, d_2) be metric spaces and let α, β be positive. A (not necessarily continuous) function $\phi : X_1 \to X_2$ is said to be an (α, β)**-rough isometry** if the following hold.

1. (ϕ **roughly preserves distances.**) $\alpha^{-1} d_1(x, y) - \beta \leq d_2(\phi(x), \phi(y)) \leq \alpha d_1(x, y) + \beta$ for all $x, y \in X_1$.

2. (ϕ **is almost surjective.**) For every $x_2 \in X_2$, there exists $x_1 \in X_1$ such that $d_2(\phi(x_1), x_2) \leq \beta$.

See [173, §2.6] for further background on rough isometries. We write d_G for the graph distance on V.

Corollary 10.5.17. *Let G be a uniformly transient, simple, 3-connected, proper plane network with bounded codegrees and bounded local geometry. Let d_G denote the graph distance on V, let $(P, P^\dagger)$ be a double circle packing of G in $\mathbb{D}$, and let $z(v)$ be the centre of the disc in P corresponding to the vertex v. Then there exist positive constants $\alpha = \alpha(\mathbf{M}, \mathbf{p})$ and $\beta = \beta(\mathbf{M}, \mathbf{p})$ such that z is an (α, β)-rough isometry from (V, d_G) to $(\mathbb{D}, d_{\mathbb{H}})$.*

Proof. Proposition 10.5.16 implies that for every vertex v of G, $r_{\mathbb{H}}(v)$ is bounded both above and away from zero by positive constants. Almost surjectivity is immediate. For each two vertices u and v in G, the shortest graph distance path between them induces a curve in $\mathbb{D}$ (by going along the hyperbolic geodesics between the centres of the circles in the path) whose hyperbolic length is $\succeq d_G(u, v)$.

Conversely, let γ be the hyperbolic geodesic between $z(u)$ and $z(v)$, and consider the set W of vertices w of G such that either $P(w)$ intersects γ or $P^\dagger(f)$ intersects γ for some face f incident to w. Let d be the length of γ. Since all circles in $(P, P^\dagger)$ have a uniform upper bound on their hyperbolic radii, we deduce that all circles in W are contained in a hyperbolic neighbourhood of constant thickness about γ, and hence the total area of these circles is $\preceq d$. Since the radii of the circles are also bounded away from zero, we deduce that the cardinaility of W is also $\preceq d$. Since W contains a path in G from u to v, we deduce that $d_G(u, v)$ is $\preceq d$ as required. □

We summarise the situation for uniformly transient graphs in the following corollary, which follows immediately by combining Corollary 10.5.17 with Theorems 10.1.3–10.1.5 and Corollary 10.1.6. We write diam_G for the graph distance diameter of a set of vertices in G.

Corollary 10.5.18 (Graph distance exponents)**.** *Let G be a uniformly transient, simple, 3-connected, proper plane network with bounded codegrees and bounded local geometry, and let $\mathfrak{F}$ be the free uniform spanning forest of G. Let $\mathbf{p} > 0$ be a uniform lower bound on the escape probabilities of G. Then there exist positive constants $k_1 = k_1(\mathbf{M}, \mathbf{p})$ and $k_2 = k_2(\mathbf{M}, \mathbf{p})$ such that*

$$\begin{array}{rcccl} k_1 R^{-1} & \leq & \mathsf{FUSF}_G(\operatorname{diam}_G(\Gamma_{\mathfrak{F}}(x,y)) \geq R) & \leq & k_2 R^{-1}, \\ k_1 R^{-1} & \leq & \mathsf{WUSF}_G(\operatorname{diam}_G(\operatorname{past}_{\mathfrak{F}}(e)) \geq R) & \leq & k_2 R^{-1}, \\ k_1 R^{-1/2} & \leq & \mathsf{WUSF}_G(|\operatorname{past}_{\mathfrak{F}}(e)| \geq R) & \leq & k_2 R^{-1/2}, \text{ and} \\ k_1 R^{-1/2} & \leq & \mathsf{FUSF}_G(|\Gamma_{\mathfrak{F}}(x,y)| \geq R) & \leq & k_2 R^{-1/2} \end{array}$$

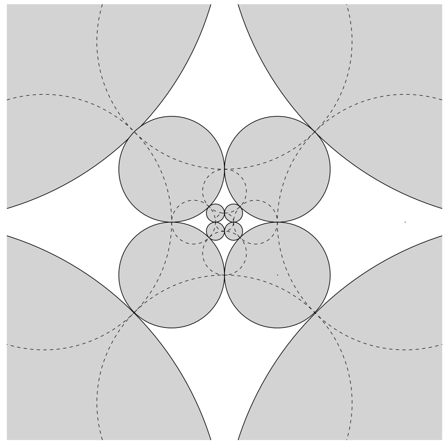

Figure 10.6: The double circle packing of $\mathbb{N} \times \mathbb{Z}_4$.

for every edge $e = (x, y)$ of G and every $R \geq 1$.

10.6 Closing remarks and open problems

10.6.1 Remarks

Remark 10.6.1 (Non-universality in the parabolic case.). Unlike in the CP hyperbolic case, the exponents governing the behaviour of the USTs of CP parabolic, simple, 3-connected proper plane graphs with bounded codegrees and bounded local geometry are not universal, and need not exist in general.

