Random graphs

the local convergence point of view

N.C.

October 17, 2018

This is the current draft of the lecture notes for the Master 2 course "graphes aléatoires" given during 2015 at the IHÉS and in 2016 at Orsay. The final part of the proof of Theorem 25 has been written by Léo Miolane. Thanks are due to the students of this course for having tested this material and in particular to Henri Elad-Altman, Pierre Rousselin and Guangqu Zheng for their numerous comments. We are grateful to Lasse Leskelä for pointing an inaccuracy in the definition of local convergence.

Contents

Chapter 1

Introduction

The theory of random graphs has been an intense research area in the last few decades. It is motivated by the desire to model real-life random graphs such as the internet, digital social networks... One of its aims is also to make rigorous some of the predictions of theoretical physicists in statistical and quantum physics. But above all, it furnishes a lot of appealing mathematical concepts and problems related to many fields such as combinatorics, ergodic theory, group theory, percolation... It is impossible to survey the monstrous literature on the subject and the purpose of this course is only to give a very restrictive view of the numerous applications of one particular concept: the local convergence of so-called unbiased random graphs.

1.1 Graphs, examples, questions

1.1.1 Graphs

In these notes, a *graph* (sometimes also called a non-oriented multi-graph) is a pair $\mathsf{g} = (\mathrm{V}(\mathsf{g}), \mathrm{E}(\mathsf{g}))$, where $V = \mathrm{V}(\mathsf{g})$ is the set of vertices of g and $E = \mathrm{E}(\mathsf{g})$ is the set of edges of g which is a multiset over $V \times V$ i.e. where repetitions are allowed. The graph is *simple* if they are no multiple edges nor loops. If $x, y \in V$ and $\{x, y\} \in E$ we say that x and y are neighbors and that x and y are

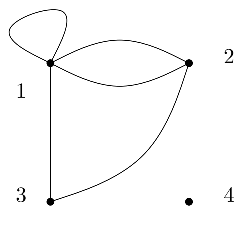

Figure 1.1: An example of a graph $\mathsf{g} = (V, E)$ with vertex set $V = \{1, 2, 3, 4\}$ and edge set $E = \{\!\{\{1,1\}, \{1,2\}, \{1,2\}, \{1,3\}, \{3,2\}\}\!\}$. The vertex degrees of $1, 2, 3, 4$ are respectively $5, 3, 2, 0$.

adjacent to the edge $\{x, y\}$. The *degree* of a vertex $x \in V$ denoted by $\deg(x)$ is the number of half-edges adjacent to x, otherwise said it is the number of edges adjacent to x where loops are counted twice. The graph distance on g is denoted by $\mathrm{d}_{\mathrm{gr}}^{\mathrm{g}}$ or d_{gr} when there is no ambiguity and is defined by

$$\mathrm{d}_{\mathrm{gr}}(x,y) = \text{minimal number of edges to cross to go from } x \text{ to } y.$$

By convention we put $\mathrm{d}_{\mathrm{gr}}(x,y) = \infty$ if there is no path linking x to y in g. The equivalence classes from the relation $x \sim y \iff \mathrm{d}_{\mathrm{gr}}(x,y) < \infty$ are the connected components of g. We say that g is connected if it has only one connected component.

Proposition 1. *Let $\mathrm{g} = (V, E)$ be a connected graph on n vertices. Then we must have $\#E \geq n-1$. If $\#E = n-1$ then g is a tree, meaning that is has no non trivial cycle.*

Proof. We can suppose that the vertex set of g is $\{1, 2, 3, ..., n\}$. We start with the vertex 1, since g is connected there exists an edge adjacent to 1 of the form $\{1, i_1\}$. If $i_1 = 1$ then this edge is a loop otherwise $i_1 \neq 1$. We then throw this edge and pick a new edge adjacent to either 1 or i_1. Iteratively after throwing k edges we have explored a part of the connected component of 1 which has at most $k+1$ vertices. Since g is connected it follows that $\#E \geq n-1$. In case of equality this means that during the exploration process we have never found an edge linking two vertices already explored, in other words no non trivial cycle has been created and g is thus a tree. □

Graph equivalence. If g and g' are two graphs we say that g and g' are equivalent if they represent the same graph up to renaming the vertex set. Formally this means that there exists a bijection $\phi : \mathrm{V}(\mathrm{g}) \to \mathrm{V}(\mathrm{g}')$ which maps the multiset $\mathrm{E}(\mathrm{g})$ to $\mathrm{E}(\mathrm{g}')$: such a function is called a homomorphism of graph (automorphism if $\mathrm{g} = \mathrm{g}'$). Obviously we want to identify two equivalent graphs and thus consider the set of equivalence classes of graphs. This quotient space is monstrous even in the case of countable graphs. The two following exercises should discourage the reader to further dive into the realm of general infinite graphs:

Exercise 1 (Infinite Erdös–Rényi random graph, Rado graph). (∗) Consider the random graph G whose vertex set is $\mathbb{N}$ and where independently for each $i, j \in \mathbb{N}$ the edge $i \leftrightarrow j$ is present in the graph with probability 1/2 (in other words, it is a Bernoulli bond percolation with parameter 1/2 on the graph $(\mathbb{N}, \mathbb{N} \times \mathbb{N})$). Show that almost surely two samples of this random graph are homomorphic.

Open Question 1 (Hadwiger–Nelson problem). *Consider the graph g whose vertex set is $\mathbb{R}^2$ and edge set is $\{\!\{\{x,y\} : |x-y| = 1\}\!\}$. The chromatic number of g is the minimal number of colors necessary to paint the vertices of the graph so that no neighboring vertices have the same color. It is known that the chromatic number of g is in $\{4,5,6,7\}$ but its precise value is unknown and may depend on the axiom of choice.*

1.1.2 Regular graphs

In this section we give several classes of "regular" graphs in order to motivate the forthcoming definitions of stationary or unimodular random graphs which in a sense are generalizations of the following notions.

Definition 1 (Cayley graph). *A Cayley graph encodes the structure of a group. Given a (countable) group* Gr *and a finite symmetric generating set* $S = \{s_1, s_1^{-1}, s_2, s_2^{-1}, \dots, s_k, s_k^{-1}\}$ *of* Gr*, we form the simple graph whose vertices are the elements of the group* Gr *and an edge links* x *to* y *if there exists* $s \in S$ *such that* $x = sy$ *in* Gr*.*

In particular the Cayley graph depends not only upon the group Gr but also on the symmetric generating set we have chosen. For example, here are two Cayley graphs for the group $(\mathbb{Z}, +)$ for the symmetric generating sets $\{-1, +1\}$ and for $\{-3, -2, 2, 3\}$:

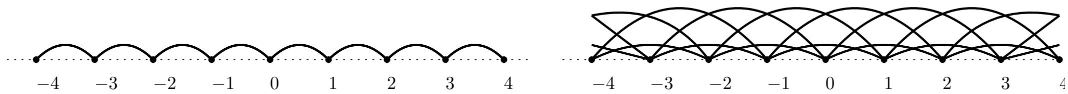

Figure 1.2: Two Cayley graphs of the same group $(\mathbb{Z}, +)$

In particular Cayley graphs are connected and "regular" in the sense that they look the same from any vertex. This is formalized in the following notion:

Definition 2 (Transitive graphs). *A graph* $\mathsf{g} = (V, E)$ *is (vertex)-transitive if for any* $u, v \in V$ *there is a graph automorphism that takes* u *to* v *(in pedantic terms the automorphism group of* g *acts transitively on its vertices).*

Clearly, any Cayley graph is vertex transitive, but the reciproque is false as we will see later. In fact, as it will turn out, Cayley graphs possess on top of transitivity an additional "inverse" property which is not always present in infinite vertex transitive graphs.

Exercise 2. Show that the Petersen graph 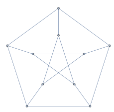 is transitive but is not a Cayley graph.

We could go on and continue weakening the notion of regularity by introducing quasi-transitive graphs (where the action of the automorphism group has only finitely many orbits)... However, since we will be later interested in random graphs any deterministic notion of regularity is doomed. The reader should however keep in mind the last two notions since they will play a great role in these notes.

1.1.3 Examples of random graphs

We give informally several models of random graphs as well as a few questions (some of them still open as of today) to show the diversity of the thematic.

Percolation. Given a finite or countable graph $\mathfrak{g}$, we call Bernoulli bond (resp. site) percolation on $\mathfrak{g}$ with parameter $p \in (0,1)$ the random graph obtained from $\mathfrak{g}$ by keeping independently each edge (resp. vertex) with probability p and erasing it with probability $1-p$.

- When the underlying graph is the complete graph on n vertices (there is an edge between any pair of distinct vertices, so $n(n-1)/2$ edges in total) and when we perform a Bernoulli bond percolation, we speak of the *Erdös–Rényi random graph* and denote it by $G(n,p)$. Introduced in 1959, it is the most famous and most studied model of random graph. This model is referred to as "mean-field" because the geometry of the underlying lattice is trivial in the sense that any pair of vertices are neighbors of each other.

  Question. *Understand the connectivity properties of $G(n,p)$ for very large n as a function of p.*

- When the underlying lattice is the classical lattice $\mathbb{Z}^d$ for $d \geq 2$, it corresponds to the well-known model of Bernoulli bond or site percolation on regular lattices first studied by Hammersley (1957). The main question is then the existence of an infinite cluster. The general theory shows that there is a critical percolation threshold $p_c(\mathbb{Z}^d) \in (0,1)$ such that for $p < p_c(\mathbb{Z}^d)$ there is no infinite cluster and for $p > p_c(\mathbb{Z}^d)$ there is a unique infinite cluster. The following is a well-known tantalizing problem:

  Open Question 2. *Show that there is no infinite cluster as $p = p_c(\mathbb{Z}^3)$ on $\mathbb{Z}^3$.*

Galton–Watson tree. We are given a probability distribution $\mathrm{p} = (p_k)_{k\geq 0}$ on $\{0,1,2,\ldots\}$. Informally, a Galton–Watson tree with offspring distribution p is the genealogical tree obtained by starting from a single ancestor particle and such that each particle reproduces independently of the others by following the offspring distribution p. This model of random tree is also very well-known and very well understood, see Section 1.3.1. It is one of the key characters in the theory of random graphs.

Uniform graph in a combinatorial class. Let $\mathscr{C}_n$ be a combinatorial class of graphs with finite size. We can for example take $\mathscr{C}_n$ to be the set of all graphs with n edges, the set of all graphs with n edges and $\alpha(n)$ vertices, the set of all trees with n vertices, the set of all planar graphs with n edges... Provided that $\#\mathscr{C}_n$ is finite (which is the case in the above examples) then we can consider $C_n \in \mathscr{C}_n$, a random variable uniformly distributed over $\mathscr{C}_n$.

Open Question 3. *Understand the large scale geometry of C_n when $\mathscr{C}_n$ is for example the set of all planar graphs with n edges.*

Dynamical random graphs. We can also consider a sequence of growing random graphs which are built recursively. The (very interesting) prototype is the Barabasi–Albert model [10] which was introduced to model the degree distribution of many real networks such as the internet.

The initial graph G_1 is just a single vertex. Then inductively for $n \geq 1$ we choose in G_n a vertex proportionally to its degree (the number of adjacent edges) and we attach a new leaf to this vertex. Hence, the vertices with higher degree get more likely chosen : rich get richer.

Question. *Understand the empirical degree distribution of the vertices in G_n.*

Configuration model. The configuration model is designed so that the degree sequences of the vertices is fixed a priori. More precisely, let $\mathrm{d} = (d_1, ..., d_n)$ be a sequence of integers so that $D = \sum d_i$ is even. We will then consider a random (multi) graph on n vertices $\{1, 2, ..., n\}$ such that the vertex i has degree d_i. One way to sample such a graph is to start with those n vertices to which we attach d_i half-edges or stubs which are labeled from 1 to D. We then pair those half-edges in the most natural way: we couple the half-edge number 1 with a uniformly chosen half-edge among those numbered $2, 3, ..., D$. We then merge the two stubs involved and create a true edge. We then iterate the procedure with the remaining stubs. We denote by $G(\mathrm{d})$ the random graph obtained. This graph may not be simple (it may contain multiple edges or loops).

Miscellaneous. There are many more models of random graphs. We give a final example to re-interpret a well-known model in terms of a random graph. On $\mathbb{Z}^2$ consider a simple random walk started from the origin $(0,0)$ until it escapes from the square $[-n, n]^2$. The trace (the edges and vertices visited by the walk) can thus be considered as a random subgraph G_n of $[-n, n]^2$. Although the "shape" of this trace is very well understood (thanks to the convergence of the random walk towards Brownian motion), its geometry seen as graph is still elusive:

Open Question 4 (K. Burdzy)**.** *In the above random graph prove (or disprove!) that the distance between 0 and $\partial[-n, n]^2$ is of order $n^{1+o(1)}$.*

1.1.4 In search of a limit model

In some of the above examples, we have a model of a random graph of finite size n. It would be desirable to define a limiting infinite model. Indeed, it is common in mathematics that asymptotic questions on models of growing size could often be resolved by a direct analysis on an appropriate infinite limiting model. Let us give tree examples of such questions:

Degree distribution. If (G_n) is a sequence of model of (random) graphs whose size tends to infinity, one may want to understand the degree distribution i.e. the random vector

$$\mathrm{p}_n(k) = \frac{\#\{u \in \mathrm{V}(G_n) : \deg(u) = k\}}{\#\mathrm{V}(G_n)}.$$

Spanning trees. A spanning tree of a (connected) graph $\mathfrak{g}$ is a connected subgraph of $\mathfrak{g}$ without cycle which spans all the vertices of $\mathfrak{g}$. Now if (G_n) is a sequence of model of (random) graphs whose size tends to infinity, one would like to understand the asymptotic number of spanning trees

$$\frac{\log \#\mathrm{SpanTrees}(G_n)}{|G_n|}, \qquad n \to \infty.$$

Matching number. A matching on g is a subset of mutually non-adjacent edges on g. We denote by $\nu(\mathrm{g})$ the largest size of a matching on g. Again if (G_n) is a sequence of model of (random) graphs whose size tends to infinity, one would like to understand the asymptotic

$$\frac{\nu(G_n)}{|G_n|}, \qquad n \to \infty.$$

All these asymptotic enumeration problems can indeed be answered [35, 1] by looking at an appropriate limit of the G_n, the so-called Benjamini-Schramm or local limit. The goal of this course is to describe this limit which is well suited to the analysis of so-called "dilute" random graphs where the average number of edges per vertex (or mean degree) remains typically bounded.

1.2 Local convergence topology

From now on, unless explicitly mentioned, all the graphs g considered are

- countable i.e. $\mathrm{E}(\mathrm{g})$ is countable,
- locally finite i.e. $\deg(x) < \infty$ for all $x \in \mathrm{V}(\mathrm{g})$,
- connected.

1.2.1 Local topology

Imagine that we are given a metric space $(\mathcal{E}, \delta)$ such that for any $x \in \mathcal{E}$ and for any $r \in \{0, 1, 2, 3, ...\}$, there is a notion of *restriction of radius* r in x that we denote by

$$[x]_r \in \mathcal{E}.$$

We suppose that for any $r \geq 0$ the map $x \mapsto [x]_r$ is continuous for δ and that different restrictions of radius r are compatible in the sense that $[[x]_{r'}]_r = [x]_r$ for any $r' \geq r$. We assume that for any $r \geq 0$, the set $\{[x]_r : x \in \mathcal{E}\}$ is separable and complete for δ. We also suppose that for any sequence of coherent elements $x_0, x_1, \ldots$ satisfying $[x_i]_j = x_j$ for any $0 \leq j \leq i$, there exists a unique "*infinite*" element $x \in \mathcal{E}$ such that

$$[x]_r = x_r, \qquad \text{for all } r \geq 0. \tag{1.1}$$

Then we endow the space $\mathcal{E}$ with a distance, called the *local distance* defined as

$$\mathrm{d}_{\mathrm{loc}}(x, y) = \sum_{r \geq 0} 2^{-r} \min\Big(1, \delta([x]_r, [y]_r)\Big).$$

In other words, a sequence x_n of elements of $\mathcal{E}$ converges towards x for the local distance if and only if for any $r \geq 0$, the sequence $[x_n]_r$ converges to $[x]_r$ for the metric δ as $n \to \infty$. Then we have the following:

Theorem 2 (*The local topology is Polish*)

The space $(\mathcal{E}, \mathrm{d}_{\mathrm{loc}})$ is a Polish space, that is metric, separable and complete. Furthermore, a subset $\mathcal{A} \subset \mathcal{E}$ is pre-compact (meaning that its closure is compact) if and only if

$$\text{for every } r \geq 0 \text{ we have } \{[x]_r : x \in \mathcal{A}\} \text{ is pre-compact.}$$

Proof. Distance. Let us first show that $\mathrm{d}_{\mathrm{loc}}$ is a distance. The symmetry is easy, as well as the triangle inequality. The separability is also easy since we suppose that if $[x]_r = [y]_r$ for all $r \geq 0$ then we have $x = y$ by (1.1). **Separation.** For any $x \in \mathcal{E}$, we have $\mathrm{d}_{\mathrm{loc}}(x, [x]_r) \leq 2^{-r}$ and we supposed that the set $\{[x]_r : x \in \mathcal{E}\}$ of all restrictions of radius r is separable. Hence we make a union over r of dense countable sets in $\{[x]_r : x \in \mathcal{E}\}$ which end-up being countable and dense for $\mathrm{d}_{\mathrm{loc}}$. **Completeness.** If (x_n) is a Cauchy sequence for $\mathrm{d}_{\mathrm{loc}}$ then for every r, the restriction $[x_n]_r$ is again Cauchy and, by completeness of $\{[x]_r : x \in \mathcal{E}\}$, thus converges for δ to a certain element $y_r \in \{[x]_r : x \in \mathcal{E}\}$. By continuity of $x \mapsto [x]_r$ we deduce that $[y_{r'}]_r = y_r$ for any $r' \geq r$ and so by the coherence property (1.1) we can define a unique element $y \in \mathcal{E}$ such that $y_r = [y]_r$. It is then clear that $x_n \to y$ for $\mathrm{d}_{\mathrm{loc}}$. **Characterization of the compacts.** The condition in the theorem is clearly necessary for $\mathcal{A}$ to be pre-compact for otherwise there exists $r_0 \geq 0$ and a sequence (x_n) in $\mathcal{A}$ whose restrictions of radius r_0 are all at distance ε for each other. Such a sequence cannot admit a convergent subsequence. Conversely, a subset $\mathcal{A}$ satisfying the condition of the theorem is easily seen to be pre-compact for $\mathrm{d}_{\mathrm{loc}}$: just cover with restrictions of radius 2^{-r} centered on an 2^{-r}-net for δ of $\mathcal{A}$ to get a 2^{-r}-net for $\mathrm{d}_{\mathrm{loc}}$. □

Examples :

- We can consider the space $(\mathcal{C}(\mathbb{R}_+, \mathbb{R}), \|\|_\infty)$. This space is not Polish but satisfies the above assumptions once we put for the restriction of radius r

$$[f]_r(x) = f(x)1_{x \leq r} + f(x)(r + 1 - x)1_{r \leq x \leq r+1}.$$

  The local topology we create coincides with the uniform convergence of all compact sets of $\mathbb{R}_+$ and is now Polish.