Indeed, consider the double circle packing of the proper plane quadrangulation with underlying graph $\mathbb{N}\times\mathbb{Z}_4$, pictured in Figure 10.6. Let the packing be normalised to be symmetric under rotation by $\pi/2$ about the origin and to have $r(0,0) = 1$. It is possible to compute that $r(i,j) = (3+2\sqrt{2})^i$ and hence that $|z(i,j)|$ is comparable to $(3+2\sqrt{2})^i$ for every $(i,j) \in \mathbb{N} \times \mathbb{Z}_4$. Suppose that the edges connecting (i,j) to $(i \pm 1, j)$ are given weight 1 for every $(i,j) \in \mathbb{N} \times \mathbb{Z}_4$, while the edges connecting (i,j) to $(i, j \pm 1)$ are given weight c for each $(i,j) \in \mathbb{N} \times \mathbb{Z}_4$. It can be computed that the probability that a walk started at $(i,0)$ hits $(0,0)$ without ever changing its second coordinate is $a(c)^i := (1 + c - \sqrt{c^2 + 2c})^i$. Let $e = ((0,0),(0,1))$. By running Wilson's algorithm rooted at $(0,0)$ starting from the vertices $(i,0)$ and $(i,1)$, we see that

$$\begin{aligned}\mathsf{UST}\left(\mathrm{past}_T(e) \cap \{i\} \times \mathbb{Z}_4 \neq \emptyset\right) &\geq \mathbf{P}_{(i,0)}(\tau_{(0,0)} < \tau_{\mathbb{N}\times\{1,2,3\}})\mathbf{P}_{(i,1)}(\tau_{(0,1)} < \tau_{\mathbb{N}\times\{0,2,3\}}) \\ &\qquad\qquad \cdot \mathbf{P}_{(0,1)}(X_1 = (0,0)) \\ &= \frac{c}{2c+1} a(c)^{2i}.\end{aligned}$$

the right-hand side is exactly the probability that the random walk from $(i,0)$ hits $(0,0)$ without ever changing its second coordinate, and that the random walk from $(i,1)$ hits $(0,1)$ without ever changing its second coordinate and then steps to $(0,0)$.

Let $q(c)$ be the probability that a random walk started at (i,j) visits every vertex of $\{i\} \times \mathbb{Z}_4$ before changing its vertical coordinate, which tends to one as $c \to \infty$. Let Y be a loop-erased random walk from $(0,0)$. It can be computed that the probability that a random walk started from (i,j) visits $\{0\} \times \mathbb{Z}_4$ before hitting the trace of Y is at most $b(c)^i := (1-q(c))^i$. Thus, by Wilson's algorithm and a union bound,

$$\mathsf{UST}(\mathrm{past}_T(e) \cap \{i\} \times \mathbb{Z}_4 \neq \emptyset) \leq 4b(c)^i.$$

It follows that there exist positive constants $k(c)$, $\alpha(c)$ and $\beta(c)$ such that $\alpha(c) \to 0$ as $c \to 0$, $\beta(c) \to \infty$ as $c \to \infty$, and

$$k(c)^{-1} R^{-\alpha(c)} \leq \mathsf{UST}\left(\mathrm{diam}_{\mathbb{C}}(\mathrm{past}_T(e)) \geq R\right) \leq k(c) R^{-\beta(c)}.$$

Thus, by varying c, we obtain CP parabolic proper plane networks with bounded codegrees and bounded local geometry with different exponents governing the diameter of the pasts of edges in their USTs: If c is large the diameter has a light tail, while if c is small the diameter has a heavy tail. Furthermore, by varying the weight of $((i,j),(i,j\pm 1))$ as a function of i in the above example (i.e., making c small at some scales and large at others), it is possible to construct a simple, 3-connected, CP parabolic proper plane network G with bounded codegrees and bounded local geometry such that

$$\frac{\log \mathsf{UST}_G\left(\mathrm{diam}_{\mathbb{C}}(\mathrm{past}_T(e)) \geq R\right)}{\log(R)}$$

does not converge as $R \to \infty$ for some edge e of G. The details are left to the reader.

Similar constructions show that the behaviour of $\mathsf{WUSF}_G(\mathrm{diam}_{\mathbb{H}}(\mathrm{past}_{\mathfrak{F}}(e)) \geq R \cdot r_{\mathbb{H}}(x))$ is not universal over simple, 3-connected, CP hyperbolic proper plane network G with bounded codegrees and bounded local geometry in the regime that $r_{\mathbb{H}}(x)$ is small.

10.6.2 Open problems

It is natural to ask to what extent the assumption of planarity in Theorem 10.1.1 can be relaxed. Part (1) of the following question was suggested by R. Lyons.

Question 10.6.2. *Let G be a bounded degree proper plane graph.*

1. *Let H be a finite graph. Is the free uniform spanning forest of the product graph $G \times H$ connected almost surely?*

2. *Let G' be a bounded degree graph that is rough isometric to G. Is the the free uniform spanning forest of G connected almost surely?*

Without the assumption of planarity, connectivity of the FUSF is not preserved by rough isometries; this can be seen from an analysis of the graphs appearing in [47, Theorem 3.5].

Bibliography

[1] Memories of Oded Schramm. available at `http://odedschramm.wordpress.com/`. Retrieved August 2017.

[2] Dov Aharonov. The sharp constant in the ring lemma. *Complex Variables Theory Appl.*, 33(1-4):27–31, 1997.

[3] M. Aizenman, H. Kesten, and C. M. Newman. Uniqueness of the infinite cluster and continuity of connectivity functions for short and long range percolation. *Comm. Math. Phys.*, 111(4):505–531, 1987.

[4] Michael Aizenman and David J Barsky. Sharpness of the phase transition in percolation models. *Communications in Mathematical Physics*, 108(3):489–526, 1987.

[5] Michael Aizenman and Simone Warzel. The canopy graph and level statistics for random operators on trees. *Math. Phys. Anal. Geom.*, 9(4):291–333 (2007), 2006.

[6] David Aldous and James Allen Fill. Reversible markov chains and random walks on graphs, 2002. Unfinished monograph, recompiled 2014.

[7] David Aldous and Russell Lyons. Processes on unimodular random networks. *Electron. J. Probab.*, 12:no. 54, 1454–1508, 2007.

[8] David J. Aldous. The random walk construction of uniform spanning trees and uniform labelled trees. *SIAM J. Discrete Math.*, 3(4):450–465, 1990.

[9] A. D. Aleksandrov and V. A. Zalgaller. *Intrinsic geometry of surfaces.* Translated from the Russian by J. M. Danskin. Translations of Mathematical Monographs, Vol. 15. American Mathematical Society, Providence, R.I., 1967.

[10] V. Anantharam and P. Tsoucas. A proof of the Markov chain tree theorem. *Statist. Probab. Lett.*, 8(2):189–192, 1989.

[11] Alano Ancona. Positive harmonic functions and hyperbolicity. In *Potential theory—surveys and problems (Prague, 1987)*, volume 1344 of *Lecture Notes in Math.*, pages 1–23. Springer, Berlin, 1988.