- We can consider the space $\mathcal{D}$ made of all possible (finite and infinite) words on a finite alphabet. This discrete space is endowed with the trivial distance $\delta(x, y) = 1_{x \neq y}$. If w is such a word, then we can set $[w]_r$, the restriction of radius r to be the word made of the first r letters of w. It is easy to see that all the conditions are satisfied.

- More generally when the space $\mathcal{E}$ is discrete it will always be endowed with the trivial distance. Then the above process enables to put a topology on $\mathcal{E}$ turning its rough metric into a Polish topology called the local topology. In this lecture notes we will apply this construction of the local topology to the following discrete structures:

  – The space $\mathcal{G}^\bullet$ of pointed graphs,

– The space $\vec{\mathcal{G}}$ of rooted graphs,

– The space $\mathcal{G}^{\leftrightarrow}$ of graphs with a path on it,

– The space $\mathcal{M}^{\bullet}$ of pointed planar maps,

– The space $\vec{\mathcal{M}}$ of rooted planar maps,

– The space $\mathcal{T}$ of rooted plane trees,

We will always write $\mathrm{d_{loc}}$ for the local distance and it will be clear from the context what we exactly mean. It can also be applied to graphs (non necessary locally-finite) carrying labels on the edges [5], or even more exotic structure such as equivalence classes of locally compact metric spaces endowed with the Gromov-Hausdorff distance [21]. Before giving examples of applications below, let us recall basic notions in probability theory. A random variable X is a measurable function from the underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with values in the Polish space $(\mathcal{E}, \mathrm{d_{loc}})$ endowed with the Borel σ-field denoted by $\mathcal{B}_{\mathrm{doc}}$. Hence we have the natural notion of convergence in distribution, namely if $(X_n)_{n\geq 0}$ is a sequence of random variables, then X_n converges in distribution (for the local topology) towards a random variable X if for any bounded continuous function $F : \mathcal{E} \to \mathbb{R}$ we have

$$\mathbb{E}[F(X_n)] \xrightarrow[n\to\infty]{} \mathbb{E}[F(X)].$$

We recall the fundamental concept of tightness for random variables with values in a Polish space: A family $(X_i)_{i\in I}$ of random variables is tight for $\mathrm{d_{loc}}$ if for any $\varepsilon > 0$ there exists a compact $A_\varepsilon \subset \mathcal{E}$ such that for any $i \in I$ we have

$$\mathbb{P}(X_i \in A_\varepsilon) \geq 1 - \varepsilon.$$

In our case, by the construction of the local topology it is easy to check:

Proposition 3. *A family $(X_i : i \in I)$ of random variables with values in $\mathcal{E}$ is tight if and only if for any $r \geq 1$ the family $([X_i]_r : i \in I)$ is tight.*

Proposition 4 (Characterization of the law and of the convergence for the local topology)**.**

(i) Let X_1 and X_2 be two random variables with values in $(\mathcal{E}, \mathrm{d_{loc}})$ such that for any $A \in \mathcal{B}_{\mathrm{loc}}$ and any $r \geq 0$ we have $\mathbb{P}([X_1]_r \in A) = \mathbb{P}([X_2]_r \in A)$, then $X_1 = X_2$ in distribution.

(ii) A sequence of random variables $(X_n)_{n\geq 0}$ converges in distribution towards X_∞ if and only if for every $r \geq 0$ and every $A \in \mathcal{B}_{\mathrm{loc}}$ we have $\mathbb{P}([X_n]_r \in A) \to \mathbb{P}([X_\infty]_r \in A)$ as $n \to \infty$.

Proof. (i) We consider the family of events

$$\mathcal{M} = \Big\{\{x \in \mathcal{E} : [x]_r \in A\} : A \in \mathcal{B}_{\mathrm{loc}}, r \geq 0\Big\}.$$

It is easy to see that the family $\mathcal{M}$ generates the Borel σ-field on $\mathcal{E}$ and moreover that $\mathcal{M}$ is stable under finite intersections. It follows from the monotone class theorem that two random variables agreeing on $\mathcal{M}$ have the same law.

(ii) For the second point, we have seen above that the sets $\{[x]_r \in A\}$ are stable under finite intersections and it is easy to see that any open sets of the local topology can be written as a countable union of those sets. The result then follows from [15, Theorem 2.2]. □

In particular, we deduce that for a sequence of random variables $(X_n)_{n\geq 0}$ to converge it is necessary and sufficient that $(X_n)_{n\geq 0}$ be tight and that for every $r \geq 0$ and every $A \in \mathcal{B}_{\text{loc}}$ we have $\mathbb{P}([X_n]_r \in A)$ converges. (The two conditions are necessary, exercise!).

1.2.2 Two examples: random pointed graphs and plane trees

Plane trees. We recall the formalism for plane trees as found in [39]. Let

$$\mathcal{U} = \bigcup_{n=0}^{\infty} (\mathbb{N}^*)^n$$

where $\mathbb{N}^* = \{1, 2, \ldots\}$ and $(\mathbb{N}^*)^0 = \{\varnothing\}$ by convention. An element u of $\mathcal{U}$ is thus a finite sequence of positive integers. We let $|u|$ be the length of the word u. If $u, v \in \mathcal{U}$, uv denotes the concatenation of u and v. If v is of the form uj with $j \in \mathbb{N}$, we say that u is the *parent* of v or that v is a *child* of u. More generally, if v is of the form uw, for $u, w \in \mathcal{U}$, we say that u is an *ancestor* of v or that v is a *descendant* of u.

Definition 3. *A (locally finite)* plane tree τ *is a (finite or infinite) subset of* $\mathcal{U}$ *such that*

1. $\varnothing \in \tau$ *(*$\varnothing$ *is called the* root *of* τ*),*
2. *if* $v \in \tau$ *and* $v \neq \varnothing$*, the parent of* v *belongs to* τ
3. *for every* $u \in \mathcal{U}$ *there exists* $k_u(\tau) \in \{0, 1, 2, \ldots\}$ *such that* $uj \in \tau$ *if and only if* $j \leq k_u(\tau)$*.*

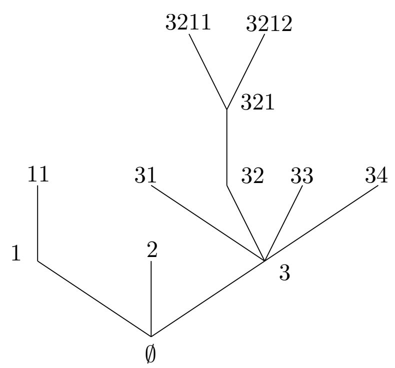

Figure 1.3: A finite plane tree.

A plane tree can be seen as a graph, in which an edge links two vertices u, v such that u is the parent of v or vice-versa. Notice that with our definition, vertices of infinite degree are not allowed since k_u cannot be infinite. This graph is of course a tree in the graph-theoretic sense, and has a natural embedding in the plane, in which the edges from a vertex u to its children $u1, \dots, uk_u(\tau)$ are drawn from left to right. All the trees considered in these pages are plane trees. The integer $|\tau|$ denotes the number of edges of τ and is called the size of τ.

We denote by $\mathcal{T}$ the set of all (finite or infinite) plane trees (sometimes called rooted plane trees). If $\tau \in \mathcal{T}$ is a plane tree, the notion of restriction $[\tau]_r$ of radius r we use is just the plane tree obtained by keeping all the vertices of τ which are at generation less than r from the origin, that is $[\tau]_r = \{u \in \tau : |u| \leq r\}$. Using Theorem 2 this gives rise to the local topology on $\mathcal{T}$ which is then Polish. A random plane tree will thus be seen as a random variable with values in $(\mathcal{T}, \mathrm{d_{loc}})$.

Random pointed graphs.

Definition 4. *A pointed graph* $\mathsf{g}^\bullet$ *is a pair* (g, ρ) *where* g *is a (countable, locally finite, connected) graph and* $\rho \in \mathrm{V}(\mathsf{g})$ *is a reference vertex sometimes called the origin of the graph. Two pointed graphs* $\mathsf{g}^\bullet = (\mathsf{g}, \rho)$ *and* $\mathsf{h}^\bullet = (\mathsf{h}, \varrho)$ *are equivalent if there exists a graph homomorphism between* $\mathsf{g}^\bullet$ *and* $\mathsf{h}^\bullet$ *which sends* ρ *onto* ϱ *(we speak of pointed graph homomorphism).*

In what follows we will obviously identify equivalent graphs and so formally work on the space of equivalence classes of pointed graphs. We will implicitly make this identification and later speak of pointed graphs (instead of equivalence classes of pointed graphs). We introduce $\mathcal{G}^\bullet$ the set of all (equivalence classes) of (locally finite, countable, connected) pointed graphs. If $\mathsf{g}^\bullet$ is a pointed graph, we denote by $[\mathsf{g}^\bullet]_r$, the restriction of radius r around the origin of $\mathsf{g}^\bullet$ to be the (equivalence class of the) graph obtained from $\mathsf{g}^\bullet = (\mathsf{g}, \rho)$ by keeping only those vertices which are at distance less than r from ρ and the edges between them; the resulting graph being pointed at ρ. The compatibility relations of the restrictions are easy to check. The fact that there are only countably many restrictions of radius r is also easy to see. It requires a bit of thought to show that if $\mathsf{g}_1^\bullet, \mathsf{g}_2^\bullet, \dots$ is a sequence of combatible graphs in the sense that $[\mathsf{g}_j^\bullet]_r = \mathsf{g}_r^\bullet$ for $r \leq j$ then there exists a unique (equivalence class of a) infinite pointed graph $\mathsf{g}^\bullet$ whose restrictions of radius r is $\mathsf{g}_r^\bullet$ but this is also true. Hence we can apply the above procedure to endow $\mathcal{G}^\bullet$ with a local distance making it a Polish space. In this case, it is also easy to check that the pre-compact subset $\mathcal{A} \subset \mathcal{G}^\bullet$ are those satisfying for every $r \geq 0$ we have

$$\sup_{\mathsf{g}^\bullet \in \mathcal{A}} \sup_{x \in \mathrm{V}([\mathsf{g}^\bullet]_r)} \deg(x) < \infty.$$

From the previous characterization of pre-compact sets in $\mathcal{G}^\bullet$, it follows that for any $M > 0$, the subset of pointed graphs where the degree of any vertex is bounded by M is a compact set.

Exercise 3. Show that the following family of pointed graphs is not compact for the local topology on pointed graphs

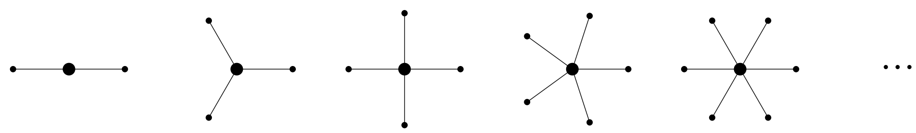

Exercise 4. Compute the local limit of the following 6 sequences of pointed graphs:

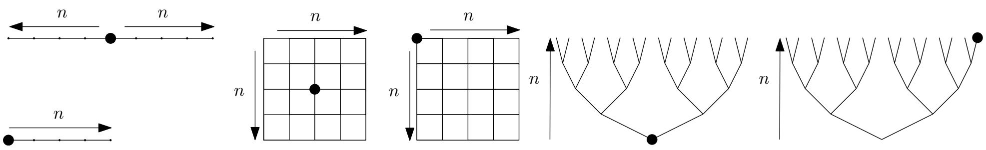

Exercise 5. Show that a family of random pointed graphs $(G_i^\bullet)_{i\in I}$ is tight if and only if for any $r \geq 0$ the family of random variables

$$\max_{x\in \mathrm{V}([G_i^\bullet]_r)} \deg(x), \qquad i \in I$$

is tight as a family of real-valued random variables.

A random pointed graph will be, in these notes, a random variable $G^\bullet$ taking values in the Polish space $(\mathcal{G}^\bullet, \mathrm{d_{loc}})$ endowed with the Borel σ-field.

Exercise 6. Let g be a deterministic connected graph and let $x \in \mathrm{V(g)}$. We launch a simple random walk on g starting from x, and denote by G_n the random graph obtained by keeping those vertices and edges that have been visited by the random walker before time n. Show that (G_n, x) converges in distribution as $n \to \infty$ and that its limit is (g, x) if and only if g is recurrent.

1.3 The prototype: Erdös–Renyi → Poisson Galton–Watson tree

In this section we describe the most famous local convergence result, in a sense, the prototypical example which underlies much of the theory on the Erdös–Renyi random graph. It says that in the appropriate range of parameters we have the local convergence of $G(n,p)$ towards a Poisson Galton–Watson tree.

1.3.1 Reminder on Galton–Watson trees

We first recall briefly the basics on Galton–Watson trees. If τ is a plane tree, for any vertex $u \in \tau$, we denote the shifted tree at u by $\sigma_u(\tau) := \{v \in \tau : uv \in \tau\}$. I

Definition 5. *Let* $\mathrm{p} = (p_{0,1}, p_2, ...)$ *be a probability measure on* $\mathbb{N}$ *such that* $p_1 < 1$. *The law of the Galton–Watson tree with offspring distribution* p *(abbreviated by* p-*GW) is the unique probability measure* GW_p *on* $\mathcal{T}$ *such that:*

1. $\mathsf{GW}_\mathrm{p}(k_\varnothing(\tau) = j) = p_j$ *for* $j \geq 0$,
2. *for every* $j \geq 1$ *with* $p_j > 0$, *conditionally on* $\{k_\varnothing(\tau) = j\}$, *the subtrees* $\sigma_1(\tau), \dots, \sigma_j(\tau)$ *are i.i.d. with distribution* GW_p.

It is easy to check from the above definition that for any finite tree $\tau_0 \in \mathcal{T}$ we have

$$\mathsf{GW}_\mathrm{p}(\tau = \tau_0) = \prod_{u\in\tau_0} p_{k_u(\tau_0)},$$

but the above display does not characterize the measure GW_{p} because in general a random tree τ under GW_{p} may very well be infinite... as shown in the next result:

Theorem 5 (***Extinction probabilities for Galton–Watson trees***)

We introduce $\mathrm{F}_{\mathrm{p}}(z) = \sum_{k\geq 0} p_k z^k$ *the generating function of the offspring distribution* p *and suppose that* $p_1 \neq 1$. *Then the probability that a* p*-GW tree is finite (a.k.a. extinction probability) is the smallest solution* $z \in [0,1]$ *to*

$$\mathrm{F}_{\mathrm{p}}(z) = z. \tag{1.2}$$

In particular, this probability is strictly less than 1 *if and only if the mean of the offspring distribution* $m = \sum_{k\geq 0} k\, p_k$ *is strictly larger than* 1. *When* $m < 1$ *the offspring distribution is called subcritical, critical when* $m = 1$ *and supercritical when* $m > 1$.

Proof: We denote by E_n the event on which the height of τ is strictly less than n (in other words, the genealogy is extinct at time n). The probability that τ is finite is then the probability of the event

$$E_\infty = \lim_{n\to\infty} E_n = \bigcup_{n\geq 0} E_n.$$

For $n \geq 0$, we can write E_{n+1} has the intersection of the event $E_n(\sigma_1(\tau)), ..., E_n(\sigma_{k_\varnothing(\tau)}(\tau))$ of extinction at time n of the subtrees of the children of the origin vertex. By the definition of the Galton–Watson measure we can thus write the recursion

$$\mathsf{GW}_{\mathrm{p}}(E_{n+1}) = \sum_{k\geq 0} p_k \mathsf{GW}_{\mathrm{q}}(E_n)^k = \mathrm{F}_{\mathrm{p}}\Big(\mathsf{GW}_{\mathrm{q}}(E_n)\Big).$$

We deduce that $x_n = \mathsf{GW}_{\mathrm{p}}(E_n)$ is a sequence defined by the function recursion $x_n = \mathrm{F}_{\mathrm{p}}(x_{n+1})$ and started at $x_0 = 0$. Using the (strict) convexity of F_{p} and the fact that $\mathrm{F}_{\mathrm{p}}(1) = 1$ we conclude that x_n tends to the smallest fixed point of F_{p} on $[0,1]$. This proves the first point of the theorem. For the second point we notice that m is the derivative of F_{p} at 1 and so by convexity F_{p} is strictly above $y = x$ on $[0,1)$ if $m = 1$ and crosses it once if $m > 1$. □

Exercise 7. We say that an infinite tree τ contains an infinite binary tree (starting at the root) if it is possible to find a subset of vertices B of τ containing the origin $\varnothing$ and such that each $u \in \mathrm{B}$ has exactly two children in B. Then the probability that τ contains no infinite binary tree (starting at the root) is the smallest solution $z \in [0,1]$ to

$$z = \mathrm{F}_{\mathrm{p}}(z) + (1-z)\mathrm{F}'_{\mathrm{p}}(z).$$

Application: in the case $p_1 = (1-p)$ and $p_3 = p$ with $p \in [0,1]$ show that there is no infinite binary tree in τ if and only if $p < \frac{8}{9}$ and that in the critical case $p = \frac{8}{9}$ this probability is in fact positive (contrary to the above case for survival of the tree).

Of course if τ is a plane tree, one can associate with it a pointed graph which we denote by $\tau^\bullet$ obtained by considering the graph made by the vertices of τ and putting an edge between each parent and their children. The reference point of $\tau^\bullet$ is $\varnothing$. This application $\mathcal{T} \to \mathcal{G}_\bullet$ is easily checked to be continuous for the respective local distances.

1.3.2 Local convergence of Erdös–Rényi

Consider the complete graph on n vertices and perform a Bernoulli bond percolation with parameter $p \in (0,1)$. The graph obtained may have many connected components and we denote by $G^\bullet(n,p)$ the random graph made of the connected component of 1 pointed at the vertex 1. For $\lambda > 0$, we also denote by $T_\lambda^\bullet$ the random graph obtained from a Galton–Watson tree with offspring distribution Poisson(λ) pointed at the ancestor vertex.

Theorem 6

For any $\lambda > 0$ we have the following convergence in distribution for $\mathrm{d_{loc}}$

$$G^\bullet(n,\lambda/n) \xrightarrow[n\to\infty]{} T_\lambda^\bullet.$$

Let us sketch an intuitive proof of the last result. First, the number of vertices adjacent to the vertex 1 has a Bin($n-1,\lambda/n$) distribution and the latter converges towards a Poisson(λ) distribution. Hence the 1-neighborhood of 1 in $G(n,\lambda/n)$ looks like the 1-neighborhood of the origin in a Poisson(λ)-Galton–Watson tree. We then pass to the neighbors of 1: Each of these vertices is linked to 1 by definition and the number of other vertices to which it is linked follows a Bin($n-2,\lambda/n$) law which is also converging towards Poisson(λ). The point here is to remark that these variables are roughly independent and that it is very unlikely that when exploring the 2-neighborhood of 1 we discover a cycle, i.e. a vertex linked to two different neighbors of 1. Hence, the 2-neighborhood of 1 looks like the 2-neighborhood of a Poisson(λ)-Galton–Watson tree. We then proceed similarly to explore the $3,4,5,\ldots$-neighborhoods of the vertex 1. Although the above sketch can be made rigorous (see [17]) and constitutes the basis of the analysis of the Erdös–Rényi model around criticality, we give another proof using the rigidity of plane tree structure.

Proof. Fix a pointed tree $\mathsf{t}^\bullet = (\mathsf{t},\rho)$ of height at most r. We will prove that

$$\mathbb{P}([G^\bullet(n,\lambda/n)]_r = \mathsf{t}^\bullet) \xrightarrow[n\to\infty]{} \mathbb{P}([T_\lambda^\bullet]_r = \mathsf{t}^\bullet). \tag{1.3}$$

To this end it will be easier to add an additional ordering on the vertices of the tree in order to break possible symmetries. We thus consider a *plane* tree t such that its pointed-graph version is indeed $\mathsf{t}^\bullet$ (there are generally several choices for this plane tree). We will also use the numbering of the vertices of the graph $G^\bullet(n,\lambda/n)$ to our advantage. We will write $[G^\bullet(n,\lambda/n)]_r \equiv \mathsf{t}$ if we can map the vertices of the restriction of radius r around 1 to the vertices of $[\mathsf{t}]_r$ such that $1 \to \varnothing$ in a way which preserves the graph structure and such that the order of the "children" in t coincides with the order of the vertices in $G^\bullet(n,\lambda/n)$. With this notation in hands we have

$$\mathbb{P}\Big([G^\bullet(n,\lambda/n)]_r \equiv \mathsf{t}\Big) = \mathsf{N}_{n,\mathsf{t}} \cdot \underbrace{\prod_{u\in\mathsf{t}:|u|<r} \left(\frac{\lambda}{n}\right)^{k_u(\mathsf{t})}}_{\text{edges present}} \cdot \underbrace{\left(1-\frac{\lambda}{n}\right)^{n-1-\deg(x)}}_{\text{edges absent}},$$

where $\mathsf{N}_{n,\mathsf{t}}$ is the number of ways to assign different numbers in $\{1,2,\dots,n\}$ to the vertices of t so that the ancestor gets label 1 and so that the numbers assigned of the children of a given vertex are increasing from left to right. It is easy to see that we have the asymptotic

$$\mathsf{N}_{n,\mathsf{t}} \underset{n\to\infty}{\sim} \prod_{u\in\mathsf{t}:|u|<r} \frac{n^{k_u(\mathsf{t})}}{(k_u(\mathsf{t}))!}.$$

Combining the last displays we get that for any plane tree t of height at most r we have

$$\mathbb{P}([G^\bullet(n,\lambda/n)]_r \equiv \mathsf{t}) \to \mathbb{P}([T_\lambda]_r = \mathsf{t}), \qquad \text{as } n\to\infty.$$

Removing the orientation and summing over all trees yielding to a given tree neighborhood we get (1.3). Using Proposition 4 it thus suffices to prove that for any pointed graph $\mathsf{g}^\bullet$ which is not a tree we have $\mathbb{P}([G^\bullet(n,\lambda/n)]_r = \mathsf{g}^\bullet) \to 0$ as $n\to\infty$. But this is easy since by Fatou's lemma we have

$$\begin{aligned}\limsup_{n\to\infty} \mathbb{P}([G^\bullet(n,\lambda/n)]_r \text{ is not a tree}) &= 1 - \liminf_{n\to\infty} \sum_{\mathsf{t}^\bullet,\ \text{tree}} \mathbb{P}([G^\bullet(n,\lambda/n)]_r = \mathsf{t}^\bullet)\\ &\le 1 - \sum_{\mathsf{t}^\bullet,\ \text{tree}} \mathbb{P}([T^\bullet_\lambda]_r = \mathsf{t}^\bullet)\\ &= 0.\end{aligned}$$

□

The last result together with Theorem 5 is an indication that in $G(n,\lambda/n)$, a dramatic change appears at $\lambda = 1$. Indeed when $\lambda < 1$ a Galton–Watson tree with Poisson(λ) offspring distribution is almost surely finite, whereas for $\lambda > 1$ the later has a positive probability $q_\lambda > 0$ of being infinite. This phase transition in the underlying Galton–Watson reflects into a phase transition in the behavior of the Erdös–Rényi random graph: When $\lambda < 1$ then all the connected components in $G(n,\lambda/n)$ are of size $O(\log n)$ with high probability as $n\to\infty$, whereas for $\lambda > 1$ there exists a unique giant connected component of $G(n,\lambda/n)$ of size roughly $q_\lambda \cdot n$ and all the other components are of size $O(\log n)$ with high probability as $n\to\infty$. Strictly speaking, Theorem 6 does not permit to deduce the above results but it illustrates the philosophy underlying the proof of this theorem. We refer to [44] for much more about the description of this phase transition in the $G(n,p)$ model. As we said, this model is often referred to as "mean-field" and this property translates in our context by the fact that the local limit is a tree. This property is shared by the complete graph but also by other lattices of "sufficiently high dimension":

Exercise 8 (Hypercube). Suppose that we replace the complete graph in the definition of the Erdös-Rényi model by the hypercube $\{0,1\}^n$ with an edge between two vertices having only one coordinate which differ. Denote $H^\bullet(n,\frac{\lambda}{n})$ the component of $(0,0,\dots,0)$ in a percolation on $\{0,1\}^n$ with parameter $\frac{\lambda}{n}$ pointed at this vertex. Show that as above

$$H^\bullet(n,\frac{\lambda}{n}) \xrightarrow[n\to\infty]{} T^\bullet_\lambda.$$

Exercise 9 (High dimension). Variation on the above exercise: Let $Z^\bullet(d, \frac{\lambda}{2d})$ be the component of the origin in a percolation of parameter $\frac{\lambda}{2d}$ in $\mathbb{Z}^d$ with the standard edge set (again pointed at the origin). Show again that

$$Z^\bullet(d, \frac{\lambda}{2d}) \xrightarrow[d\to\infty]{} T_\lambda^\bullet.$$

Bibliographical notes. A vast majority of the literature on random graphs deals with the Erdös-Rényi model and its generalizations. The configuration model has been introduced by Bollobas [16]. We refer the interested reader to the wonderful book-to-be of Van der Hofstad [44] which also contains a lot of information on dynamical random graphs such as the Barabasi-Albert model. There is a different theory to make sense of "dense" random graphs developed recently by Lovász and co-authors [34] where the number of edges is typically of order the number of vertices squared (think of $G(n, p)$ when p is bounded away from 0 as $n \to \infty$). Although implicit in many earlier works, the local convergence of random graphs has been formally introduced by Benjamini–Schramm in [14]. The formalism for plane trees has been introduced by Neveu [39], see [3] for nice lecture notes on Galton–Watson trees and their local limits. Exercise 7 is a result of Dekking [22]. See [17, 38] for lecture notes available on the web which also treat random graphs under the local convergence point of view as done in this course.

Chapter 2

Unimodular random graphs

Roughly speaking, a unimodular random graph is a random graph where the origin plays no special role. We start with an easy example to make the reader grasp the key concept. Consider the case of the line graph $\{-n, -n+1, ..., -1, 0, 1, ..., n\}$ with the obvious edge set. If this graph is pointed at $-n$ or at n, it converges locally as $n \to \infty$ towards the line graph $\mathbb{N}$; whereas if it is pointed at the vertex 0 then it converges towards the line $\mathbb{Z}$. Which of the above limits make more sense? The key idea is to let the randomness choose the pointed vertex by distinguishing a vertex of the graph uniformly at random.

2.1 Mass-transport principle

2.1.1 Finite case

Definition 6. *Let $G^\bullet$ be a finite (connected) random pointed graph. We say that $G^\bullet$ is* uniformly pointed *if the reference point is uniform over the vertices of G, namely if for any measurable $f : \mathcal{G}^\bullet \to \mathbb{R}_+$ we have*

$$\mathbb{E}[f(G^\bullet)] = \mathbb{E}\left[\frac{1}{\#\mathrm{V}(G)} \sum_{x \in \mathrm{V}(G)} f(G, x)\right].$$

Exercise 10. Compute, if they exist, the local limit of the following uniformly pointed random graphs:

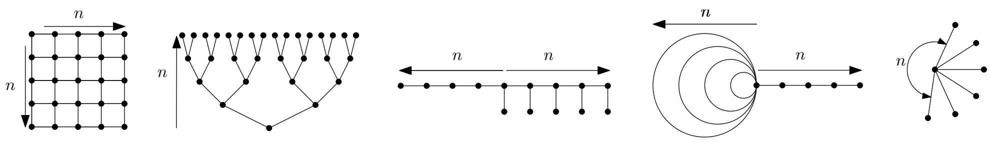

2.1.2 Unimodular random graphs

The problem with the last definition is that it a priori does not make sense for infinite graphs. We will thus give an equivalent definition which can be extended to the infinite case. To do so, we first introduce the set $\mathcal{G}^{\bullet\bullet}$ of (equivalence classes) of doubly-pointed graphs (g, x, y). More precisely, two doubly-pointed graphs $(\mathbf{g}, x, y)$ and $(\mathbf{g}', x', y')$ are identified if there exists a graph homomorphism from $\mathbf{g}$ to $\mathbf{g}'$ sending x to x' and y to y'. This set is endowed with a local topology for the notion of restriction $[(\mathbf{g}, x, y)]_r$ given by the graph made of all the vertices and edges which are at distance less than r from either x or y subject also to the condition that $\mathrm{d_{gr}}(x, y) \le r$ and where the resulting graph is doubly pointed at x and y. Hence we can apply the general construction of Section 1.2.1 to get that $(\mathcal{G}^{\bullet\bullet}, \mathrm{d_{loc}})$ is a Polish space.

Exercise 11. Show that the projection $\pi : \mathcal{G}^{\bullet\bullet} \to \mathcal{G}^{\bullet}$ which forgets the second distinguished point is continuous with respect to the local distances.

A Borel function $f : \mathcal{G}^{\bullet\bullet} \to \mathbb{R}_+$ (thus invariant by homomorphism of doubly-pointed graph) is called a *transport function*: $f(\mathbf{g}, x, y)$ is interpreted as a quantity, a mass say, that the vertex x sends to the vertex y in the graph $\mathbf{g}$.

Definition 7 (Unimodular random graph)**.** *A random pointed graph* (G, ρ) *is* unimodular *if it obeys the Mass-Transport Principle (MTP) i.e. if for any transport function f we have*

$$\mathbb{E}\left[\sum_{x \in \mathrm{V}(G)} f(G, \rho, x)\right] = \mathbb{E}\left[\sum_{x \in \mathrm{V}(G)} f(G, x, \rho)\right]. \tag{2.1}$$

The preceding equation can be interpreted by saying that the average mass the reference point in G sends in total is equal to the average mass it receives from other vertices. The transport of mass using f is a fair game on average.

Remark 1. The terminology "unimodular" comes from group theory: If $\mathbf{g}$ is a graph we denote by Γ the group of all its automorphisms (homomorphisms from $\mathbf{g}$ to $\mathbf{g}$). When Γ is locally compact, we know by general theory that there exists a left-invariant measure on it (called the Haar measure). The graph $\mathbf{g}$ is unimodular if this left-invariant measure is also right-invariant.

As promised, the notion of unimodular random graph coincides with the notion of uniformly pointed random graph in the finite case:

Proposition 7. *A random finite pointed graph $G^{\bullet}$ is uniformly pointed if and only if it is unimodular.*

Proof. Suppose that $G^{\bullet}$ is uniformly pointed and let f be a transport function. Then noticing that $\sum_{x \in \mathrm{V}(\mathbf{g})} f(\mathbf{g}, \rho, x) =: F(\mathbf{g}, \rho)$ and $\sum_{x \in \mathrm{V}(\mathbf{g})} f(\mathbf{g}, x, \rho) =: F'(\mathbf{g}, \rho)$ are measurable functions for the single-pointed local topology we have

$$\mathbb{E}\left[\sum_{x \in \mathrm{V}(G)} f(G, \rho, x)\right] = \mathbb{E}[F(G, \rho)] = \mathbb{E}\left[\frac{1}{\#\mathrm{V}(G)} \sum_{x \in \mathrm{V}(G)} F(G, x)\right] = \mathbb{E}\left[\frac{1}{\#\mathrm{V}(G)} \sum_{x, y \in \mathrm{V}(G)} f(G, x, y)\right],$$

$$\mathbb{E}\left[\sum_{x \in \mathrm{V}(G)} f(G, x, \rho)\right] = \mathbb{E}[F'(G, \rho)] = \mathbb{E}\left[\frac{1}{\#\mathrm{V}(G)} \sum_{x \in \mathrm{V}(G)} F'(G, x)\right] = \mathbb{E}\left[\frac{1}{\#\mathrm{V}(G)} \sum_{x, y \in \mathrm{V}(G)} f(G, y, x)\right].$$

Since the right-most quantities are equal we conclude that $G^\bullet$ obeys the MTP. Conversely, if $G^\bullet$ is unimodular and almost surely finite, we choose a transport function of the form $f(\mathrm{g},x,y) = \frac{1}{\#\mathrm{V}(\mathrm{g})}h(\mathrm{g},x)$ where $h : \mathcal{G}^\bullet \to \mathbb{R}_+$ is a measurable function. We obtain by applying the MTP that

$$\mathbb{E}[h(G,\rho)] = \mathbb{E}\left[\sum_{x\in\mathrm{V}(G)} f(G,\rho,x)\right] = \mathbb{E}\left[\sum_{x\in\mathrm{V}(G)} f(G,x,\rho)\right] = \mathbb{E}\left[\frac{1}{\#\mathrm{V}(G)}\sum_{x\in\mathrm{V}(G)} h(G,x)\right],$$

which indeed shows that $G^\bullet$ is uniformly pointed. □

2.1.3 Mass-Transport principle and local limits

One nice feature with the notion of unimodularity is that it is preserved by taking local limit:

Theorem 8

Let $G_n^\bullet = (G_n, \rho_n)$ be a sequence of unimodular random graphs converging in distribution for $\mathrm{d_{loc}}$ towards $G_\infty^\bullet$. Then $G_\infty^\bullet = (G_\infty, \rho_\infty)$ is unimodular.

Proof. We start with a warmup. If f is a transport function with finite range, i.e. such that $f(\mathrm{g},x,y)$ is zero as soon as x and y are at at least distance r_0 and that $f(\mathrm{g},x,y)$ only depends on $[(\mathrm{g},x,y)]_{r_0}$ then it follows that for every $k \geq 0$ the functions

$$F_k(\mathrm{g},\rho) = \sum_{x\in\mathrm{V}(\mathrm{g})} \big(k \wedge f(\mathrm{g},\rho,x)\big) \mathbb{1}_{\#\mathrm{V}([(\mathrm{g},\rho,x)]_{r_0})\leq k} \text{ and } F_k'(\mathrm{g},\rho) = \sum_{x\in\mathrm{V}(\mathrm{g})} \big(k \wedge f(\mathrm{g},x,\rho)\big) \mathbb{1}_{\#\mathrm{V}([(\mathrm{g},x,\rho)]_{r_0})\leq k},$$

are both bounded continuous functions for the local topology. Hence, applying the mass-transport principle on $G_n^\bullet$ we have that

$$\mathbb{E}[F_k(G_n^\bullet)] = \mathbb{E}[F_k'(G_n^\bullet)].$$

By the local convergence of $G_n^\bullet$ to $G_\infty^\bullet$ we thus get that $\mathbb{E}[F_k(G_\infty^\bullet)] = \mathbb{E}[F_k'(G_\infty^\bullet)]$. Letting $k \to \infty$ we get by monotone convergence that

$$\mathbb{E}\left[\sum_{x\in\mathrm{V}(G_\infty)} f(G_\infty,\rho_\infty,x)\right] = \mathbb{E}\left[\sum_{x\in\mathrm{V}(G_\infty)} f(G_\infty,x,\rho_\infty)\right].$$

The mass-transport principle is thus satisfied for all transport functions depending only on a finite range around the first point. However, there are transport functions which are not simple function of that sort, for example consider $f(\mathrm{g},x,y) = \mathbb{1}_{x\sim y}\mathbb{1}_{\#\mathrm{V}(\mathrm{g})=\infty}$ for which the condition $\#\mathrm{V}(\mathrm{g}) = \infty$ is not continuous for the local topology! Proving the general result necessitates a bit of abstract measure theory. We proceed as follows. Let $r_0, k \geq 0$ and denote by

$$D_{r_0,k} = \{(\mathrm{g},x,y) : \mathrm{d_{gr}}(x,y) \leq r_0 \text{ and } \#\mathrm{V}([(\mathrm{g},x,y)]_{r_0}) \leq k\} \subset \mathcal{G}^{\bullet\bullet}.$$

We then introduce the family of measurable sets

$$\mathcal{M}_{r_0,k} = \left\{A \subset \mathcal{G}^{\bullet\bullet} \text{ measurable} : \mathbb{E}\left[\sum_{x\in\mathrm{V}(G_\infty)} \mathbb{1}_{(G_\infty,\rho_\infty,x)\in A\cap D_{r_0,k}}\right] = \mathbb{E}\left[\sum_{x\in\mathrm{V}(G_\infty)} \mathbb{1}_{(G_\infty,x,\rho_\infty)\in A\cap D_{r_0,k}}\right]\right\}.$$

By the above warmup, all the elementary sets $A = \{(\mathrm{g}, x, y) : [(\mathrm{g}, x, y)]_r = \mathrm{g}_0^{\bullet\bullet}\}$ when $\mathrm{g}_0^{\bullet\bullet} \in \mathcal{G}^{\bullet\bullet}$ is a finite bi-pointed graph are in $\mathcal{M}_{r_0,k}$ and those sets generate the Borel σ-field of $\mathcal{G}^{\bullet\bullet}$ and are stable under finite intersection. We claim that $\mathcal{M}_{r_0,k}$ is a monotone class: the stability under monotone union is clear using the monotone convergence theorem, and the stability under difference follows from the fact that

$$\mathbb{E}\left[\sum_{x \in \mathrm{V}(G_\infty)} \mathbb{1}_{(G_\infty,\rho_\infty,x)\in A^c\cap D_{r_0,k}}\right] = \mathbb{E}\left[\sum_{x \in \mathrm{V}(G_\infty)} \mathbb{1}_{(G_\infty,\rho_\infty,x)\in D_{r_0,k}}\right] - \mathbb{E}\left[\sum_{x \in \mathrm{V}(G_\infty)} \mathbb{1}_{(G_\infty,\rho_\infty,x)\in A\cap D_{r_0,k}}\right],$$

and similarly when the roles of ρ_∞ and x are exchanged. We notice that the first expectation in the right-hand side is finite (less than k from the definition of $D_{r_0,k}$). It follows that $\mathcal{M}_{r_0,k}$ is the Borel σ-field of $\mathcal{G}^{\bullet\bullet}$. Sending $r_0 \to \infty$ and $k \to \infty$, we deduce from the monotone convergence theorem that $G_\infty^\bullet$ obeys the mass-transport principle for any indicator function. Since positive functions are almost sure increasing limits of sum of indicator functions the full MTP follows from another application of the monotone convergence theorem. □

Exercise 12. Recall the notation $G^\bullet(n, p)$ for the connected component of 1 where the vertex 1 is distinguished in an Erdös–Rényi random graph over n vertices and parameter p. Show that $G^\bullet(n, p)$ is a unimodular random graph and deduce that Poisson-Galton–Watson trees (pointed at the ancestor) are also unimodular. Show that a supercritical Poisson-Galton–Watson trees (pointed at the ancestor) conditioned to be infinite is still unimodular.