[12] Alano Ancona, Russell Lyons, and Yuval Peres. Crossing estimates and convergence of Dirichlet functions along random walk and diffusion paths. *Ann. Probab.*, 27(2):970–989, 1999.

[13] James Anderson. *Hyperbolic geometry.* Springer Science & Business Media, 2006.

[14] E. M. Andreev. Convex polyhedra in Lobačevskiĭ spaces. *Mat. Sb. (N.S.)*, 81 (123):445–478, 1970.

[15] O. Angel. Growth and percolation on the uniform infinite planar triangulation. *Geom. Funct. Anal.*, 13(5):935–974, 2003.

[16] O. Angel, A. Nachmias, and G. Ray. Random walks on stochastic hyperbolic half planar triangulations. *ArXiv e-prints*, August 2014.

[17] Omer Angel, Martin T. Barlow, Ori Gurel-Gurevich, and Asaf Nachmias. Boundaries of planar graphs, via circle packings. *Ann. Probab.*, 44(3):1956–1984, 2016.

[18] Omer Angel, Guillaume Chapuy, Nicolas Curien, and Gourab Ray. The local limit of unicellular maps in high genus. *Elec. Comm. Prob*, 18:1–8, 2013.

[19] Omer Angel and Nicolas Curien. Percolations on random maps I: half-plane models. *Ann. Inst. H. Poincaré*, 2013. To appear.

[20] Omer Angel, Tom Hutchcroft, Asaf Nachmias, and Gourab Ray. Hyperbolic and parabolic unimodular random maps. *arXiv preprint arXiv:1612.08693*, 2016.

[21] Omer Angel, Tom Hutchcroft, Asaf Nachmias, and Gourab Ray. Unimodular hyperbolic triangulations: circle packing and random walk. *Inventiones Mathematicae*, 206(1):229–268, 2016.

[22] Omer Angel and Gourab Ray. The half plane uipt is recurrent. *Probability Theory and Related Fields.* to appear.

[23] Omer Angel and Gourab Ray. Classification of half-planar maps. *Ann. Probab.*, 43(3):1315–1349, 2015.

[24] Omer Angel and Oded Schramm. Uniform infinite planar triangulations. *Comm. Math. Phys.*, 241(2-3):191–213, 2003.

[25] Tonći Antunović and Ivan Veselić. Sharpness of the phase transition and exponential decay of the subcritical cluster size for percolation on quasi-transitive graphs. *Journal of Statistical Physics*, 130(5):983–1009, 2008.

[26] André Avez. Entropie des groupes de type fini. *C. R. Acad. Sci. Paris Sér. A-B*, 275:A1363–A1366, 1972.

[27] André Avez. Théorème de Choquet-Deny pour les groupes à croissance non exponentielle. *C. R. Acad. Sci. Paris Sér. A*, 279:25–28, 1974.

[28] M. T. Barlow, D. A. Croydon, and T. Kumagai. Subsequential scaling limits of simple random walk on the two-dimensional uniform spanning tree. *Ann. Probab.*, 45(1):4–55, 2017.

[29] Martin T. Barlow. Loop erased walks and uniform spanning trees. 34:1–32, 2016.

[30] Martin T Barlow and Antal A Járai. Geometry of the uniform spanning forest components in high dimensions. *arXiv preprint arXiv:1602.01505*, 2016.

[31] Martin T Barlow and Robert Masson. Spectral dimension and random walks on the two dimensional uniform spanning tree. *Communications in Mathematical Physics*, 305(1):23–57, 2011.

[32] Martin T. Barlow, Yuval Peres, and Perla Sousi. Collisions of random walks. *Ann. Inst. Henri Poincaré Probab. Stat.*, 48(4):922–946, 2012.

[33] Laurent Bartholdi and Bálint Virág. Amenability via random walks. *Duke Mathematical Journal*, 130(1):39–56, 2005.

[34] Alan F. Beardon and Kenneth Stephenson. Circle packings in different geometries. *Tohoku Math. J. (2)*, 43(1):27–36, 1991.

[35] Alan F Beardon and Kenneth Stephenson. Circle packings in different geometries. *Tohoku Mathematical Journal, Second Series*, 43(1):27–36, 1991.

[36] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Group-invariant percolation on graphs. *Geom. Funct. Anal.*, 9(1):29–66, 1999.

[37] I. Benjamini, E. Paquette, and J. Pfeffer. Anchored expansion, speed, and the hyperbolic Poisson Voronoi tessellation. *ArXiv e-prints*, September 2014.

[38] Itai Benjamini and Nicolas Curien. Ergodic theory on stationary random graphs. *Electron. J. Probab.*, 17:no. 93, 20, 2012.

[39] Itai Benjamini and Nicolas Curien. Simple random walk on the uniform infinite planar quadrangulation: subdiffusivity via pioneer points. *Geom. Funct. Anal.*, 23(2):501–531, 2013.

[40] Itai Benjamini, Nicolas Curien, and Agelos Georgakopoulos. The Liouville and the intersection properties are equivalent for planar graphs. *Electron. Commun. Probab.*, 17:no. 42, 5, 2012.

[41] Itai Benjamini, Harry Kesten, Yuval Peres, and Oded Schramm. Geometry of the uniform spanning forest: transitions in dimensions $4, 8, 12, \ldots$. *Ann. of Math. (2)*, 160(2):465–491, 2004.

[42] Itai Benjamini and Gady Kozma. A resistance bound via an isoperimetric inequality. *Combinatorica*, 25(6):645–650, 2005.

[43] Itai Benjamini, Russell Lyons, Yuval Peres, and Oded Schramm. Critical percolation on any nonamenable group has no infinite clusters. *Ann. Probab.*, 27(3):1347–1356, 1999.

[44] Itai Benjamini, Russell Lyons, Yuval Peres, and Oded Schramm. Uniform spanning forests. *Ann. Probab.*, 29(1):1–65, 2001.

[45] Itai Benjamini, Russell Lyons, and Oded Schramm. Percolation perturbations in potential theory and random walks. In *Random walks and discrete potential theory (Cortona, 1997)*, Sympos. Math., XXXIX, pages 56–84. Cambridge Univ. Press, Cambridge, 1999.

[46] Itai Benjamini, Russell Lyons, and Oded Schramm. Unimodular random trees. 2015.

[47] Itai Benjamini and Oded Schramm. Harmonic functions on planar and almost planar graphs and manifolds, via circle packings. *Invent. Math.*, 126(3):565–587, 1996.

[48] Itai Benjamini and Oded Schramm. Random walks and harmonic functions on infinite planar graphs using square tilings. *Ann. Probab.*, 24(3):1219–1238, 1996.

[49] Itai Benjamini and Oded Schramm. Percolation in the hyperbolic plane, 2001.

[50] Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6:no. 23, 13 pp. (electronic), 2001.

[51] Itai Benjamini and Oded Schramm. Percolation beyond $\mathbb{Z}^d$, many questions and a few answers [mr1423907]. In *Selected works of Oded Schramm. Volume 1, 2*, Sel. Works Probab. Stat., pages 679–690. Springer, New York, 2011.

[52] Sandeep Bhupatiraju, Jack Hanson, and Antal A Járai. Inequalities for critical exponents in d-dimensional sandpiles. *arXiv preprint arXiv:1602.06475*, 2016.

[53] N. L. Biggs, Bojan Mohar, and John Shawe-Taylor. The spectral radius of infinite graphs. *Bull. London Math. Soc.*, 20(2):116–120, 1988.

[54] Ian Biringer and Jean Raimbault. The topology of invariant random surfaces. http://arxiv.org/abs/1411.0561.

[55] David Blackwell. On transient markov processes with a countable number of states and stationary transition probabilities. *The Annals of Mathematical Statistics*, pages 654–658, 1955.

[56] Lewis Bowen. Periodicity and circle packings of the hyperbolic plane. *Geom. Dedicata*, 102:213–236, 2003.

[57] Lewis Bowen. Couplings of uniform spanning forests. *Proceedings of the American Mathematical Society*, pages 2151–2158, 2004.

[58] Graham R. Brightwell and Edward R. Scheinerman. Representations of planar graphs. *SIAM J. Discrete Math.*, 6(2):214–229, 1993.

[59] Andrei Broder. Generating random spanning trees. In *Foundations of Computer Science, 1989., 30th Annual Symposium on*, pages 442–447. IEEE, 1989.

[60] R. L. Brooks, C. A. B. Smith, A. H. Stone, and W. T. Tutte. The dissection of rectangles into squares. *Duke Math. J.*, 7:312–340, 1940.

[61] LEJ Brouwer. On the structure of perfect sets of points. In *KNAW, Proceedings*, volume 12, pages 1909–1910.

[62] Robert Burton and Robin Pemantle. Local characteristics, entropy and limit theorems for spanning trees and domino tilings via transfer-impedances. *Ann. Probab.*, 21(3):1329–1371, 1993.

[63] Robert Burton and Robin Pemantle. Local characteristics, entropy and limit theorems for spanning trees and domino tilings via transfer-impedances. *Ann. Probab.*, 21(3):1329–1371, 1993.

[64] Robert M Burton and Michael Keane. Density and uniqueness in percolation. *Communications in Mathematical Physics*, 121(3):501–505, 1989.

[65] Peter Buser. A note on the isoperimetric constant. *Ann. Sci. École Norm. Sup. (4)*, 15(2):213–230, 1982.

[66] Thomas Keith Carne. A transmutation formula for Markov chains. *Bull. Sci. Math. (2)*, 109(4):399–405, 1985.

[67] Jiří Černý and Augusto Quadros Teixeira. *From random walk trajectories to random interlacements*, volume 23 of *Ensaios Matemáticos [Mathematical Surveys]*. Sociedade Brasileira de Matemática, Rio de Janeiro, 2012.

[68] Jeff Cheeger. A lower bound for the smallest eigenvalue of the Laplacian. pages 195–199, 1970.

[69] Dmitry Chelkak. Robust discrete complex analysis: a toolbox. *Ann. Probab.*, 44(1):628–683, 2016.

[70] Dayue Chen, Yuval Peres, and Gabor Pete. Anchored expansion, percolation and speed. *Annals of probability*, pages 2978–2995, 2004.

[71] XinXing Chen and DaYue Chen. Two random walks on the open cluster of $\mathbb{Z}^2$ meet infinitely often. *Sci. China Math.*, 53(8):1971–1978, 2010.

[72] Ching Chou. Elementary amenable groups. *Illinois Journal of Mathematics*, 24(3):396–407, 1980.

[73] N. Curien, L. Ménard, and G. Miermont. A view from infinity of the uniform infinite planar quadrangulation. *ALEA Lat. Am. J. Probab. Math. Stat.*, 10(1):45–88, 2013.

[74] Nicolas Curien. A glimpse of the conformal structure of random planar maps. *Comm. Math. Phys.*, 333(3):1417–1463, 2015.

[75] Nicolas Curien. Planar stochastic hyperbolic triangulations. *Probab. Theory Related Fields*, 165(3-4):509–540, 2016.

[76] Nicolas Curien and Grégory Miermont. Uniform infinite planar quadrangulations with a boundary. *Random Structures Algorithms*, 47(1):30–58, 2015.

[77] Yves Derriennic. Quelques applications du théorème ergodique sous-additif. In *Conference on Random Walks (Kleebach, 1979) (French)*, volume 74 of *Astérisque*, pages 183–201, 4. Soc. Math. France, Paris, 1980.

[78] Deepak Dhar. Self-organized critical state of sandpile automaton models. *Phys. Rev. Lett.*, 64(14):1613–1616, 1990.

[79] Jian Ding, James R Lee, and Yuval Peres. Markov type and threshold embeddings. *Geometric and Functional Analysis*, 23(4):1207–1229, 2013.

[80] J. Dodziuk and W. S. Kendall. Combinatorial Laplacians and isoperimetric inequality. In *From local times to global geometry, control and physics (Coventry, 1984/85)*, volume 150 of *Pitman Res. Notes Math. Ser.*, pages 68–74. Longman Sci. Tech., Harlow, 1986.

[81] Jozef Dodziuk. Difference equations, isoperimetric inequality and transience of certain random walks. *Trans. Amer. Math. Soc.*, 284(2):787–794, 1984.