We can now state perhaps the most interesting open problem in the field. Its resolution would imply quite a few famous conjectures in group theory as this would imply that every group is "sofic" .

Open Question 5 (Aldous–Lyons)**.** *Show that every unimodular random graph is a local limit in distribution of uniformly pointed random graphs.*

This conjecture has been proved in special cases for example when the limiting random graph does not grow too fast [4] or the case of unimodular random trees [19, 24, 13, 18].

2.2 Examples

We now give a few example of unimodular random graphs. The first one actually being a counter-example.

2.2.1 The grand-father graph

The following graph is an example of a vertex-transitive graph (recall Definition 2) which is not unimodular. We start with a k-regular tree t with $k \geq 3$ (in the following $k = 3$) given together with a distinguished infinite ray (i.e. an infinite self-avoiding path converging to a point on the boundary). This ray enables us to speak about a limit point "∞" at its extremity. Hence, if x and y are neighbors in the graph, then one of the two vertices x or y is closer to "∞" than the other. The furthest of the two points is called the parent of the other and this induces a genealogical order on the vertices of the tree. In the original graph t we then add all the edges linking a vertex to its four grand-parents (hence the name of the graph). The graph gf obtained is clearly vertex-transitive.

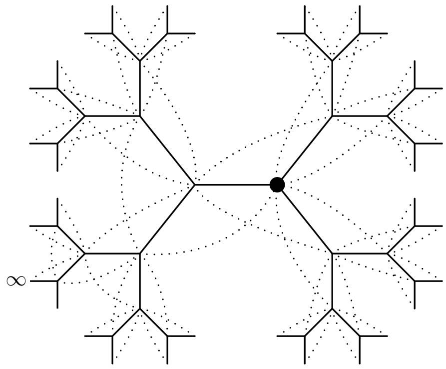

However, this graph (pointed at any vertex) is not unimodular: consider the transport function $f(\mathsf{g}, x, y) = 1$ if y is a parent of x and 0 otherwise. We let the reader check that this is indeed a transport function (this is not trivial, we have to show that the graph structure of gf enables to recover the distinguished ray hence the genealogical order). Then the MTP is violated since

$$2 = \mathbb{E}\left[\sum_{x \in \mathrm{V}(\mathsf{gf})} f(\mathsf{gf}, \rho, x)\right] \neq \mathbb{E}\left[\sum_{x \in \mathrm{V}(\mathsf{gf})} f(\mathsf{gf}, x, \rho)\right] = 1.$$

2.2.2 Cayley graphs

Recall the definition of Cayley graphs (Definition 1).

Proposition 9. *Any Cayley graph (pointed anywhere) is unimodular.*

Proof. Let g be the Cayley graph of (Gr, S) pointed at the identity e of the group. Then for any $x, y \in \mathsf{Gr}$ there is a homomorphism of bi-pointed graph $(\mathsf{g}, x, y) \to (\mathsf{g}, e, yx^{-1})$ where e is the identity of the group: this is the multiplication by x^{-1} on the right. Hence, since a transport function f is invariant under homomorphism of bi-pointed graphs, we have $f(\mathsf{g}, x, y) = \tilde{f}(yx^{-1})$ for some function $\tilde{f} : \mathsf{Gr} \to \mathbb{R}_+$. Hence we have

$$\mathbb{E}\left[\sum_{x \in \mathsf{Gr}} f(\mathsf{g}, e, x)\right] = \sum_{x \in \mathsf{Gr}} f(\mathsf{g}, e, x) = \sum_{x \in \mathsf{Gr}} \tilde{f}(x) = \sum_{x \in \mathsf{Gr}} \tilde{f}(x^{-1}) = \mathbb{E}\left[\sum_{x \in \mathsf{Gr}} f(\mathsf{g}, x, e)\right],$$

where we used the fact that $x \mapsto x^{-1}$ is an involution of the group Gr. □

As a consequence of the last proposition we deduce that although being vertex-transitive, the grand-father graph is not a Cayley graph of any group (do you see another way to prove it?).

Exercise 13. Show that $\mathbb{Z}^d$ for $d \geq 1$ with the standard edge set as well as infinite regular trees of even degree are unimodular.

Exercise 14. Show that the infinite three-regular tree although not a Cayley graph in the above sense (at least not for a symmetric set of generators) is nevertheless unimodular.

2.2.3 Construction from existing unimodular random graphs

We can obtain new unimodular random graphs from existing ones by modifications which do not depend on a base point. We present here the case of bond percolation but the result can be adapted to many other situations such as invariant percolation (see Section 2.3.2). Let $G^\bullet = (G, \rho)$ be a unimodular random graph and conditionally on $G^\bullet$ perform a bond percolation on G with parameter $p \in (0,1)$ (i.e. keep each edge independently with probability p). Denote by $\mathcal{C}^\bullet(\rho) = (\mathcal{C}(\rho), \rho)$ the cluster of the origin ρ pointed at ρ. Hence $\mathcal{C}^\bullet(\rho)$ is a random pointed graph.

Proposition 10. *The random graph $\mathcal{C}^\bullet(\rho)$ is unimodular.*

Proof. We directly verify the mass-transport principle (2.1): Fix a transport function f and compute

$$\mathbb{E}\left[\sum_{x \in \mathrm{V}(\mathcal{C}(\rho))} f(\mathcal{C}(\rho), \rho, x)\right] = \mathbb{E}\left[\sum_{x \in \mathrm{V}(G)} \mathrm{E}_G\left[f(\mathcal{C}(\rho), \rho, x) 1_{x \leftrightarrow \rho}\right]\right],$$

where E_{g} is the probability measure underlying a bond percolation on the graph g and $x \leftrightarrow y$ means that x and y are in the same cluster after performing the percolation. We just have to realize that $F(\mathrm{g}, x, y) = \mathrm{E}_{\mathrm{g}}\left[f(\mathcal{C}(x), x, y) 1_{x \leftrightarrow y}\right]$ is a transport function and so applying the MTP with the initial graph (G, ρ) we get that

$$\mathbb{E}\left[\sum_{x \in \mathrm{V}(G)} \mathrm{E}_G\left[f(\mathcal{C}(\rho), \rho, x) 1_{x \leftrightarrow \rho}\right]\right] = \mathbb{E}\left[\sum_{x \in \mathrm{V}(G)} \mathrm{E}_G\left[f(\mathcal{C}(x), x, \rho) 1_{\rho \leftrightarrow x}\right]\right].$$

Noting that on the event $\{\rho \leftrightarrow x\}$ the clusters $\mathcal{C}(\rho)$ and $\mathcal{C}(x)$ containing respectively x and ρ are the same it remains just to apply Fubini's theorem to arrive at the desired equality. □

Exercise 15. Let (G, ρ) be a unimodular random graph. Let $A \subset \mathcal{G}^\bullet$ be an invariant event in the sense that if $(\mathrm{g}, x) \in A$ then $(\mathrm{g}, y) \in A$ for any $y \in \mathrm{g}$. Give examples of such A and prove that if $\mathbb{P}((G, \rho) \in A) > 0$ then (G, ρ) conditioned on being in A is again unimodular.

2.3 A few applications

We now present a couple of results which hold for any unimodular random graphs. We will see indeed that these random graphs cannot behave too widely and should be thought of as the stochastic analog of "regular" or "homogeneous" graphs. We focus mainly on the concept of ends. Then we present the initial use of the Mass-Transport [27] which does not directly lie in our framework but illustrates the power of the technique. We start with a useful proposition:

Proposition 11 (Everything shows at the origin). *Let $G^\bullet$ be a unimodular random graph and $A \subset \mathcal{G}^\bullet$ be a Borel set such that $\mathbb{P}(G^\bullet \in A) = 0$. Then the probability that there exists a vertex $x \in \mathrm{V}(G)$ such that $(G, x) \in A$ is equal to zero.*

Proof. Simply consider the transport function $f(\mathrm{g}, x, y) = \mathbb{1}_{(\mathrm{g},x)\in A}$. The mass-transport principle entails that

$$0 = \mathbb{P}((G, \rho) \in A) = \mathbb{E}\left[\sum_{x\in\mathrm{V}(G)} f(G, \rho, x)\right] = \mathbb{E}\left[\sum_{x\in\mathrm{V}(G)} f(G, x, \rho)\right] = \mathbb{E}\left[\sum_{x\in\mathrm{V}(G)} \mathbb{1}_{(G,x)\in A}\right].$$

□

Proposition 12 (If it happens, it happens a lot). *Let (G, ρ) be a unimodular random graph which is almost surely infinite. Then for any $A \subset \mathcal{G}^\bullet$ Borel we have*

$$\#\{x \in \mathrm{V}(G) : (G, x) \in A\} \in \{0, \infty\} \quad a.s.$$

Proof. Fix a measurable subset $A \subset \mathcal{G}^\bullet$ and for $\mathrm{g} \in \mathcal{G}^\bullet$ denote by $A_g = \{x \in \mathrm{V}(\mathrm{g}) : (\mathrm{g}, x) \in A\}$. We then consider the transport function

$$f(x, y, \mathrm{g}) = \frac{1}{\#A_\mathrm{g}} \mathbb{1}_{y\in A_\mathrm{g}} \mathbb{1}_{0<\#A_\mathrm{g}<\infty}.$$

In other words, on the event when A_g is non-empty and finite, each vertex x splits a unit mass between all the vertices of A_g and otherwise does nothing. Applying the mass-transport principle (2.1) we deduce that

$$\mathbb{P}(\#A_G \in \{1, 2, ...\}) = \mathbb{E}[\infty \cdot \mathbb{1}_{(G,\rho)\in A} \mathbb{1}_{0<\#A_\mathrm{g}<\infty}].$$

Since the left-hand side is bounded by one we deduce that the event $\{(G, \rho) \in A$ and $0 < \#A_\mathrm{g} < \infty\}$ has zero probability. By the same display it follows that $\mathbb{P}(0 < \#A_G < \infty) = 0$ also. This is the desired statement. □

Heuristically speaking, when the random graph G possesses some vertex x such that seen from x the graph G has a certain property, then there are infinitely such vertices and even with "positive" density, whatever it means. Beware, this is not a $0-1$ law for the event in question: if the graph (G, ρ) is equal to $\mathbb{Z}^2$ with probability $1/2$ and to $\mathbb{Z}$ with probability $1/2$ and if A is the event $\{\deg(\rho) = 2\}$ then $\mathbb{P}(A) = 1/2$.

2.3.1 Ends of unimodular random graphs

Definition 1: Let g be a graph and $\mathrm{k}_1 \subset \mathrm{k}_2 \subset \dots$ an increasing sequence of finite subgraphs of g which exhausts g, that is $\bigcup_{i\geq 0} \mathrm{k}_i = \mathrm{g}$. An *end* of g is a nested sequence $\cdots \subset U_3 \subset U_2 \subset U_1$ where U_i is an infinite connected component of $\mathrm{g}\backslash\mathrm{k}_i$. A priori, the number of ends depends on the sequence $(\mathrm{k}_i)_{i\geq 1}$ however it is an exercise to see that it does not.

Definition 2: A ray in g is an infinite self-avoiding path (i.e. a sequence of distinct neighboring edges). We say that two rays r and r' are equivalent if there exists a third ray r'' which shares an infinite number of edges with both r and r'. This is an equivalence relation. The number of ends of g is the cardinality of the quotient space of the space of rays by the above equivalence relation.

Exercise 16. Show that the two notions coincide, that is give the same number of ends for every graph. Show that:

- any finite graph has 0 end,
- any infinite graph has at least one end,
- $\mathbb{Z}$ has two ends,
- $\mathbb{Z}^d$ for $d \geq 2$ has one end
- the complete k-ary tree with $k \geq 3$ has uncountably many ends.

There is also a natural way to put a topology on the space of ends of a graph to turn it into a topological space (see wikipedia "ends of graphs" for example). In particular, the space of ends of the full k-ary tree is homeomorphic to a Cantor set, i.e. a closed set with no interior and no isolated points.

Theorem 13 (***The degree tells us a lot***)

Let (G, ρ) be a unimodular random graph.

1. *If G is almost surely finite then we have $\mathbb{E}[\deg(\rho)] \geq 2 \cdot \mathbb{E}[1 - \#\mathrm{V}(G)^{-1}]$.*
2. *If G is almost surely infinite then $\mathbb{E}[\deg(\rho)] \geq 2$.*
3. *If G is almost surely infinite and $\mathbb{E}[\deg(\rho)] = 2$ then G is almost surely a tree with 1 or 2 ends.*

Proof. For the first point we consider the transport function $f(\mathrm{g}, x, y) = \deg(x)/\#\mathrm{V}(\mathrm{g})$. Applying (2.1) we get that

$$\mathbb{E}[\deg(\rho)] = \mathbb{E}\left[\sum_{x \in \mathrm{V}(G)} \frac{\deg(x)}{\#\mathrm{V}(G)}\right].$$

But in any graph the sum of all the degrees of the vertices is equal to twice the number of edges and since G is connected we have by Proposition 1 that $\#\mathrm{E}(G) \geq \#\mathrm{V}(G) - 1$. Combining these observations we get point 1.

For the second point we consider the transport function f defined by $f(\mathrm{g}, x, y) = 1$ whenever x and y are neighbors and linked by a single edge whose suppression isolates x in a finite connected component. We then apply the mass-transport principle (2.1) and split the cases depending on the degree of ρ. Remark first that ρ receives deterministically always strictly less than $\deg(\rho)$ unit of mass for otherwise g would be finite, furthermore:

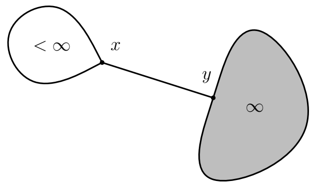

- when $\deg(\rho) = 1$ (we say that ρ is a leaf) then ρ sends a unit of mass to its unique neighbor and receives nothing,
- when $\deg(\rho) \geq 2$ and ρ happens to send a unit of mass through some edge then it cannot receive mass through this same edge and thus receives in total less than $\deg(\rho) - 1$ unit of mass,
- when $\deg(\rho) \geq 2$ and ρ sends no mass at all then after a few drawings, one can convince oneself that ρ receives less than $\deg(\rho) - 2$ unit of mass.

In all cases we have deterministically

$$\deg(\rho) + \mathrm{Sent}(\rho) - \mathrm{Received}(\rho) \geq 2.$$

The mass-transport principle (2.1) precisely says that the averages of the last two terms of the left-hand side cancel so that $\mathbb{E}[\deg(\rho)] \geq 2$ as desired. The last point of theorem corresponds to the saturation of the above inequalities. More precisely, if

$$F(\mathrm{g}, x) = \deg(x) + \sum_{y \in \mathrm{V}(\mathrm{g})} f(\mathrm{g}, x, y) - \sum_{y \in \mathrm{V}(\mathrm{g})} f(\mathrm{g}, y, x)$$

then we deterministically have $F(\mathrm{g}, x) \geq 2$ by the above point and thus $\mathbb{E}[\deg(\rho)] = 2$ implies that $F(G, \rho) = 2$ almost surely. By Proposition 11 we deduce that almost surely we have $F(G, x) = 2$ simultaneously for all $x \in \mathrm{V}(G)$. However it is easy to see that if there is a non-trivial cycle in the graph G then $F(G, x) > 2$ for any vertex x on this cycle. The graph is thus a tree. Also, when x is a point where at least three infinite paths merge then we have $F(\mathrm{g}, x) > 2$, hence there are no such points and the tree has 1 or 2 ends. □

Exercise 17. Recall that a vertex x in a graph g is a leaf if it has degree 1 (that is adjacent to a unique edge which is not a loop). Show that in a unimodular random graph (G, ρ) we have

$$\mathbb{P}(\rho \text{ is a leaf}) = \mathbb{E}\left[\#\text{leaves neighboring } \rho\right].$$

The next result is a generalization of a well-known theorem for Cayley graphs:

Theorem 14 (*Number of ends*)

The number of ends of a unimodular random graph necessarily belongs to $\{0, 1, 2, \infty\}$ *(but* *can be a random variable).*

Proof. We can suppose that (G, ρ) is almost surely infinite, otherwise consider (G, ρ) conditioned on the event $\#\mathrm{V}(G) = \infty$ which is still a unimodular random graph by Exercise 15. For $r \geq 0$ we call a vertex $x \in \mathrm{V}(\mathtt{g})$ an r-trifurcation if $\mathtt{g} \backslash [(\mathtt{g}, x)]_r$ contains at least 3 infinite connected components. It is an easy exercise to see that if a locally finite infinite graph has $k \in \{3, 4, \dots\}$ ends then there must exists a r-trifurcation for some $r \geq 1$. Also $\mathtt{g}$ has infinitely many ends if and only if

$$\#\{r\text{-trifurcations}\} \xrightarrow[r\to\infty]{} \infty.$$

However by Proposition 12 for any $r \geq 1$, in a unimodular infinite random graph (G, ρ) the number of r-trifurcations is either 0 or ∞ almost surely. We deduce from the last two geometric remarks that the number of ends of (G, ρ) is necessarily in $\{1, 2, \infty\}$. □

2.3.2 Invariant percolation on trees

Let $\mathtt{t}^\bullet = (\mathtt{t}, \rho)$ be the four-regular infinite tree pointed at any vertex. An invariant (site) percolation on $\mathtt{t}$ is a probability measure on $\Omega : \{0,1\}^{\mathrm{V}(\mathtt{t})}$ which is invariant under any graph homomorphism of $\mathtt{t}$, in words, it is a random bicoloring of the vertices $v \in \mathrm{V}(\mathtt{t})$ in black if $\omega(v) = 1$ and white if $\omega(v) = 0$ such that the law of coloring does not depend on the distinguished pointed vertex. For example a Bernoulli site percolation would do the job, but the interesting case is that this percolation may exhibit dependence between sites! By invariance, the percolation density

$$p = \mathbb{E}[\omega(v)] \in [0, 1],$$

does not depend on $v \in \mathrm{V}(\mathtt{t})$. The result we want to prove is the following:

Theorem 15 (*High density automatically implies percolation*)

If p is close enough to 1 then there exists an infinite cluster with positive probability.