[82] Alexander Drewitz, Balázs Ráth, and Artëm Sapozhnikov. *An introduction to random interlacements*. Springer Briefs in Mathematics. Springer, Cham, 2014.

[83] Hugo Duminil-Copin and Vincent Tassion. A new proof of the sharpness of the phase transition for bernoulli percolation and the ising model. *Communications in Mathematical Physics*, pages 1–21, 2015.

[84] E. B. Dynkin. The boundary theory of Markov processes (discrete case). *Uspehi Mat. Nauk*, 24(2 (146)):3–42, 1969.

[85] Gábor Elek. On the limit of large girth graph sequences. *Combinatorica*, 30(5):553–563, 2010.

[86] Gábor Elek and Gábor Lippner. Sofic equivalence relations. *J. Funct. Anal.*, 258(5):1692–1708, 2010.

[87] Inessa Epstein and Nicolas Monod. Nonunitarizable representations and random forests. *Int. Math. Res. Not. IMRN*, (22):4336–4353, 2009.

[88] P. Erdös and S. J. Taylor. Some intersection properties of random walk paths. *Acta Math. Acad. Sci. Hungar.*, 11:231–248, 1960.

[89] Robert Fitzner and Remco van der Hofstad. Nearest-neighbor percolation function is continuous for $d > 10$. *arXiv preprint arXiv:1506.07977*, 2015.

[90] P. J. Fitzsimmons and Thomas S. Salisbury. Capacity and energy for multiparameter Markov processes. *Ann. Inst. H. Poincaré Probab. Statist.*, 25(3):325–350, 1989.

[91] Harry Furstenberg. A Poisson formula for semi-simple Lie groups. *Ann. of Math. (2)*, 77:335–386, 1963.

[92] Harry Furstenberg. Boundaries of Lie groups and discrete subgroups. In *Actes du Congrès International des Mathématiciens (Nice, 1970), Tome 2*, pages 301–306. Gauthier-Villars, Paris, 1971.

[93] Harry Furstenberg. Random walks and discrete subgroups of Lie groups. In *Advances in Probability and Related Topics, Vol. 1*, pages 1–63. Dekker, New York, 1971.

[94] D. Gaboriau. Invariant percolation and harmonic Dirichlet functions. *Geom. Funct. Anal.*, 15(5):1004–1051, 2005.

[95] Damien Gaboriau. What is ... cost? *Notices Amer. Math. Soc.*, 57(10):1295–1296, 2010.

[96] Damien Gaboriau and Russell Lyons. A measurable-group-theoretic solution to von Neumann's problem. *Invent. Math.*, 177(3):533–540, 2009.

[97] A Gandolfi, MS Keane, and CM Newman. Uniqueness of the infinite component in a random graph with applications to percolation and spin glasses. *Probability Theory and Related Fields*, 92(4):511–527, 1992.

[98] Christophe Garban. Quantum gravity and the KPZ formula [after Duplantier-Sheffield]. *Astérisque*, (352):Exp. No. 1052, ix, 315–354, 2013. Séminaire Bourbaki. Vol. 2011/2012. Exposés 1043–1058.

[99] Christophe Garban, Gábor Pete, and Oded Schramm. The fourier spectrum of critical percolation. *Acta Mathematica*, 205(1):19–104, 2010.

[100] Agelos Georgakopoulos. The boundary of a square tiling of a graph coincides with the Poisson boundary. *Invent. Math.*, 203(3):773–821, 2016.

[101] Peter Gerl. Random walks on graphs with a strong isoperimetric property. *J. Theoret. Probab.*, 1(2):171–187, 1988.

[102] John R Gilbert, Joan P Hutchinson, and Robert Endre Tarjan. A separator theorem for graphs of bounded genus. *Journal of Algorithms*, 5(3):391–407, 1984.

[103] James T. Gill and Steffen Rohde. On the Riemann surface type of random planar maps. *Rev. Mat. Iberoam.*, 29(3):1071–1090, 2013.

[104] G. R. Grimmett, H. Kesten, and Y. Zhang. Random walk on the infinite cluster of the percolation model. *Probab. Theory Related Fields*, 96(1):33–44, 1993.

[105] Geoffrey Grimmett. *The random-cluster model*, volume 333 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2006.

[106] Geoffrey R Grimmett. Percolation (grundlehren der mathematischen wissenschaften). 2010.

[107] Ori Gurel-Gurevich and Asaf Nachmias. Recurrence of planar graph limits. *Ann. of Math. (2)*, 177(2):761–781, 2013.

[108] Ori Gurel-Gurevich, Asaf Nachmias, and Juan Souto. Recurrence of multiply-ended planar triangulations. http://arxiv.org/abs/1506.00221.

[109] Olle Häggström. Random-cluster measures and uniform spanning trees. *Stochastic Process. Appl.*, 59(2):267–275, 1995.

[110] Olle Häggström. Infinite clusters in dependent automorphism invariant percolation on trees. *Ann. Probab.*, 25(3):1423–1436, 1997.

[111] Olle Häggström. Uniform and minimal essential spanning forests on trees. *Random Structures & Algorithms*, 12(1):27–50, 1998.

[112] Olle Häggström and Johan Jonasson. Uniqueness and non-uniqueness in percolation theory. *Probab. Surv.*, 3:289–344, 2006.

[113] Olle Häggström and Yuval Peres. Monotonicity of uniqueness for percolation on Cayley graphs: all infinite clusters are born simultaneously. *Probab. Theory Related Fields*, 113(2):273–285, 1999.

[114] Olle Häggström, Yuval Peres, and Roberto H. Schonmann. Percolation on transitive graphs as a coalescent process: relentless merging followed by simultaneous uniqueness. In *Perplexing problems in probability*, volume 44 of *Progr. Probab.*, pages 69–90. Birkhäuser Boston, Boston, MA, 1999.

[115] Alan Hammond, Gábor Pete, and Oded Schramm. Local time on the exceptional set of dynamical percolation and the incipient infinite cluster. *The Annals of Probability*, 43(6):2949–3005, 2015.

[116] Lowell J. Hansen. On the Rodin and Sullivan ring lemma. *Complex Variables Theory Appl.*, 10(1):23–30, 1988.

[117] Takashi Hara and Gordon Slade. Mean-field behaviour and the lace expansion. In *Probability and phase transition*, pages 87–122. Springer, 1994.

[118] Theodore E Harris. A lower bound for the critical probability in a certain percolation process. In *Proc. Cambridge Philos. Soc*, volume 56, page 3. Cambridge Univ Press, 1960.

[119] Allen Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002.

[120] Zheng-Xu He. Rigidity of infinite disk patterns. *Annals of Mathematics*, 149:1–33, 1999.

[121] Zheng-Xu He and O. Schramm. Hyperbolic and parabolic packings. *Discrete Comput. Geom.*, 14(2):123–149, 1995.

[122] Zheng-Xu He and Oded Schramm. Fixed points, Koebe uniformization and circle packings. *Ann. of Math. (2)*, 137(2):369–406, 1993.

[123] Zheng-Xu He and Oded Schramm. The c-convergence of hexagonal disk packings to the riemann map. *Acta Mathematica*, 180(2):219–245, 1998.

[124] W. Hebisch and L. Saloff-Coste. Gaussian estimates for Markov chains and random walks on groups. *Ann. Probab.*, 21(2):673–709, 1993.

[125] Christopher Hoffman. Recurrence of simple random walk on $\mathbb{Z}^2$ is dynamically sensitive. *ALEA Lat. Am. J. Probab. Math. Stat.*, 1:35–45, 2006.

[126] Alexander E. Holroyd, Lionel Levine, Karola Mészáros, Yuval Peres, James Propp, and David B. Wilson. Chip-firing and rotor-routing on directed graphs. In *In and out of equilibrium. 2*, volume 60 of *Progr. Probab.*, pages 331–364. Birkhäuser, Basel, 2008.

[127] Tom Hutchcroft. Wired cycle-breaking dynamics for uniform spanning forests. *Annals of Probability.* to appear.

[128] Tom Hutchcroft. Interlacements and the wired uniform spanning forest. 2015.

[129] Tom Hutchcroft. Critical percolation on any quasi-transitive graph of exponential growth has no infinite clusters. *Comptes Rendus Mathematique*, 354(9):944 – 947, 2016.

[130] Tom Hutchcroft and Asaf Nachmias. Uniform spanning forests of planar graphs. http://arxiv.org/abs/1603.07320.

[131] Tom Hutchcroft and Asaf Nachmias. Indistinguishability of trees in uniform spanning forests. *Probability Theory and Related Fields*, pages 1–40, 2016.

[132] Tom Hutchcroft and Yuval Peres. Boundaries of planar graphs: A unified approach. *arXiv preprint arXiv:1508.03923*, 2015.

[133] Tom Hutchcroft and Yuval Peres. Collisions of random walks in reversible random graphs. *Electron. Commun. Probab.*, 20:6 pp., 2015.

[134] Tom Hutchcroft and Yuval Peres. The component graph of the uniform spanning forest: Transitions in dimensions $9, 10, 11, \ldots$. *arXiv preprint arXiv:1702.05780*, 2017.

[135] Antal A. Járai. Sandpile models. Preprint (2014).

[136] Antal A. Járai. Incipient infinite percolation clusters in 2D. *Ann. Probab.*, 31(1):444–485, 2003.

[137] Antal A. Járai and Frank Redig. Infinite volume limit of the abelian sandpile model in dimensions $d \geq 3$. *Probab. Theory Related Fields*, 141(1-2):181–212, 2008.

[138] Antal A. Járai and Nicolás Werning. Minimal configurations and sandpile measures. *J. Theoret. Probab.*, 27(1):153–167, 2014.

[139] Johan Jonasson and Oded Schramm. On the cover time of planar graphs. *Electron. Comm. Probab.*, 5:85–90 (electronic), 2000.

[140] Kate Juschenko and Nicolas Monod. Cantor systems, piecewise translations and simple amenable groups. *Ann. of Math. (2)*, 178(2):775–787, 2013.

[141] V Kaimanovich. Boundary and entropy of random walks in random environment. *Prob. Theory and Math. Stat*, 1:573–579, 1990.

[142] Vadim A. Kaimanovich. Measure-theoretic boundaries of Markov chains, 0-2 laws and entropy. In *Proceedings of the Conference on Harmonic Analysis and Discrete Potential Theory (Frascati*, pages 145–180. Plenum, 1991.

[143] Vadim A. Kaimanovich. Dirichlet norms, capacities and generalized isoperimetric inequalities for Markov operators. *Potential Anal.*, 1(1):61–82, 1992.

[144] Vadim A Kaimanovich. Hausdorff dimension of the harmonic measure on trees. *Ergodic Theory and Dynamical Systems*, 18(03):631–660, 1998.

[145] Vadim A Kaimanovich. Random walks on sierpiński graphs: hyperbolicity and stochastic homogenization. In *Fractals in Graz 2001*, pages 145–183. Springer, 2003.

[146] Vadim A Kaimanovich, Yuri Kifer, and B-Z Rubshtein. Boundaries and harmonic functions for random walks with random transition probabilities. *Journal of Theoretical Probability*, 17(3):605–646, 2004.

[147] Vadim A Kaimanovich and Anatoly M Vershik. Random walks on discrete groups: boundary and entropy. *The annals of probability*, pages 457–490, 1983.

[148] Vadim A Kaimanovich and Wolfgang Woess. Boundary and entropy of space homogeneous markov chains. *Annals of probability*, pages 323–363, 2002.

[149] V. A. Kaĭ manovich and A. M. Vershik. Random walks on discrete groups: boundary and entropy. *Ann. Probab.*, 11(3):457–490, 1983.

[150] Richard Kenyon. The asymptotic determinant of the discrete Laplacian. *Acta Math.*, 185(2):239–286, 2000.

[151] Harry Kesten. Full Banach mean values on countable groups. *Math. Scand.*, 7:146–156, 1959.

[152] Harry Kesten. Symmetric random walks on groups. *Trans. Amer. Math. Soc.*, 92:336–354, 1959.

[153] Harry Kesten. The incipient infinite cluster in two-dimensional percolation. *Probab. Theory Related Fields*, 73(3):369–394, 1986.