This theorem may seem awkward at first glance. To understand its power we will construct a counterexample on $\mathbb{Z}^2$ that is an invariant dependent percolation with density as close to 1 as possible but with no infinite cluster. To do this, imagine that we tile $\mathbb{Z}^2$ with squares of size $n \times n$ and color the vertices on the boundary of these squares in black (there are $4n - 4$ such vertices per big square) and all the others in white (there are $n^2 - 4n + 4$ such vertices per square). This gives a periodic coloring of $\mathbb{Z}^2$. To transform it into an invariant percolation we just need to pick the distinguished vertex ρ of $\mathbb{Z}^2$ uniformly at random inside a fixed big square.

Exercise 18. Check that the resulting random configuration of colors on $(\mathbb{Z}^2, \rho)$ is invariant in law under all graph homomorphisms of $\mathbb{Z}^2$.

Clearly the resulting percolation has no infinite cluster, but the density of the percolation is

$$\frac{n^2 - 4n + 4}{n^2} \xrightarrow[n\to\infty]{} 1.$$

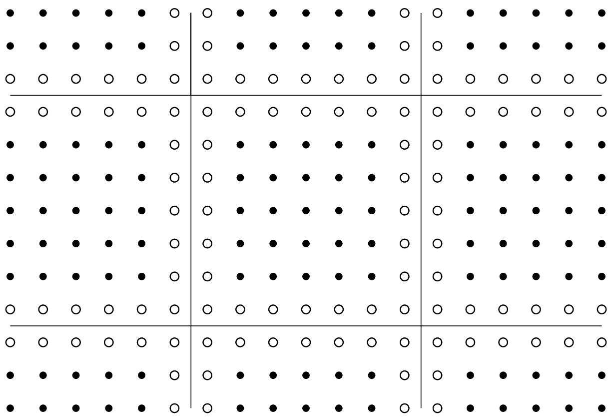

Hiding behind this construction is the fact that in $\mathbb{Z}^2$ there are large sets whose boundary size is negligible with respect to their size. This fact is not true in regular trees (or more generally in non-amenable graphs) and we can check that for any connected set A of vertices in $\mathtt{t}$, if we denote by ∂A the set of vertices adjacent to A but not in A then we have for some constant $c > 0$ independent of A,

$$\frac{\#\partial A}{\#A} \geq c. \tag{2.2}$$

Proof of Theorem 15. Given the tree $\mathtt{t}$ and a percolation ω on it, we will define a transport function f_ω on the vertices of $\mathtt{t}$ as follows: Recall that $\mathcal{C}(x)$ is the cluster of x then for all $y \in \mathrm{V}(\mathtt{t})$ we put

$$\begin{cases} \text{if } \omega(x) = 0 & \text{then} \quad f_\omega(x,y) = 0 \\ \text{if } \omega(x) = 1 \text{ and } \#\mathcal{C}(x) = \infty & \text{then} \quad f_\omega(x,y) = 0 \\ \text{if } \omega(x) = 1 \text{ and } \#\mathcal{C}(x) < \infty & \text{then} \quad f_\omega(x,y) = \frac{1}{|\partial\mathcal{C}(x)|}\mathbb{1}_{y\in\partial\mathcal{C}(x)}. \end{cases}$$

In words, if $\omega(x) = 0$ or if x is in an infinite cluster then x sends no mass at all. Otherwise, if x is in a finite connected component for ω he sends mass 1 which is spread over all the neighbor vertices of $\mathcal{C}(x)$. Recall that $\mathtt{t}$ can be seen as the Cayley graph of the free group over two elements (see Exercise 13) we deduce that $\mathbb{E}[f_\omega(x,y)]$ in fact only depends on xy^{-1} and by invariance of the percolation and involution invariance of the group we get that (exercise!)

$$\mathbb{E}\left[\sum_{x\in\mathrm{V}(\mathtt{t})} f_\omega(\rho,x)\right] = \mathbb{E}\left[\sum_{x\in\mathrm{V}(\mathtt{t})} f_\omega(x,\rho)\right]. \tag{2.3}$$

This is the version of (2.1) that we will use. On the left-hand side of the last display we have $\mathbb{P}(\omega(x) = 1 \text{ and } \#\mathcal{C}(x) < \infty)$. We will now bound the right-hand side. Remark first that to receive mass, the vertex ρ must be white $\omega(\rho) = 0$ and must lie on the boundary of a finite cluster. Deterministically ρ can be on the boundary of at most 4 finite clusters A_1, A_2, A_3 and A_4 and all the vertices in A_i sends a mass to ρ equal to

$$\frac{\#A_i}{\#\partial A_i}.$$

By (2.2) and (2.3) we thus deduce that

$$\mathbb{P}(\omega(\rho) = 1 \text{ and } \#\mathcal{C}(\rho) < \infty) \leq \frac{4}{c} \cdot \mathbb{P}(\omega(\rho) = 0).$$

If the density $p = \mathbb{P}(\omega(\rho) = 1)$ is close enough to 1 then the left-hand side becomes smaller than p and thus ρ is in an infinite cluster with positive probability as desired. □

2.4 Abstracting the setup: Measured equivalence relations

We developed in this chapter the notion of unimodular random graph in the context of pointed random graphs. Obviously we could abstract the setup: We only needed a notion of structure (here graphs) with a base point (here a distinguished vertex) which we then want to pick at random. More precisely recall the notation of Section 1.2.1: we have an abstract metric space $(\mathcal{E}, \delta)$ and suppose that we have an equivalence relation $\sim$ on $\mathcal{E}$ so that the equivalence class of x is denoted by x. Suppose also that for each class x we have a measure μ_{x}, non necessarily finite, on the equivalence class x. Provided that some compatibility conditions are required on the local topology and the measures we will say that a random variable X with values in $(\mathcal{E}, \mathrm{d_{doc}})$ is uniformly based if μ_{X} is almost surely finite and

$$\mathbb{E}[f(X)] = \mathbb{E}\left[\frac{1}{\mu_{\mathrm{X}}(\mathcal{E})} \int \mathrm{d}\mu_{\mathrm{X}}(x) f(x)\right].$$

Similarly, when the measure μ can be infinite we can speak of a unimodular random variable X if we have

$$\mathbb{E}\left[\int \mathrm{d}\mu_{\mathrm{X}}(y) f(X, y)\right] = \mathbb{E}\left[\int \mathrm{d}\mu_{\mathrm{X}}(x) f(y, X)\right],$$

for every function $f : \mathcal{E} \times \mathcal{E} \to \mathbb{R}_+$. This concept is (related to) the theory of "measured equivalence relations". Let us see how to apply this setup in the above setting: We were dealing with random pointed graphs $\mathcal{E} = \mathcal{G}^\bullet$ and the equivalence relation is defined by $(\mathbf{g}, \rho) \sim (\mathbf{g}', \rho')$ if and only if $(\mathbf{g}', \rho')$ is obtained from $\mathbf{g}$ by changing the origin point. In this case $\mu_{\mathrm{g}^\bullet}$ is morally obtained as the counting measure obtained by changing the base point

$$\mu_{\mathrm{g}^\bullet} = \sum_{x \in \mathrm{V}(g)} \delta_{(g, x)}.$$

There is however a difficulty here in the presence of symmetries in the graph. Indeed, think of $\mathbf{g}$ as been a vertex-transitive graph, then its equivalence class g contains a unique point and the transport $f(\mathbf{g}, x, y)$ for two vertices $x, y \in \mathrm{V}(g)$ cannot be seen as a $f((\mathbf{g}, x), (\mathbf{g}, y))$ for a function $f : \mathcal{G}^\bullet \times \mathcal{G}^\bullet \to \mathbb{R}_+$! There are two multiple ways to bypass this difficulty: considering *rigid* graphs (that are graphs without non trivial isomorphisms see [30, Section 1E]) or add independent uniform labels $\in [0, 1]$ on the graphs (see [4, Example 9.9]) so that we are dealing pointed label graphs which now have also no non trivial isomorphisms.

Biliographical references. Olle Haggstrom [27] first used a form of the Mass-transport principle to study (dependent) invariant percolation on regular trees and show Theorem 15. This triggered the systematic study of percolation on Cayley graph using the mass-transport principle as done by Benjamini, Lyons, Peres and Schramm [12]. The most general form of unimodular random graph and the MTP (as well as the local topology) have been introduced by Benjamini and Schramm [14]. Most of this chapter is adapted from the wonderful survey paper [4] which regroups and extends the results of [12] to random graphs. We refer to [4] for original references and pointers. The Open Question 5 is due to Aldous & Lyons [4]. The answer is known to be true in the case of trees [24, 13]. The reader interested in measured equivalence relations should consult [4, 30, 40]

Chapter 3

The random walk point of view

In this chapter we develop another point of view on unimodular random graphs. Roughly speaking we will show that unimodular random graphs are those random graphs where the landscape viewed from a particle performing a simple random walk is stationary in distribution.

3.1 Stationary (and reversible) random graphs

A rooted graph is a pair $\vec{\mathrm{g}} = (\mathrm{g}, \vec{e})$ where g is a (locally finite, connected) graph and $\vec{e}$ is a distinguished oriented edge (that is given with a direction) that we call the root edge. If $\mathrm{g} = (\mathrm{V}(\mathrm{g}), \mathrm{E}(\mathrm{g}))$ is a graph, we denote by $\vec{\mathrm{E}}(\mathrm{g})$ the set of all oriented edges of the graph, which is obtained informally by duplicating each non-oriented edge (including self-loops) into two oriented edges. If $\vec{e} \in \vec{\mathrm{E}}(\mathrm{g})$ we denote by $\vec{e}_*$ the origin vertex of the oriented edge. Two rooted graphs $\vec{\mathrm{g}}_1 = (\mathrm{g}_1, \vec{e}_1)$ et $\vec{\mathrm{g}}_2 = (\mathrm{g}_2, \vec{e}_2)$ are equivalent if there exists a graph homomorphism $\mathrm{g}_1 \to \mathrm{g}_2$ which sends $\vec{e}_1$ onto $\vec{e}_2$. As usual we will implicitly identify such graphs and work on $\vec{\mathcal{G}}$, the space of equivalence classes of (locally finite, connected) rooted graphs. For $r \geq 0$, the restriction of radius r in $\vec{\mathrm{g}}$, denoted by $[\vec{\mathrm{g}}]_r$, is obtained by keeping only those vertices and edges of g which are at distance less than or equal to r from the origin of the root edge and keeping the root edge as the distinguished oriented edge. As previously, we consider the local distance $\mathrm{d}_{\mathrm{loc}}$ on $\vec{\mathcal{G}}$ defined by the procedure of Section 1.2.1. If $\vec{\mathrm{g}} = (\mathrm{g}, \vec{e})$ is a rooted graph, we denote by $\pi_\bullet(\vec{\mathrm{g}})$ the pointed graph obtained by distinguished in g the origin vertex of the root edge. Conversely, if $G^\bullet = (G, \rho)$ is a deterministic or random pointed graph, we denote by $\pi_\to(G^\bullet)$ the *random* rooted graph obtained by distinguishing an oriented edge emanating from ρ uniformly at random, conditionally on $G^\bullet$. Although $\pi_\to$ contains some additional randomness, the mapping $\pi_\bullet \circ \pi_\to$ is the identity on $\mathcal{G}^\bullet$.

Exercise 19. Show that $\pi_\bullet : \vec{\mathcal{G}} \to \mathcal{G}^\bullet$ is continuous.

3.1.1 Uniformly rooted random graphs

We start with the exact analog of Definition 4 in the context of rooted graphs:

Definition 8. *Let $\vec{G} = (G, \vec{e})$ an almost surely finite random rooted graph. We say that $\vec{G}$ is uniformly rooted if the root edge is chosen uniformly on $\vec{\mathrm{E}}(G)$ or more precisely if for all Borel function $f : \vec{\mathcal{G}} \to \mathbb{R}_+$ we have*

$$\mathbb{E}[f(\vec{G})] = \mathbb{E}\left[\frac{1}{\#\vec{\mathrm{E}}(G)} \sum_{\vec{\sigma} \in \vec{\mathrm{E}}(G)} f(G, \vec{\sigma})\right].$$

We now explain how to pass from a uniformly rooted graph to a uniformly pointed graph and vice-versa. This uses the concept of biased random variable. Recall that if X, Y are two random variables defined on a common probability space such that X takes its values in some abstract space (E, d) and such that Y is positive real-valued then we can construct a new random variable $\tilde{X}$ whose law is the law of *the random variable X biased by Y* characterized by

$$\mathbb{E}[f(\tilde{X})] = \frac{1}{\mathbb{E}[Y]}\mathbb{E}[f(X) \cdot Y],$$

for every Borel function $f : E \to \mathbb{R}_+$. Of course, this definition requires that Y has a finite non zero expectation. Equivalently, the law of $\tilde{X}$ is the distribution with Radon–Nikodym derivative equal to $Y(\omega)/\mathbb{E}[Y]$ with respect to the law of X.

Exercise 20. Let $U \in [0, 1]$ be a uniform random variable. Compute the law of U biased by itself.

Exercise 21. Let $(X_i)_{i \in I}$ be a family of positive real-valued random variables all of mean 1. For each $i \in I$ denote by $\tilde{X}_i$ a random variable with the law of the variable X_i biased by itself. Prove that $(\tilde{X}_i)_{i \in I}$ is tight if and only if $(X_i)_{i \in I}$ is uniformly integrable.

Exercise 22 (Change of measure via a martingale). Let $(X_n)_{n \geq 0}$ be a simple symmetric random walk on $\mathbb{Z}$ starting from $X_0 = 1$. We denote by τ the first hitting time of 0 by X . Everyone knows that $(X_{n \wedge \tau})_{n \geq 0}$ is a martingale for the canonical filtration. For every $n \geq 0$ we denote by $(\tilde{X}_k^{(n)})_{k \leq n}$ the law of $(X_k)_{k \leq n}$ biased by $X_{n \wedge \tau}$.

1. Show that for any $0 \leq m \leq n$ the law of $(\tilde{X}_k^{(n)})_{k \leq m}$ does not depend on n. Hence, by coherence we can define the law of $(\tilde{X}_k)_{k \geq 0}$.

2. Show that $(\tilde{X}_k)_{k \geq 0}$ is a Markov chain whose probability transitions are $q(i, i+1) = (i+1)/(2i)$ and $p(i, i-1) = (i-1)/(2i)$ for $i \geq 1$. (For those who know, recognize that $\tilde{X}$ has the law of the h-transform of $(X_n)_{n \geq 0}$ with the harmonic function $h(i) = i$)

In our case, if $\vec{G}$ is a random uniformly rooted graph then the graph $\underline{G}^{\bullet}$ obtained from $\pi_{\bullet}(\vec{G})$ by biasing by $\deg(\vec{e}_*)^{-1}$ is a random uniformly pointed graph (we denoted by $\vec{e}_*$ the origin vertex of the root edge $\vec{e}$). Indeed, if f is a positive Borel function on $\mathcal{G}\bullet$ we have

$$\begin{aligned}
\mathbb{E}[f(\underline{G}^\bullet)] &= \frac{1}{\mathbb{E}[\deg^{-1}(\vec{e}_*)]}\mathbb{E}\left[\deg^{-1}(\vec{e}_*)f(\pi_\bullet(\vec{G}))\right] \\
&\underset{\text{def.8}}{=} \frac{1}{\mathbb{E}[\deg^{-1}(\vec{e}_*)]}\mathbb{E}\left[\frac{1}{\#\vec{E}(G)}\sum_{\vec{\sigma}\in\vec{E}(G)}\deg^{-1}(\vec{\sigma}_*)f(G,\vec{\sigma}_*)\right] \\
&= \frac{1}{\mathbb{E}[\deg^{-1}(\vec{e}_*)]}\mathbb{E}\left[\frac{1}{\#\vec{E}(G)}\sum_{x\in\mathrm{V}(G)}f(G,x)\right] \\
&= \frac{1}{\mathbb{E}[\deg^{-1}(\vec{e}_*)]}\mathbb{E}\left[\frac{1}{\#\vec{E}(G)}\sum_{\vec{\sigma}\in\vec{E}(G)}\deg^{-1}(\vec{\sigma}_*)\frac{1}{\#\mathrm{V}(G)}\sum_{x\in\mathrm{V}(G)}f(G,x)\right].
\end{aligned}$$

The function $F(G,\vec{\sigma}) = \deg^{-1}(\vec{\sigma}_*)\frac{1}{\#\mathrm{V}(G)}\sum_{x\in\mathrm{V}(G)}f(G,x)$ is measurable function from $\vec{\mathcal{G}}$ to $\mathbb{R}_+$ hence by Definition 8 we get that the last chain of equalities goes on with

$$\frac{\mathbb{E}\left[F(G,\vec{e})\right]}{\mathbb{E}[\deg^{-1}(\vec{e}_*)]} = \frac{1}{\mathbb{E}[\deg^{-1}(\vec{e}_*)]}\mathbb{E}\left[\deg^{-1}(\vec{e}_*)\frac{1}{\#\mathrm{V}(G)}\sum_{x\in\mathrm{V}(G)}f(G,x)\right] = \mathbb{E}\left[\frac{1}{\#\mathrm{V}(\underline{G})}\sum_{x\in\mathrm{V}(\underline{G})}f(\underline{G},x)\right].$$

Exercise 23. Show that conversely, if $G^\bullet$ is a random uniformly pointed graph with $\mathbb{E}[\deg(\rho)] < \infty$, then the random rooted graph $\pi_\rightarrow(G^\bullet)$ biased by $\deg(\rho)$ is uniformly rooted.

Exercise 24. Let $G^\bullet$ be a uniformly pointed graph. Show that the rooted random graph $\pi_\rightarrow(G^\bullet)$ biased by $\deg(\rho)/\#\vec{E}(G)$ is uniformly rooted. Note that $\mathbb{E}[\deg(\rho)/\#\vec{E}(G)]$ is always finite and the above procedure may be used even if $\mathbb{E}[\deg(\rho)] = \infty$ as long as the graph is almost surely finite.

3.1.2 Invariance along the random walk

We face the same problem as in Section 2.1.1: the notion of uniformly rooted graph cannot trivially be extended to infinite graphs because of the presence of $\frac{1}{\#\overline{\mathrm{E}}(G)}$. We could do the same trick as in the last chapter and define "edge-unimodular" random graphs as those which satisfy the "edge-mass transport principle":

$$\mathbb{E}\left[\sum_{\vec{\sigma}\in\overline{\mathrm{E}}(G)}f(G,\vec{e},\vec{\sigma})\right] = \mathbb{E}\left[\sum_{\vec{\sigma}\in\overline{\mathrm{E}}(G)}f(G,\vec{\sigma},\vec{e})\right],$$

for any transport function f which associates a mass to any bi-rooted graph $(g,\vec{e},\vec{\sigma})$. We will however following a different concept yielding to a more general concept.