[154] G. Kirchhoff. Ueber die auflsung der gleichungen, auf welche man bei der untersuchung der linearen vertheilung galvanischer strme gefhrt wird. *Annalen der Physik*, 148(12):497–508, 1847.

[155] Achim Klenke. *Probability theory*. Universitext. Springer-Verlag London, Ltd., London, 2008. A comprehensive course, Translated from the 2006 German original.

[156] Paul Koebe. *Kontaktprobleme der konformen Abbildung*. Hirzel, 1936.

[157] Maxim Krikun. Local structure of random quadrangulations. arXiv:math/0512304, 2005.

[158] Manjunath Krishnapur and Yuval Peres. Recurrent graphs where two independent random walks collide finitely often. *Electron. Comm. Probab.*, 9:72–81 (electronic), 2004.

[159] Takashi Kumagai. *Random walks on disordered media and their scaling limits*, volume 2101 of *Lecture Notes in Mathematics*. Springer, Cham, 2014. Lecture notes from the 40th Probability Summer School held in Saint-Flour, 2010, École d'Été de Probabilités de Saint-Flour. [Saint-Flour Probability Summer School].

[160] Sergei K. Lando and Alexander K. Zvonkin. *Graphs on surfaces and their applications*, volume 141 of *Encyclopaedia of Mathematical Sciences*. Springer-Verlag, Berlin, 2004. With an appendix by Don B. Zagier, Low-Dimensional Topology, II.

[161] Gregory F. Lawler. A self-avoiding random walk. *Duke Math. J.*, 47(3):655–693, 1980.

[162] Gregory F Lawler. *Conformally invariant processes in the plane*. Number 114. American Mathematical Soc., 2008.

[163] Gregory F. Lawler, Oded Schramm, and Wendelin Werner. Conformal invariance of planar loop-erased random walks and uniform spanning trees. *Ann. Probab.*, 32(1B):939–995, 2004.

[164] James R Lee. Conformal growth rates and spectral geometry on distributional limits of graphs. *arXiv preprint arXiv:1701.01598*, 2017.

[165] F.T. Leighton and R.L. Rivest. The markov chain tree theorem. 1983. M.I.T. Laboratory for Computer Science, Technical Report, MIT/LCL/TM-249.

[166] Thomas M. Liggett. *Interacting particle systems*. Classics in Mathematics. Springer-Verlag, Berlin, 2005. Reprint of the 1985 original.

[167] Richard J. Lipton and Robert Endre Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.*, 9(3):615–627, 1980.

[168] Russell Lyons. Random walks and the growth of groups. *Comptes rendus de l'Académie des sciences. Série 1, Mathématique*, 320(11):1361–1366, 1995.

[169] Russell Lyons. Random complexes and l^2-Betti numbers. *J. Topol. Anal.*, 1(2):153–175, 2009.

[170] Russell Lyons, Benjamin J. Morris, and Oded Schramm. Ends in uniform spanning forests. *Electron. J. Probab.*, 13:no. 58, 1702–1725, 2008.

[171] Russell Lyons, Robin Pemantle, and Yuval Peres. Ergodic theory on Galton-Watson trees: speed of random walk and dimension of harmonic measure. *Ergodic Theory Dynam. Systems*, 15(3):593–619, 1995.

[172] Russell Lyons and Yuval Peres. Poisson boundaries of lamplighter groups: proof of the kaimanovich-vershik conjecture. *arXiv preprint arXiv:1508.01845*, 2015.

[173] Russell Lyons and Yuval Peres. *Probability on Trees and Networks*. Cambridge University Press, New York, 2016. Available at `http://pages.iu.edu/~rdlyons/`.

[174] Russell Lyons, Yuval Peres, and Oded Schramm. Markov chain intersections and the loop-erased walk. *Ann. Inst. H. Poincaré Probab. Statist.*, 39(5):779–791, 2003.

[175] Russell Lyons, Yuval Peres, and Oded Schramm. Minimal spanning forests. *Ann. Probab.*, 34(5):1665–1692, 2006.

[176] Russell Lyons and Oded Schramm. Indistinguishability of percolation clusters. *Ann. Probab.*, 27(4):1809–1836, 1999.

[177] Russell Lyons and Andreas Thom. Invariant coupling of determinantal measures on sofic groups. *Ergodic Theory and Dynamical Systems*, 36(02):574–607, 2016.

[178] S. N. Majumdar and D. Dhar. Equivalence between the abelian sandpile model and the $q \to 0$ limit of the potts model. *Physica A*, (185):129–145, 1992.

[179] Al Marden and Burt Rodin. On thurston's formulation and proof of andreev's theorem. In *Computational methods and function theory*, pages 103–115. Springer, 1990.

[180] Robert Masson. The growth exponent for planar loop-erased random walk. Technical report, University of Chicago, Department of Mathematics, 2008.

[181] Péter Mester. Invariant monotone coupling need not exist. *The Annals of Probability*, 41(3A):1180–1190, 2013.

[182] Gary L. Miller, Shang-Hua Teng, William Thurston, and Stephen A. Vavasis. Separators for sphere-packings and nearest neighbor graphs. *J. ACM*, 44(1):1–29, January 1997.

[183] John Milnor. Growth of finitely generated solvable groups. *Journal of Differential Geometry*, 2(4):447–449, 1968.

[184] Bojan Mohar and Carsten Thomassen. *Graphs on surfaces*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, 2001.

[185] Ben Morris. The components of the wired spanning forest are recurrent. *Probab. Theory Related Fields*, 125(2):259–265, 2003.

[186] C St JA Nash-Williams. Random walk and electric currents in networks. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 55, pages 181–194. Cambridge Univ Press, 1959.

[187] C. M. Newman and L. S. Schulman. Infinite clusters in percolation models. *J. Statist. Phys.*, 26(3):613–628, 1981.

[188] J. R. Norris. *Markov chains*, volume 2 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 1998. Reprint of 1997 original.

[189] Rudolf Peierls. On ising's model of ferromagnetism. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 32, pages 477–481. Cambridge Univ Press, 1936.

[190] Robin Pemantle. Choosing a spanning tree for the integer lattice uniformly. *Ann. Probab.*, 19(4):1559–1574, 1991.