If $\vec{g} = (g,\vec{e})$ is a fixed rooted graph, we denote by $\mathrm{P}_{\vec{g}}$ the law of the simple random walk on g starting from the target of the root edge in g. More precisely this yields a sequence $\vec{E}_0, \vec{E}_1, \dots$ of oriented edges where $\vec{E}_0 = \vec{e}$ and recursively for $i \geq 0$ we choose independently of the past the next oriented edge $\vec{E}_{i+1}$ uniformly among all the oriented edges emanating from the target vertex of $\vec{E}_i$. If $\vec{G} = (G, \vec{e})$ is a random rooted graph the random walk on $\vec{G}$ is the law of $(G, (\vec{E}_i)_{i\geq 0})$ under the probability

$$\int \mathrm{d}\mathbb{P}(\vec{G}) \int \mathrm{dP}_{\vec{G}}((\vec{E}_i)_{i\geq 0}).$$

Exercise 25. Let $\vec{G}_n$ be a sequence of random rooted graphs converging in distribution for the local distance towards $\vec{G}_\infty$. If conditionally on $\vec{G}_n$ we denote by $(\vec{E}_i^{(n)})_{i\geq 0}$ the oriented edges traversed by a simple random walk on $\vec{G}_n$ show that for any $k \geq 0$ we have

$$(G_n, \vec{E}_k^{(n)}) \xrightarrow[n\to\infty]{(d)} (G_\infty, \vec{E}_k^{(\infty)}),$$

where $(\vec{E}_i^{(\infty)})_{i\geq 0}$ is a simple random walk on $\vec{G}_\infty$ (Hint: use Skorokhod's representation theorem to assume that $\vec{G}_n$ converge almost surely towards $\vec{G}_\infty$).

Definition 9. *Let $\vec{G} = (G, \vec{e})$ be a random rooted graph and denote by $(\vec{E}_i)_{i\geq 0}$ the sequence of edges visited by a simple random walk on it. We say that $\vec{G}$ is* stationary *(or* stationary *along simple random walk) if for every $k \geq 0$ the law of $(G, \vec{E}_k)$ is the same as that of $\vec{G}$. It is furthermore* reversible *if on top of it we have $(G, \vec{E}_0) = (G, \overleftarrow{E}_0)$ in law, where $\overleftarrow{e}$ is the edge $\vec{e}$ with reversed orientation.*

In the context of random walk in random environment (here the random environment is the underlying random graph) we often speak of a stationary environment seen from the particle. In other words, when the random walk displaces, it sees at each step the same surrounding random graph in distribution.

We begin with a few elementary remarks:

- By an easy induction we deduce that $\vec{G}$ is stationary if and only if $(G, \vec{E}_1) = (G, \vec{E}_0)$ in distribution.

- In the definition of reversibility, we first ask for stationarity: there are examples of random rooted graphs such that $(G, \vec{E}_0) = (G, \overleftarrow{E}_0)$ in distribution but which are not stationary (for example consider two copies of $\mathbb{Z}$ joined by an oriented edge).

- For a stationary and reversible random graph we get that $(G, \vec{E}_0) = (G, \vec{E}_k) = (G, \overleftarrow{E}_k)$ for every k. Also conditionally on the graph, the probability that $\vec{E}_0 = \vec{e}_0, \ldots, \vec{E}_k = \vec{e}_k$ for a fixed path $\gamma = (\vec{e}_0, \ldots, \vec{e}_k)$ is given by

  $$\left(\prod_{i=1}^{k} \deg((\vec{e}_i)_*) \right)^{-1}.$$

  Combining this with the above remark we deduce that for a stationary and reversible random graph we have

  $$(G, \vec{E}_0, \ldots, \vec{E}_k) \overset{(d)}{=} (G, \overleftarrow{E}_k, \ldots, \overleftarrow{E}_0), \tag{3.1}$$

  or in words that the first k steps of a random walk have the same law seen from either tip on a stationary and reversible random graph.

- We will often use that for a stationary random rooted graph $\vec{G} = (G, \vec{E}_0)$ we have

$$\pi_{\rightarrow} \circ \pi_{\bullet}(\vec{G}) = \vec{G}, \quad \text{in distribution.} \tag{3.2}$$

  In words, this means that we can start the simple random walk from the origin of the root edge instead of the target of the root edge and get the same distribution. To see this, denote $(\vec{E}_0, \vec{E}_1)$ and $(\vec{E}_0, \vec{E}_1')$ two independent one-step simple random walks on $\vec{G}$. Since these two walks start with the same oriented edges we clearly have

$$\pi_{\rightarrow} \circ \pi_{\bullet}((G, \vec{E}_1)) \overset{(d)}{=} (G, \vec{E}_1').$$

  But by stationarity we also have $(G, \vec{E}_1') = (G, \vec{E}_1) = \vec{G}$ in distribution. Combining the two statements we indeed get (3.2).

As promised, in the case of almost surely finite random graphs, stationarity along the simple random walk is equivalent to uniform rooting.

Proposition 16. *Let $\vec{G}$ be an almost surely finite random rooted graph. Then $\vec{G}$ is uniformly rooted if and only if $\vec{G}$ is stationary along the simple random walk.*

Proof. We first recall a well-known fact. If g is a fixed finite connected graph then the invariant probability measure for the oriented edges visited by simple random walk is nothing but the uniform measure on $\vec{\mathrm{E}}(\mathrm{g})$. In other words, if we pick a uniform oriented edge $\vec{E}_0$ of g and then perform $k \geq 0$ steps of random walk starting from the extremity of $\vec{E}_0$ then the distribution of the last oriented edge visited is again uniform over $\vec{\mathrm{E}}(\mathrm{g})$. It easily follows from this observation that if $(\vec{E}_k)_{k\geq 0}$ is a simple random walk on a random uniformly rooted graph $\vec{G}$ then $(G, \vec{E}_k) = (G, \vec{E}_0)$ in distribution as desired.

To show the converse we again recall that if $\vec{\mathrm{g}}$ is a finite connected rooted graph and $(\vec{E}_i)_{i\geq 0}$ has law $\mathrm{P}_{\vec{\mathrm{g}}}$ then by the classical ergodic theorem for recurrent Markov chains on a finite state space we have the almost sure weak convergence

$$\frac{1}{n}\sum_{k=0}^{n-1} \delta_{\vec{E}_k} \xrightarrow[n\to\infty]{a.s.} \frac{1}{\#\vec{\mathrm{E}}(\mathrm{g})} \sum_{\vec{\sigma}\in\vec{\mathrm{E}}(\mathrm{g})} \delta_{\vec{\sigma}}.$$

Then, if $f : \vec{\mathcal{G}} \to \mathbb{R}_+$ is a bounded Borel function then we have

$$S_f(\vec{\mathrm{g}}, n) = \frac{1}{n}\sum_{k=0}^{n-1} f(\mathrm{g}, \vec{E}_k) \xrightarrow[n\to\infty]{\mathrm{P}_{\vec{\mathrm{g}}}-a.s.} \frac{1}{\#\vec{\mathrm{E}}(\mathrm{g})} \sum_{\vec{\sigma}\in\vec{\mathrm{E}}(\mathrm{g})} f(\mathrm{g}, \vec{\sigma}) = U_f(\vec{\mathrm{g}}).$$

If now $\vec{G}$ is a finite stationary random graph then $\int \mathrm{d}\mathbb{P}(\vec{G}) \int \mathrm{dP}_{\vec{G}}(\vec{E}_k) f(G, \vec{E}_k) = \mathbb{E}[f(\vec{G})]$ so

$$\mathbb{E}[f(\vec{G})] \underset{\text{stationarity}}{=} \mathbb{E}\left[\int \mathrm{dP}_{\vec{G}} S_f(\vec{G}, n)\right] \xrightarrow[n\to\infty]{\text{dom. conv.}} \mathbb{E}[U_f(\vec{G})],$$

which proves that $\vec{G}$ is indeed uniformly rooted. □

Notice that a uniformly rooted random graph is automatically reversible and by the above result any stationary random graph which is almost surely finite is automatically reversible. However, this is not always the case in the infinite setting: Recall the grand-father graph gf of Section 2.2.1. Since gf is transitive, it is easy to see that $\pi_{\to}(\mathsf{gf})$ is stationary. However, $\pi_{\to}(\mathsf{gf})$ is not reversible since the root edge is more likely to link a vertex to one of its two grand-fathers (probability 1/3) rather than to its unique grand-son (probability 1/6).

We have the analog of Theorem 8:

Theorem 17

Let $\vec{G}_n$ be a sequence of stationary (resp. stationary and reversible) random graphs which converges locally in distribution towards $\vec{G}_\infty$. Then $\vec{G}_\infty$ is also stationary (resp. stationary and reversible).

Proof. Let $A_r \subset \vec{\mathcal{G}}$ be a Borel subset of rooted graphs such that $\vec{\mathfrak{g}} \in A_r$ only depends on the restriction of radius r around the root edge of $\vec{\mathfrak{g}}$. Then by stationarity of the graphs $\vec{G}_n$ we get that for any $k \geq 0$ we have

$$\mathbb{E}\left[\mathrm{P}_{\vec{G}_n}((G_n, \vec{E}_k^{(n)}) \in A_r)\right] = \mathbb{E}\left[\mathrm{P}_{\vec{G}_n}((G_n, \vec{E}_0^{(n)}) \in A_r)\right]$$

where $\vec{E}_i^{(n)}$ are the oriented edges traversed by the simple random walk on $\vec{G}_n$. By our assumption, the function 1_{A_r} is a bounded continuous function for the local topology hence using Exercise 25 we can pass to the limit in the last display and deduce that

$$\mathbb{E}\left[\mathrm{P}_{\vec{G}_\infty}((G_\infty, \vec{E}_k^{(\infty)}) \in A_r)\right] = \mathbb{E}\left[\mathrm{P}_{\vec{G}_\infty}((G_\infty, \vec{E}_0^{(\infty)}) \in A_r)\right]$$

with obvious notation. Using an easy adaptation of Proposition 4 to rooted graphs we conclude that $(G_\infty, \vec{E}_0^{(\infty)}) = (G_\infty, \vec{E}_k^{(\infty)})$ in distribution. The case of stationary and reversible random graphs is treated similarly. □

The analog of the Open question 5 also holds in this context: is any stationary and reversible random graphs a local limit of uniformly rooted random graphs (this is not true for stationary random graphs only as we have seen above that there are stationary infinite random graphs which are not reversible)?

3.1.3 Involution invariance

We now make the connection between stationary and reversible random rooted graphs and unimodular random pointed graphs directly in the (possibly) infinite case. We start with the easier direction: from unimodular random graphs to stationary and reversible random graphs.

Proposition 18. *Let $G^\bullet = (G, \rho)$ be a unimodular random pointed graph. Consider first the random pointed graph $\overline{G}^\bullet = (\overline{G}, \overline{\rho})$ obtained from (G, ρ) after biasing by the degree of its origin (provided that $\mathbb{E}[\deg(\rho)] < \infty$) and then introduce $\vec{\overline{G}} = \pi_{\to}(\overline{G}^\bullet)$ obtained by distinguishing an oriented edge emanating from $\overline{\rho}$ uniformly at random. Then $\vec{\overline{G}}$ is stationary and reversible.*

Proof. We will first show that $\vec{\overline{G}} = (\overline{G}, \vec{E})$ has the same law as $(\overline{G}, \overleftarrow{E})$. For this, let $h(\mathrm{g}, \vec{e})$ be a function $\vec{\mathcal{G}} \to \mathbb{R}_+$ and denote by

$$f(\mathrm{g}, x, y) = 1_{x \sim y} \sum_{x \underset{\vec{e}}{\to} y} h(\mathrm{g}, \vec{e}), \tag{3.3}$$

the associated transport function obtained by summing over all choice of an oriented edge linking x to y in g (in particular x, y must be neighbors). Then applying the mass-transport principle we get that

$$\begin{array}{ccccc} \mathbb{E}[\deg(\rho)] \cdot \mathbb{E}\left[h(\overline{G}, \vec{E})\right] = \mathbb{E}\left[\deg(\rho) \dfrac{1}{\deg(\rho)} \displaystyle\sum_{\vec{e} \text{ s.t. } \vec{e}_* = \rho} h(G, \vec{e})\right] & = & \mathbb{E}\left[\displaystyle\sum_{x \in \mathrm{V}(G)} f(G, \rho, x)\right] \\ & & \| \\ \mathbb{E}[\deg(\rho)] \cdot \mathbb{E}\left[h(\overline{G}, \overleftarrow{E})\right] = \mathbb{E}\left[\deg(\rho) \dfrac{1}{\deg(\rho)} \displaystyle\sum_{\vec{e} \text{ s.t. } \vec{e}_* = \rho} h(G, \overleftarrow{e})\right] & = & \mathbb{E}\left[\displaystyle\sum_{x \in \mathrm{V}(G)} f(G, x, \rho)\right] \end{array}$$

But in our context this statement also implies stationary. Indeed if $(\vec{E}_0 = \overleftarrow{E}, \vec{E}_1)$ are the first two steps of a random walk on $(\overline{G}, \overleftarrow{E})$ (started from the extremity of the root edge) then it is clear that $(\overline{G}, \vec{E}_1)$ has the same distribution as $(\overline{G}, \vec{E})$ since $\vec{E}_1$ is nothing but a uniform edge emanating from $\overline{\rho}$. Since $(\overline{G}, \overleftarrow{E})$ has the same law as $(\overline{G}, \vec{E})$ we indeed deduce that $(\overline{G}, \vec{E}_0) = (\overline{G}, \vec{E}_1)$ in distribution and hence the desired stationarity. □

Theorem 19

Let $\vec{G} = (G, \vec{E})$ be a stationary and reversible random graph. We consider $\underline{\vec{G}} = (\underline{G}, \underline{\vec{E}})$ the graph obtained from $\vec{G}$ after biasing by the inverse of the degree of the origin of $\vec{E}$. Then $\pi_\bullet(\underline{\vec{G}})$ is a unimodular random graph.

Proof. We will verify the mass-transport principle. Let $f(\mathrm{g}, x, y)$ be a transport function which is zero as soon as x and y are not neighbors. We then form the function $h(\mathrm{g}, \vec{e})$ such that (3.3) holds. Then the same calculation as in the previous proposition shows that the MTP is verified in $\pi_\bullet(\underline{\vec{G}}) = (\underline{G}, \rho)$ for such functions f: if we write $\Delta = \deg(\vec{E}_*)$ then

$$\begin{array}{l} \mathbb{E}\left[h(G, \vec{E})\right] \underset{(3.2)}{=} \mathbb{E}\left[\Delta^{-1} \displaystyle\sum_{\vec{\sigma}_* = \vec{E}_*} h(G, \vec{\sigma})\right] = \mathbb{E}\left[\Delta^{-1} \displaystyle\sum_{x \in \mathrm{V}(G)} f(G, \vec{E}_*, x)\right] = \mathbb{E}\left[\Delta^{-1}\right] \mathbb{E}\left[\displaystyle\sum_{x \in \mathrm{V}(\underline{G})} f(\underline{G}, \rho, x)\right] \\ \text{(rev.)} \quad \| \\ \mathbb{E}\left[h(G, \overleftarrow{E})\right] \underset{(3.2)}{=} \mathbb{E}\left[\Delta^{-1} \displaystyle\sum_{\vec{\sigma}_* = \vec{E}_*} h(G, \overleftarrow{\sigma})\right] = \mathbb{E}\left[\Delta^{-1} \displaystyle\sum_{x \in \mathrm{V}(G)} f(G, x, \vec{E}_*)\right] = \mathbb{E}\left[\Delta^{-1}\right] \mathbb{E}\left[\displaystyle\sum_{x \in \mathrm{V}(\underline{G})} f(\underline{G}, x, \rho)\right] \end{array}$$

Actually, this suffices to imply the full MTP:

Lemma 20. *Let (G, ρ) be a random pointed graph satisfying the MTP for transport functions $f(\mathrm{g}, x, y)$ which are null as soon as x and y are not neighbors in g. Then (G, ρ) is unimodular.*

Proof. Suppose that $f(\mathrm{g}, x, y)$ is a transport function that is zero unless $\mathrm{d_{gr}}(x,y) = k$ for some $k \geq 1$ (any transport function is a sum of functions of the last sort). We denote by $\mathscr{P}(\mathrm{g}, x, y)$ the number of geodesic paths going from x to y in g and $\mathscr{P}_j(\mathrm{g}, x, y; u, v)$ the number of such paths such that the jth step links u to v where $1 \leq j \leq k$. We then form the transport functions for $1 \leq j \leq k$

$$f_j(\mathrm{g}, u, v) = \sum_{x,y \in \mathrm{V}(\mathrm{g})} f(\mathrm{g}, x, y) \frac{\mathscr{P}_j(\mathrm{g}, x, y; u, v)}{\mathscr{P}(\mathrm{g}, x, y)}.$$

Clearly these are transport functions which are null except if u and v are neighbors in g. Since the MTP is valid for such functions we have

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

But on the other hand we have the following deterministic equalities:

$$\sum_{y \in \mathrm{V}(G)} f(G, \rho, y) = \sum_{y \in \mathrm{V}(G)} f(G, \rho, y) \sum_{v : \mathrm{d_{gr}}(\rho, v) = 1} \frac{\mathscr{P}_1(G, \rho, y; \rho, v)}{\mathscr{P}(G, \rho, y)} = \sum_{v \in \mathrm{V}(G)} f_1(G, \rho, v),$$

$$\sum_{x \in \mathrm{V}(G)} f(G, x, \rho) = \sum_{x \in \mathrm{V}(G)} f(G, x, \rho) \sum_{u : \mathrm{d_{gr}}(u, \rho) = 1} \frac{\mathscr{P}_k(G, x, \rho; u, \rho)}{\mathscr{P}(G, x, \rho)} = \sum_{u \in \mathrm{V}(G)} f_k(G, u, \rho),$$

and for $1 \leq j < k$ we have

$$\begin{aligned}
\sum_{u \in \mathrm{V}(G)} f_j(G, u, \rho) &= \sum_{x,y \in \mathrm{V}(G)} f(G, x, y) \sum_{u \in \mathrm{V}(G)} \frac{\mathscr{P}_j(G, x, y; u, \rho)}{\mathscr{P}(G, x, y)} \\
&= \sum_{x,y \in \mathrm{V}(G)} f(G, x, y) \sum_{v \in \mathrm{V}(G)} \frac{\mathscr{P}_{j+1}(G, x, y; \rho, v)}{\mathscr{P}(G, x, y)} = \sum_{v \in \mathrm{V}(G)} f_{j+1}(G, \rho, v),
\end{aligned}$$

since the second sum over paths is just the proportion of those paths going from x to y through the vertex ρ in $j+1$th position. Combining the last displays yields the MTP for the function f. □

3.2 Applications

We now develop the analogs of the propositions we derived for unimodular random graphs in the "weaker" context of stationary random graphs only. When the graph is stationary and reversible one can consider its unimodular version and apply the results of Section 2.3.

3.2.1 Tightness

Proposition 21. *A family $(G_i, \vec{E}^{(i)})_{i \in I}$ of stationary random graphs is tight in $\vec{\mathfrak{G}}$ if and only if the family $(\deg(\vec{E}^{(i)}_*))_{i \in I}$ is tight.*

Proof. The characterization of pre-compact sets in $\vec{\mathcal{G}}$ is analogous to the one in $\mathcal{G}^\bullet$, i.e. they are those sets $A \subset \vec{\mathcal{G}}$ such that there exist $n_0, n_1, n_2, ..., n_r, ...$ satisfying

$$\sup\{\deg(x) : x \in [\vec{g}]_r,\ \vec{g} \in A\} \leq n_r.$$

In the following we denote by $M_r(\vec{g})$ the maximum vertex degree inside $[\vec{g}]_r$. In particular $(\vec{G}_i)_{i\in I}$ is tight if for every $\varepsilon > 0$ there exist $n_0, n_1, ..., n_r$ such that for all $i \in I$ we have $\mathbb{P}(\forall r \geq 0 : M_r(\vec{G}_i) \leq n_r) \geq 1 - \varepsilon$. This is equivalent to the seemingly weaker condition:

$$\forall \varepsilon' > 0, \forall r \geq 0, \exists n'_r \geq 0, \quad \forall i \in I, \quad \mathbb{P}(M_r(\vec{G}_i) \leq n'_r) \geq 1 - \varepsilon'.$$

But taking $\varepsilon' = \varepsilon \cdot 2^{-r}$ and performing a union bound on r, we see that the second condition is indeed equivalent to the first one. We have thus shown that $(\vec{G}_i)_{i\in I}$ is tight if and only if for every $r \geq 0$ the family $(M_r(\vec{G}_i))_{i\in I}$ is tight. The last argument is valid for any family of random rooted graphs, and we have not used stationarity. We will show now that if $(\vec{G}_i)_{i\in I}$ are stationary random graphs and $(M_0(\vec{G}_i))_{i\in I}$ is tight then $(M_r(\vec{G}_i))_{i\in I}$ is also tight for every $r \geq 0$. We treat the case $r = 1$ and leave the general case as an exercise to the reader. By assumption, for any $\varepsilon > 0$ there exists n_0 such that

$$\forall i \in I, \quad \mathbb{P}(M_0(\vec{G}_i) \geq n_0) \leq \varepsilon.$$

Now inside each $\vec{G}_i$ we perform a two-steps simple random walk started from the origin of the root edge and denote by $\vec{E}^{(i)\prime}$ the second directed edge visited. By (3.2) the graph $(G_i, \vec{E}^{(i)\prime})$ has the same distribution as $\vec{G}_i$ and in particular the origin vertex $\vec{E}^{(i)\prime}_*$ of $\vec{E}^{(i)\prime}$ is a random uniform neighbor of $\vec{E}^{(i)}_*$ and its degree has the same law as that of $\deg(\vec{E}^{(i)}_*)$. We choose $\ell \geq n_0$ large enough so that if we sample independently ℓ variables uniformly in $\{1, 2, ..., n_0\}$ then all the outcome have been seen with probability at least $1 - \varepsilon$. We deduce that in $\vec{G}_i$, if we pick ℓ independent neighbors $\rho_1, ..., \rho_\ell$ of $\vec{E}^{(i)}_*$ then on the event where $\deg(\vec{E}^{(i)}_*) \leq n_0$ the set $\{\rho_1, ..., \rho_\ell\}$ covers all the neighbors of $\vec{E}^{(i)}_*$ with probability at least $1 - \varepsilon$. Combining these observations we get that

$$\begin{aligned}
\mathbb{P}(M_1(\vec{G}_i) \geq n_1) &\leq \mathbb{P}(M_0(\vec{G}_i) \geq n_0) + \varepsilon + \ell\mathbb{P}(M_0(\vec{G}_i) \geq n_1) \\
&\leq 2\varepsilon + \ell\mathbb{P}(M_0(\vec{G}_i) \geq n_1).
\end{aligned}$$

Choosing n_1 large enough we can make the above display less than 3ε and that completes the proof for the case $r = 1$. □

Exercise 26. Using Proposition 21 and the above translation between unimodular random graphs and stationary and reversible random graphs, show that if $((G_i, \rho_i))_{i\in I}$ is a family of unimodular random pointed graphs such that $(\deg(\rho_i))_{i\in I}$ is uniformly integrable then $((G_i, \rho_i))_{i\in I}$ is tight in $\mathcal{G}^\bullet$. Show that the latter condition is not necessary and that the condition $\sup_i \mathbb{E}[\deg(\rho_i)] < \infty$ is not sufficient.

3.2.2 $0-1$ laws

We now have the analogs of Proposition 11 and 12:

Proposition 22 (Everything shows at the origin)**.** *Let $\vec{G}$ be a stationary random graph and $A \subset \vec{\mathcal{G}}$ be a Borel set such that $\mathbb{P}(\vec{G} \in A) = 0$. Then the probability that there exists an edge $\vec{e} \in \vec{\mathrm{E}}(G)$ such that $(\vec{G}, \vec{e}) \in A$ is equal to zero.*

Proof. For every $n \geq 0$ if $\vec{E}_n$ is the n-th edge visited by a simple random walk on $\vec{G}$ we have

$$\mathbb{P}((G, \vec{E}_n) \in A) \underset{\text{stat.}}{=} \mathbb{P}((G, \vec{E}_0) \in A) = 0.$$

Summing-up over all $n \geq 0$ we deduce that because the graph is connected

$$0 = \mathbb{E}\left[\sum_{n \geq 0} \mathbb{1}_{(G, \vec{E}_n) \in A}\right] = \int \mathrm{d}\mathbb{P}(\vec{G}) \sum_{\vec{e} \in \vec{\mathrm{E}}(G)} \mathbb{1}_{(G, \vec{e}) \in A} \underbrace{\int \mathrm{dP}_{\vec{G}}((\vec{E}_i)_{i \geq 0}) \sum_{i \geq 0} \mathbb{1}_{\vec{E}_i = \vec{e}}}_{>0}.$$

This proves the desired statement. □

Proposition 23 (If it happens, it happens a lot)**.** *Let $\vec{G}$ be a stationary random graph which is almost surely infinite. Then for any $A \subset \vec{\mathcal{G}}$ Borel we have*

$$\#\{\vec{e} \in \vec{\mathrm{E}}(G) : (G, \vec{e}) \in A\} \in \{0, \infty\} \quad a.s.$$

Proof. Let $A \subset \vec{\mathcal{G}}$ and for a graph $\mathfrak{g}$ denote by $\vec{\mathcal{E}}_A(\mathfrak{g})$ the set of all oriented edges $\vec{e} \in \vec{\mathrm{E}}(\mathfrak{g})$ such that $(g, \vec{e}) \in A$. We argue by contradiction and suppose that with positive probability we have

$$0 < \#\vec{\mathcal{E}}_A(G) < \infty.$$

After conditioning on the above event (which preserves stationarity by an extension Exercise 15) we can suppose that $0 < \#\vec{\mathcal{E}}_A(G) < \infty$ almost surely. Hence we will suppose that the last display happens almost surely. By the last proposition we must have $\mathbb{P}(\vec{G} \in A) > 0$ and so by stationarity

$$0 < \mathbb{P}(\vec{G} \in A) = \mathbb{E}\left[\frac{1}{n} \sum_{i=0}^{n-1} \mathbb{1}_{(G, \vec{E}_i) \in A}\right] = \mathbb{E}\left[\frac{1}{n} \sum_{i=0}^{n-1} \mathbb{1}_{\vec{E}_i \in \vec{\mathcal{E}}_A(G)}\right].$$

However for any infinite graph $\mathfrak{g}$ and any finite subset $F \subset \mathrm{V}(\mathfrak{g})$ the proportion of time spend by a simple random walk in F almost surely goes to 0 (indeed, otherwise the walk would be positive recurrent and it is not possible that simple random walk on an infinite graph is positive recurrent because the only invariant measure is proportional to the degree of vertices which is an infinite measure). Hence by dominated convergence we get that

$$0 < \mathbb{P}(\vec{G} \in A) = \mathbb{E}\left[\frac{1}{n} \sum_{i=0}^{n-1} \mathbb{1}_{(G, \vec{E}_i) \in A}\right] = \mathbb{E}\left[\frac{1}{n} \sum_{i=0}^{n-1} \mathbb{1}_{\vec{E}_i \in \vec{\mathcal{E}}_A(G)}\right] \xrightarrow[n \to \infty]{} 0,$$

which yields a contradiction. □

Using the above propositions, one can adapt the proof of Theorem 14 and deduce:

Corollary 24. *The number of ends of a stationary random graph is almost surely in $\{0, 1, 2\} \cup \{\infty\}$.*

3.2.3 Degree of the root

One now prove the analog of Theorem 13 in the weaker context of stationary random graphs. Recall that since we need to bias by the inverse degree of the origin vertex to go from stationary and reversible random graph to unimodular random graph, the analog of the equality $\mathbb{E}[\deg(\rho)] = 2$ in Section 2.3 becomes $\mathbb{E}[\deg(\rho)^{-1}] = \frac{1}{2}$.

Theorem 25 (***The (inverse) degree tells us a lot!***)

Let $(G, \vec{E})$ be an a.s. infinite stationary random graph whose origin vertex is denoted by ρ, then

$$\mathbb{E}[\deg(\rho)^{-1}] \le \frac{1}{2}.$$

Besides if the above inequality is an equality, then G is a.s. a random tree (with one or two ends, as proved below in Theorem 41).

Proof. Let $(\vec{E}_n)_{n\ge0}$ the sequence of oriented edges visited by the random walk on G and denote by $(X_n)_{n\ge0}$ their origin vertex. The idea is to consider the function $H_n = \mathrm{d_{gr}}(X_0, X_n)$ and in particular its conditional variation. More precisely, given the graph G, we can condition on the σ-field $\mathcal{F}_n$ generated by the first $n-1$ steps (so that the walk sits on X_n) and compute

$$\mathrm{E}_{(G,\vec{E})}[H_{n+1} - H_n \mid \mathcal{F}_n] = \frac{1}{\deg(X_n)} \left(P_{G,X_0}(X_n) - N_{G,X_0}(X_n)\right),$$

where $P_{\mathrm{g},x}(y)$ is the number of edges in g starting from y pointing to a vertex at a distance strictly larger than $\mathrm{d_{gr}}(x, y)$ and similarly for $N_{(\mathrm{g},x)}(y)$ after replacing strictly larger than strictly less. Clearly we must have $N_{\mathrm{g},x}(y) \ge 1$ as long as $x \ne y$ since there must be a geodesic path going from y to x in g, and so we can always write:

$$\mathrm{E}_{(G,\vec{E})}[H_{n+1} - H_n \mid \mathcal{F}_n] \quad \le \quad \frac{\deg(X_n) - 2}{\deg(X_n)} + \mathbb{1}_{X_n = X_0} \frac{2}{\deg(X_n)}. \tag{3.4}$$

Now, taking expectation with respect to the simple random walk and the graph we deduce that

$$\mathbb{E}[\mathrm{d_{gr}}(X_0, X_n)] = \sum_{i=0}^{n-1} \mathbb{E}\left[\frac{\deg(X_i) - 2}{\deg(X_i)}\right] + \mathbb{E}\left[\sum_{i=0}^{n-1} \frac{2 \cdot \mathbb{1}_{X_0 = X_i}}{\deg(X_i)}\right].$$

Using the fact that the walk is either transient or null recurrent (and never positive recurrent), we can argue as in the preceding proof and deduce that the second term in the right-hand size is asymptotically negligible in front of n. We now use the stationarity of G to see that $\mathbb{E}\left[\frac{\deg(X_n)-2}{\deg(X_n)}\right]$ does not depend on n and equals $\mathbb{E}[\frac{\deg(\rho)-2}{\deg(\rho)}]$. Since $n^{-1}\mathrm{d_{gr}}(X_0, X_n)$ is obviously non negative we get the desired estimate

$$\mathbb{E}\left[\deg(\rho)^{-1}\right] \le \frac{1}{2}.$$

Let us now prove that if the above inequality is an equality, then G is almost surely a tree. Arguing by contradiction, we suppose that $\mathbb{P}(G \text{ is not a tree}) > 0$. We can therefore find an $r \in \mathbb{N}^*$ such that $\mathbb{P}([(G, \vec{E})]_r \text{ contains a cycle}) > 0$ and more precisely one can find a finite connected rooted graph $(g_0, \vec{e}_0)$ of radius $r > 0$ such that:

- $(g_0, \vec{e}_0)$ contains a cycle $(x_1, \ldots, x_k)$ with $k \geq 1$ and $x_1, \ldots, x_k$ are all distincts vertices of g_0 at distance less than $r-1$ from $\vec{e}_0$,
- $\mathbb{P}([(G, \vec{E})]_r = (g_0, \vec{e}_0)) > 0$

A vertex x is called an *intruder* in $(\mathrm{g}, \vec{e})$ if (3.4) is not saturated, meaning that the number of half-edges in g starting from x and pointing towards a vertex at graph distance $\mathrm{d_{gr}}(\rho, \vec{e}_*) + 1$ is strictly less than $\deg(x) - 1$. Then we have the following deterministic geometric lemma:

Lemma 26. *Every cycle of any rooted connected graph contains an intruder.*

Proof. Let $(\mathrm{g}, \vec{e})$ be a connected rooted graph and $x_1, \ldots, x_k$ $(k \geq 1)$ distincts vertices of g such that the edges $(x_1, x_2), \ldots, (x_{k-1}, x_k), (x_k, x_1)$ are in $\mathsf{E}(\mathrm{g})$ and are all distincts ($x_1, \ldots, x_k$ is a cycle). The case $k = 1$ is trivial and x_1 is necessarily an intruder. We consider now $k \geq 2$.

Let us suppose that the vertices of the cycle are not intruders in $(\mathrm{g}, \vec{e})$. Therefore, all half-edges but one emanating from x_i link to a vertex at distance from the origin $\vec{e}_*$ strictly larger than x_i. Without loss of generality, we can suppose that $\mathrm{d_{gr}}(\vec{e}_*, x_1) = \min_{1 \leq i \leq k} \mathrm{d_{gr}}(\vec{e}_*, x_i)$. We have then $\mathrm{d_{gr}}(\vec{e}_*, x_2) = \mathrm{d_{gr}}(\vec{e}_*, x_1) + 1$ (if $\mathrm{d_{gr}}(\vec{e}_*, x_2) = \mathrm{d_{gr}}(\vec{e}_*, x_1)$ then x_2 is an intruder). The edge $x_2 \to x_1$ of that cycle is thus the only edge that goes from x_2 to an other vertex closer to $\vec{e}_*$. Therefore $\mathrm{d_{gr}}(\vec{e}_*, x_3) = \mathrm{d_{gr}}(\vec{e}_*, x_2) + 1$ etc. Going around the cycle we reach a contradiction. □

Coming back to the proof of the theorem, conditionally on $[(G, \vec{E}_n)]_r = (\mathrm{g}_0, \vec{e}_0)$, we know from the above lemma that the (image inside G of the) cycle $(x_1, \ldots, x_k)$ of $(\mathrm{g}_0, \vec{e}_0)$ contains an intruder x for $(G, \vec{E})$ which might depend on $\vec{E}_n$ of course. Then one can find $d > 0$ and $\varepsilon > 0$ which only depend on g_0 and not on the intruder x such that conditionally on $[(G, \vec{E}_n)]_r = (\mathrm{g}_0, \vec{e}_0)$ one of the vertices $X_n, X_{n+1}, ..., X_{n+d-1}$ visited by the SRW is the intruder x with probability at least ε. With the above notation we can now write

$$\mathrm{E}_{(G,\vec{E})}[H_{n+1} - H_n | X_0, \ldots, X_n] \leq \frac{\deg(X_n) - 2}{\deg(X_n)} - \frac{2}{\deg(X_n)} \mathbb{1}_{X_n = X_0} - \frac{1}{\deg(X_n)} \mathbb{1}_{X_n \text{ is an intruder in } (G, \vec{E})}.$$

Taking expectation with respect to the graph and summing from $k\,d \leq n < (k+1)d$ we obtain using stationarity of the underlying graph

$$\begin{aligned} \mathbb{E}[H_{(k+1)d} - H_{kd}] \quad &\leq \quad d\,\mathbb{E}\left[\frac{\deg(X_0) - 2}{\deg(X_0)}\right] - 2 \sum_{n=kd}^{(k+1)d-1} \mathbb{E}\left[\frac{\mathbb{1}_{X_n = X_0}}{\deg(X_n)}\right] \\ &\quad - \mathbb{E}\left[\sum_{n=kd}^{(k+1)d-1} \frac{\mathbb{1}_{X_n \text{ is an intruder in } (G, \vec{E}))}}{\deg(X_n)}\right] \end{aligned}$$

but by our assumptions, the third term in the right-hand side is smaller than $-\varepsilon$. After arguing as above (using the transience or null-recurrence of the walk to show that the second term is negligible as $k \to \infty$) we deduce that as $k \to \infty$ we have using our assumption $\mathbb{E}[\deg(X_0)^{-1}] = \frac{1}{2}$ that

$$\limsup_{k \to \infty} \mathbb{E}[H_{(k+1)d} - H_{kd}] \leq -\varepsilon.$$

We indeed reach a contradiction after summing-up over all k since the expectation of H_n cannot be negative. □

3.3 Connection with ergodic theory

We now reformulate Definition 9 in the framework of ergodic theory in order to be able to apply powerful tools such as the subadditive ergodic theorem, Poincaré recurrence theorem... Recall that the basic ingredients of ergodic theory are 1) a probability space $(E, \mathcal{A}, \mu)$ and 2) an action $\theta : E \to E$ measurable which preserves the measure μ in the sense that $\mu(\theta^{-1}(A)) = \mu(A)$ for all $A \in \mathcal{A}$.