[191] Yuval Peres. Probability on trees: an introductory climb. In *Lectures on probability theory and statistics (Saint-Flour, 1997)*, volume 1717 of *Lecture Notes in Math.*, pages 193–280. Springer, Berlin, 1999.

[192] Yuval Peres, Gábor Pete, and Ariel Scolnicov. Critical percolation on certain nonunimodular graphs. *New York J. Math*, 12:1–18, 2006.

[193] Yuval Peres and David Revelle. Scaling limits of the uniform spanning tree and loop-erased random walk on finite graphs. unpublished, 2004.

[194] Gábor Pete. A note on percolation on $\mathbb{Z}^d$: isoperimetric profile via exponential cluster repulsion. *Electron. Commun. Probab.*, 13:377–392, 2008.

[195] Gabor Pete. Probability and geometry on groups. http://www.math.bme.hu/ gabor/PGG.pdf Retrieved August 2017, 2017.

[196] Georg Pólya. Über eine aufgabe der wahrscheinlichkeitsrechnung betreffend die irrfahrt im straßennetz. *Mathematische Annalen*, 84(1):149–160, 1921.

[197] James Gary Propp and David Bruce Wilson. How to get a perfectly random sample from a generic Markov chain and generate a random spanning tree of a directed graph. *J. Algorithms*, 27(2):170–217, 1998. 7th Annual ACM-SIAM Symposium on Discrete Algorithms (Atlanta, GA, 1996).

[198] Gourab Ray. *Hyperbolic random maps*. PhD thesis, UBC, 2014.

[199] Ian Richards. On the classification of noncompact surfaces. *Transactions of the American Mathematical Society*, 106(2):259–269, 1963.

[200] Burt Rodin and Dennis Sullivan. The convergence of circle packings to the Riemann mapping. *J. Differential Geom.*, 26(2):349–360, 1987.

[201] L. C. G. Rogers and David Williams. *Diffusions, Markov processes, and martingales. Vol. 1*. Cambridge Mathematical Library. Cambridge University Press, Cambridge, 2000. Foundations, Reprint of the second (1994) edition.

[202] Steffen Rohde. Oded Schramm: From Circle Packing to SLE. *Ann. Probab.*, 39:1621–1667, 2011.

[203] Lucio Russo. On the critical percolation probabilities. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 56(2):229–237, 1981.

[204] Thomas S. Salisbury. Energy, and intersections of Markov chains. In *Random discrete structures (Minneapolis, MN, 1993)*, volume 76 of *IMA Vol. Math. Appl.*, pages 213–225. Springer, New York, 1996.

[205] Roberto H. Schonmann. Stability of infinite clusters in supercritical percolation. *Probab. Theory Related Fields*, 113(2):287–300, 1999.

[206] Oded Schramm. Rigidity of infinite (circle) packings. *J. Amer. Math. Soc.*, 4(1):127–149, 1991.

[207] Oded Schramm. Transboundary extremal length. *J. Anal. Math.*, 66:307–329, 1995.

[208] Oded Schramm. Scaling limits of loop-erased random walks and uniform spanning trees. *Israel J. Math.*, 118:221–288, 2000.

[209] Jason Schweinsberg. The loop-erased random walk and the uniform spanning tree on the four-dimensional discrete torus. *Probability Theory and Related Fields*, 144(3-4):319–370, 2009.

[210] Daisuke Shiraishi. Growth exponent for loop-erased random walk in three dimensions. *arXiv preprint arXiv:1310.1682*, 2013.

[211] Ryan Siders. Layered circlepackings and the type problem. *Proceedings of the American Mathematical Society*, 126(10):3071–3074, 1998.

[212] Paolo M. Soardi. Rough isometries and Dirichlet finite harmonic functions on graphs. *Proc. Amer. Math. Soc.*, 119(4):1239–1248, 1993.

[213] Jeffrey E Steif. A survey of dynamical percolation. In *Fractal geometry and stochastics IV*, pages 145–174. Springer, 2009.

[214] Ken Stephenson. Circle pack, java 2.0. `http:/www.math.utk.edu/~kens/CirclePack`.

[215] Kenneth Stephenson. *Introduction to circle packing*. Cambridge University Press, Cambridge, 2005. The theory of discrete analytic functions.

[216] Alain-Sol Sznitman. Vacant set of random interlacements and percolation. *Ann. of Math. (2)*, 171(3):2039–2087, 2010.

[217] Alain-Sol Sznitman. *Topics in occupation times and Gaussian free fields*. Zurich Lectures in Advanced Mathematics. European Mathematical Society (EMS), Zürich, 2012.

[218] A. Teixeira. Interlacement percolation on transient weighted graphs. *Electron. J. Probab.*, 14:no. 54, 1604–1628, 2009.

[219] Augusto Teixeira and Johan Tykesson. Random interlacements and amenability. *The Annals of Applied Probability*, 23(3):923–956, 2013.

[220] Carsten Thomassen. Isoperimetric inequalities and transient random walks on graphs. *Ann. Probab.*, 20(3):1592–1600, 1992.

[221] William P. Thurston. The geometry and topology of 3-manifolds. *Princeton lecture notes.*, 1978-1981.

[222] Ádám Timár. Percolation on nonunimodular transitive graphs. *The Annals of Probability*, pages 2344–2364, 2006.

[223] Ádám Timár. Indistinguishability of components of random spanning forests. *arXiv preprint arXiv:1506.01370*, 2015.

[224] N. Th. Varopoulos. Isoperimetric inequalities and Markov chains. *J. Funct. Anal.*, 63(2):215–239, 1985.

[225] Nicholas Th. Varopoulos. Long range estimates for Markov chains. *Bull. Sci. Math. (2)*, 109(3):225–252, 1985.

[226] A. M. Verˇ sik and V. A. Kaĭ manovič. Random walks on groups: boundary, entropy, uniform distribution. *Dokl. Akad. Nauk SSSR*, 249(1):15–18, 1979.

[227] B. Virág. Anchored expansion and random walk. *Geom. Funct. Anal.*, 10(6):1588–1605, 2000.

[228] David Bruce Wilson. Generating random spanning trees more quickly than the cover time. In *Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996)*, pages 296–303. ACM, New York, 1996.

[229] Wolfgang Woess. *Random walks on infinite graphs and groups*, volume 138 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2000.