3.3.1 Framework

We will work on the set of locally finite connected graphs endowed with a possibly infinite path made of oriented edges $(\mathrm{g}, (\vec{e}_i)_{a<i<b})$ where $a \in \{-\infty\} \cup \{\ldots, -2, -1\}$ and $b \in \{1, 2, \ldots\} \cup \{\infty\}$. The oriented edge $\vec{e}_0$ could be seen as the root edge of the graph. As usual, we will identify $(\mathrm{g}, (\vec{e}_i)_{i\in(a,b)})$ with $(\mathrm{g}', (\vec{e}_i{}')_{i\in(a,b)})$ if there exists a graph homomorphism $\mathrm{g} \to \mathrm{g}'$ mapping $\vec{e}_i$ to $\vec{e}_i{}'$ for every $i \in (a, b)$. The quotient space is then denoted by $\mathcal{G}^{\leftrightarrow}$ and endowed with the local distance corresponding to the notion of restriction

$$\left[\left(\mathrm{g}, (\vec{e}_i)_{a<i<b}\right)\right]_r = \left([(\mathrm{g}, \vec{e}_0)]_r, (\vec{e}_i)_{i\in(a,b)\cap(-r,r)}\right).$$

The resulting space $(\mathcal{G}^{\leftrightarrow}, \mathrm{d_{loc}})$ is again Polish by Section 1.2.1. We focus in what follows on the subspace $\mathcal{G}^{\rightarrow}$ made of those graphs with a semi-infinite path indexed by $i \in \mathbb{N}$ i.e. when $a = -1$ and $b = \infty$. This space comes with a natural shift

$$\theta(\mathrm{g}, (\vec{e}_i)_{i\geq 0}) \mapsto (\mathrm{g}, (\vec{e}_{i+1})_{i\geq 0}).$$

Recall that to any random rooted graph $(G, \vec{E})$ we associated a probability measure on $\mathcal{G}^{\rightarrow}$ by launching a simple random walk on it, more precisely by considering the distribution

$$\mu = \int \mathrm{d}\mathbb{P}(G, \vec{E}) \int \mathrm{dP}_{(G,\vec{E})}((\vec{E}_i)_{i\geq 0}),$$

where we recall that $\mathrm{P}_{(\mathrm{g},\vec{e})}$ is the law of a simple random walk starting with $\vec{e}$ in g. It is an easy matter to translate the notion of stationary into θ-invariance:

Proposition 27. *The random graph $(G, \vec{E})$ is stationary if and only if μ is θ-invariant.*

Proof. If μ is θ-invariant we in particular get that $(G, \vec{E}_0) = (G, \vec{E}_1)$ where $(\vec{E}_i)_{i\geq 0}$ is a simple random walk on $(G, \vec{E})$ and so the latter is stationary. Conversely, by the Markovian property of the simple random walk one can construct $\theta(G, (\vec{E})_{i\geq 0}) = (G, (\vec{E}_{i+1})_{i\geq 0})$ by first making one step of random walk to discover $\vec{E}_1$ then re-rooting at $\vec{E}_1$ to finally launch the rest of the walk independently of this step. Since $(G, \vec{E}_0) = (G, \vec{E}_1)$ this algorithm produces a random graph with a semi-infinite path which has the same law as μ. □

Exercise 27. (∗) If $(G, (\vec{E}_i)_{i\geq 0})$ is a stationary random graph given with a simple random walk on it, show that as $k \to \infty$ the variable $(G, (\vec{E}_{i+k})_{-k\leq i<\infty})$ converges in distribution in $\mathcal{G}^{\leftrightarrow}$ and that the resulting distribution is invariant by shifting along the path by ±1. Show that the law obtained is invariant by time reversal of the path if and only if the original graph $(G, \vec{E}_0)$ is reversible.

Let us give a recreative application using the famous Poincaré recurrence theorem:

Proposition 28. *Let $(G, \vec{E})$ be a stationary random graph and denote by $(\vec{E}_i)_{i\geq 0}$ a simple random walk on it. Then we almost surely have*

$$\liminf_{n\to\infty} \mathrm{d_{loc}}\Big((G, \vec{E}_0), (G, \vec{E}_n)\Big) = 0.$$

In words, the result says that almost surely when performing a simple random walk on a stationary random graph we will discover places where the landscape around the current oriented edge is arbitrarily close to the starting landscape. Notice that the last proposition is trivial when the simple random walk on G is almost surely recurrent.

Proof. This is an application of Poincaré recurrence theorem. We recall in the next exercise the general form from which the above result can be deduced in the metric case. □

Exercise 28. Prove Poincaré recurrence theorem: If $(X, \mathcal{A}, \mu)$ is a measurable space with a finite measure μ and $\theta : X \to X$ preserves μ then for any measurable $A \subset \mathcal{A}$ and for μ-almost all $x \in A$ there exists an infinite number of $n \geq 0$ such that $\theta^n(x) \in A$. Hint: Consider $A' = \{x \in A : \theta^n(x) \notin A, \forall n \geq 1\}$ and show that $\theta^{-k}(A')$ are pairwise disjoint subspaces for $k \geq 0$.

3.3.2 Ergodicity and applications

Recall that θ is ergodic (for the measure μ) if for any measurable set A such that $\mu(A\Delta\theta^{-1}(A)) = 0$ then $\mu(A) \in \{0, 1\}$. In words, ergodicity means that the shift operation does not stabilize any non trivial event.

Definition 10. *We say that a stationary random graph $(G, \vec{E})$ is ergodic (or its law is ergodic) if the measure μ on $\mathcal{G}^{\rightarrow}$ associated by the above means is ergodic for the shift θ.*

We will usually admit ergodicity when needed. We recall Kingman's subadditive ergodic theorem which generalizes the well-known pointwise ergodic theorem of Birkhoff.

Theorem 29 (*Kingman's subadditive ergodic theorem*)

If θ is a measure preserving transformation on a probability space $(E, \mathcal{A}, \mu)$ and $(h_n)_{n\geq 1}$ is a sequence of integrable functions satisfying for $n, m \geq 1$

$$h_{n+m}(x) \leq h_n(x) + h_m(\theta^n x),$$

then we have the following convergence almost sure and in L^1

$$\frac{h_n(x)}{n} \xrightarrow[n\to\infty]{a.s. \text{ and } L^1} h(x),$$

where $h(x)$ is a θ-invariant function (in particular constant if θ is ergodic).

When $h_n(x) = x + f(x) + f(\theta x) + \ldots + f(\theta^{n-1}x)$ for some integrable function $f : E \to \mathbb{R}$ we recover Birkhoff's famous ergodic theorem. There are many generalizations of the above subadditive ergodic theorem but this version is already sufficient to provide useful applications in the context of random graphs. In the following of this section $(G, \vec{E})$ is a stationary and ergodic random graph and we denote by μ the associated probability measure on $\mathcal{G}^{\to}$.

Speed of the random walk. We consider the function

$$h_n(\mathrm{g}, (\vec{e}_i)_{i\geq 0}) = \mathrm{d_{gr}}((\vec{e}_0)_*, (\vec{e}_n)_*),$$

where $(\vec{e}_i)_*$ is the origin vertex of $\vec{e}_i$ and $\mathrm{d_{gr}}$ is the graph distance. The triangular inequality shows that

$$\begin{aligned} h_{n+m}(\mathrm{g}, (\vec{e}_i)_{i\geq 0}) = \mathrm{d_{gr}}((\vec{e}_0)_*, (\vec{e}_{n+m})_*) &\leq \mathrm{d_{gr}}((\vec{e}_0)_*, (\vec{e}_n)_*) + \mathrm{d_{gr}}((\vec{e}_n)_*, (\vec{e}_{n+m})_*) \\ &= h_n(\mathrm{g}, (\vec{e}_i)_{i\geq 0}) + h_m(\mathrm{g}, (\vec{e}_i)_{i\geq n}) = h_n(\mathrm{g}, (\vec{e}_i)_{i\geq 0}) + h_m(\theta^n(\mathrm{g}, (\vec{e}_i)_{i\geq 0})). \end{aligned}$$

Clearly the functions h_n are bounded by n and so are integrable. We can thus apply Theorem 29 and get the existence of a constant $s \geq 0$ (for speed) such that

$$\frac{\mathrm{d_{gr}}((\vec{E}_0)_*, (\vec{E}_n)_*)}{n} \xrightarrow[n\to\infty]{a.s.} s. \tag{3.5}$$

Linear growth of the range. We consider the function

$$r_n(\mathrm{g}, (\vec{e}_i)_{i\geq 0}) = \#\{(\vec{e}_0)_*, \ldots, (\vec{e}_n)_*\}$$

be the number of different vertices visited by the walk in the first n steps. It is plain that $\#\{(\vec{e}_0)_*, \ldots, (\vec{e}_{n+m})_*\} \leq \#\{(\vec{e}_0)_*, \ldots, (\vec{e}_n)_*\} + \#\{(\vec{e}_n)_*, \ldots, (\vec{e}_{n+m})_*\}$ and the second term is $r_m(\theta^n(\mathrm{g}, (\vec{e}_i)_{i\geq 0}))$. Since the r_n are also integrable, we are in position to apply Theorem 29 again and get

$$\frac{\#\{(\vec{E}_0)_*, \ldots, (\vec{E}_n)_*\}}{n} \xrightarrow[n\to\infty]{a.s.} r, \tag{3.6}$$

for some $r \geq 0$. Actually we can in this case express exactly the constant r. Indeed, even if the graph is not ergodic, as long as it is stationary we can evaluate the expectation of the range of the random walk as follows

$$\begin{aligned} \mathbb{E}\left[\#\left\{(\vec{E}_0)_*, \ldots, (\vec{E}_n)_*\right\}\right] &= \mathbb{E}\left[\sum_{k=0}^{n} \mathbb{1}\{\text{it is the last visit to } (\vec{E}_k)_* \text{ before time } n\}\right] \\ &= \sum_{k=0}^{n} \mathbb{P}\left((\vec{E}_i)_* \neq (\vec{E}_k)_*, \forall k+1 \leq i \leq n\right) \\ &\underset{\text{stat.}}{=} \sum_{k=0}^{n} \mathbb{P}\left((\vec{E}_0)_* \neq (\vec{E}_i)_*, \forall 1 \leq i \leq k\right). \end{aligned}$$

But by dominated convergence we have $\mathbb{P}\left((\vec{E}_0)_* \neq (\vec{E}_i)_*, \forall 1 \leq i \leq k\right) \to \mathbb{P}\left((\vec{E}_0)_* \neq (\vec{E}_i)_*, \forall i \geq 1\right)$ as $k \to \infty$. Performing Cesaro's summation we deduce that

$$n^{-1}\,\mathbb{E}\left[\#\left\{(\vec{E}_0)_*, \ldots, (\vec{E}_n)_*\right\}\right] \xrightarrow[n\to\infty]{} \mathbb{P}\left((\vec{E}_0)_* \neq (\vec{E}_i)_*, \forall i \geq 1\right). \tag{3.7}$$

We deduce that if the graph is almost surely recurrent (not necessarily ergodic) then the range of the random walk grows sublinearly in probability, a fact that we will use later on. If the graph is ergodic, we can compare the last display together with (3.6) and get by dominated convergence that r is the averaged probability that a simple random walk on $\vec{G}$ does not come back to its starting point.

Another application of Kingman theorem for the function

$$\mathrm{Ent}_n(\mathrm{g}, (\vec{e}_i)_{i\geq 0}) = -\log \mathrm{P}_{(\mathrm{g}, \vec{e}_0)}(\vec{E}_n = \vec{e}_n),$$

will be presented in the next chapter in connection with the notion of Shannon entropy.

Let us make use of the above machinery to prove the following result:

Theorem 30 (***Non reversibility implies transience***)

Let $(G, \vec{E})$ be an ergodic stationary random graph. If $(G, \vec{E})$ is not reversible then it must also be almost surely transient.

Proof. We denote $(G, (\vec{E})_{i\geq 0})$ the random graph with the random walk path in $\mathcal{G}^{\to}$ obtained by launching the SRW on G. Since G is not reversible, there must exist a bounded function $F : \vec{\mathcal{G}} \to \mathbb{R}_+$ such that we have

$$\alpha = \mathbb{E}[F(G, \vec{E})] \neq \mathbb{E}[F(G, \overleftarrow{E})] = \beta.$$

The ergodic theorem hence ensures that

$$\frac{1}{n}\sum_{i=0}^{n-1} F(G, \vec{E}_i) \xrightarrow[n\to\infty]{a.s.} \alpha \quad \text{and} \quad \frac{1}{n}\sum_{i=0}^{n-1} F(G, \overleftarrow{E}_i) \xrightarrow[n\to\infty]{a.s.} \beta.$$

Now if we suppose by contradiction that G is recurrent with a positive probability, since this event is shift invariant then by ergodicity G must be almost surely recurrent. Writing τ_k for the successive return times of the simple random to the origin vertex we deduce from the last display and the dominated convergence theorem that

$$\mathbb{E}\left[\frac{1}{\tau_k}\sum_{i=0}^{\tau_k-1} F(G, \vec{E}_i)\right] \xrightarrow[k\to\infty]{} \alpha \quad \text{and} \quad \mathbb{E}\left[\frac{1}{\tau_k}\sum_{i=0}^{\tau_k-1} F(G, \overleftarrow{E}_i)\right] \xrightarrow[k\to\infty]{} \beta.$$

However, in any graph, by reversibility of the simple random walk path, the law of $(\vec{E}_0, \ldots, \vec{E}_{\tau_k-1})$ is equal to that of $(\overleftarrow{E}_{\tau_k-1}, \ldots, \overleftarrow{E}_0)$ since the density of these two paths is proportional to

$$\left(\deg(X_1) \times \cdots \times \deg(X_{\tau_k-1})\right)^{-1}$$

where $X_0, X_1, \ldots, X_{\tau_k}$ are the vertices visited by the simple random walk. Combining this observation with the last display shows that $\alpha = \beta$ which is absurd! □

Biliographical references. The connection between stationary reversible random graphs and unimodular random graphs is explicit in [4] in particular thanks to the involution invariant property (Lemma 20 above). The exposition here is adapted from [11]. The tightness criterion for stationary random graphs is taken from [13], see also [8]. Theorem 25 seems to be new. Finally the link with ergodic theory is classical and can be found e.g. in [11] where the application to the entropy of the walk was first derived following the approach of [31] in the case of fixed regular graphs. The proof of Theorem 30 is new but a stronger result has been proved in [11]: an ergodic stationary non reversible random graph (with bounded degreee) must actually have positive speed.

Chapter 4

Entropy of stationary random graphs

In this section we study the entropy of the random walk on stationary random graphs and show that we can adapt the classical criterion on groups to show that stationary random graphs have the Liouville property if and only if there asymptotic entropy is equal to 0. We first recall the basics on entropy of discrete random variables.

4.1 Entropy of discrete random variables

4.1.1 Shannon entropy

We start with the definition:

Definition 11 (Shannon entropy). *Let X be a random variable taking values in discrete space $\mathcal{E}$ whose law is described by the family of the probabilities $\{p(x) : x \in \mathcal{E}\}$. The* entropy *of the variable X is defined as*

$$\mathcal{H}(X) := \mathbb{E}[-\log(p(X))] = \sum_{x \in \mathcal{E}} -p(x) \log(p(x)).$$

Clearly the definition of $\mathcal{H}(X)$ in fact only depends on the law of the random variable and we should rather write $\mathcal{H}(p)$. However we will always make this abuse of notion and speak of the entropy of random variables. This quantity measures the “dispersion”, the “information” or the “surprise” that the random variable X generates. Indeed, if $p(x)$ happens to be very small, then we are very “informed” or “surprised” to see that $X = x$ and the quantity of information needed to describe such an event is by definition the logarithm of the last probability (think of the number of bits in base 2). Let us gather a few properties of the entropy:

Proposition 31. *The entropy is :*

1. *additive with respect to independence : if X and Y are independent then*

$$\mathcal{H}(X, Y) = \mathcal{H}((X, Y)) = \mathcal{H}(X) + \mathcal{H}(Y),$$

2. *maximal when uniform :* $\mathcal{H}(X) \le \log(\#\mathrm{Supp}(X))$ *with equality when* X *is uniformly distributed,*

3. *diminishes under mapping :* $\mathcal{H}(f(X)) \le \mathcal{H}(X)$ *for any map* f.

Proof. The first point follows from the fact that if $p(x,y)$ is the law of (X,Y) then we have $p(x,y) = p_1(x)p_2(y)$ where p_1 and p_2 are the laws of X and Y respectively and so

$$\begin{aligned}\mathcal{H}((X,Y)) &= -\sum_{x,y} p(x,y)\log(p(x,y))\\ &= -\sum_{x} p_1(x)p_2(y)\Big(\log p_1(x) + \log p_2(y)\Big)\\ &= -\sum_{x} p_1(x)\log p_1(x) - \sum_{y} p_2(y)\log p_2(y) = \mathcal{H}(X) + \mathcal{H}(Y).\end{aligned}$$

For the second point we use the concavity of $\log(\cdot)$ and Jensen's inequality to get that

$$\mathcal{H}(X) = \mathbb{E}[-\log(p(X))] = \mathbb{E}[\log(1/p(X))] \le \log \mathbb{E}[1/p(X)] = \log(n),$$

where n is the size of the support for p. By the equality case in Jensen's inequality the above is sharp if and only if $1/p(X)$ is constant, meaning that $p(X)$ is constant, or in other words that X is uniformly distributed over its support.
The third point is easily checked for simple functions which identifies two points, because

$$-(p(x)+p(y))\log(p(x)+p(y)) \le -p(x)\log p(x) - p(y)\log p(y).$$

The general case is easily deduced by induction. □

Exercise 29. Compute the entropy of a Bernoulli law, and of a geometric law. Show that among all distributions on $\{1,2,3,...\}$ the geometric law maximes the entropy for a given mean. (In general, the probability measures which maximize the entropy for a given mean energy level are measures whose density with respect to the uniform measure is proportional to $\exp(\beta \mathrm{Energy}(\cdot))$ for some $\beta \in \mathbb{R}$ tuned to reach the desired mean energy level. These are called Gibbs measures.).

Exercise 30 (Asymptotic equipartition theorem). Let $(\mathcal{E}, p)$ be a finite probability space. For $n \ge 1$ consider the product space E^n (of cardinality $(\#E)^n$) endowed with the product measure $p^{\otimes n}$. For any $\varepsilon > 0$ consider

$$\mathcal{S}_\varepsilon(n) = \inf\{\#A : A \subset E^n \text{ such that } p^{\otimes n}(A) > 1-\varepsilon\}.$$

Show that for any $\varepsilon \in (0,1)$ we have $n^{-1}\log \#\mathcal{S}_\varepsilon(n) \to \mathcal{H}(p)$. Reprove that $\mathcal{H}(p) \le \log(\#E)$.

4.1.2 Conditional entropy

The fact that the entropy decreases by mapping shows that the entropy of (X,Y) is always larger than or equal to the entropy of X plus the entropy of Y whatever the correlations that may exist between X and Y. We can thus define the *conditional entropy*

$$\mathcal{H}(X|Y) = \mathcal{H}(X,Y) - \mathcal{H}(Y) \ge 0.$$

If p, p_1, p_2 denotes the law of $(X, Y), X, Y$ respectively, a simple calculation shows that

$$\begin{aligned}
\mathcal{H}(X|Y) &= -\sum_{x,y} p(x,y) \log p(x,y) - \sum_{x,y} p(x,y) \log p_2(y) \\
&= -\sum_{x,y} p(x,y) \log \frac{p(x,y)}{p_2(y)} \\
&= \sum_{y} p_2(y) \left(-\sum_{x} p(x|y) \log p(x|y) \right) = \mathbb{E}_Y[\mathcal{H}(\mathcal{L}(X|Y))].
\end{aligned}$$

Definition 12 (Conditional entropy)**.** *On a probability space with underlying probability measure $\mathbb{P}$ let $\mathcal{F}$ be an arbitrary σ-field and X be a discrete random variable. We define the conditional entropy of X with respect to $\mathcal{F}$ as*

$$\mathcal{H}(X|\mathcal{F}) = \mathbb{E}\left[\mathcal{H}(\mathcal{L}(X|\mathcal{F}))\right] = \mathbb{E}\left[-\sum_x \mathbb{P}(X = x \mid \mathcal{F}) \log \mathbb{P}(X = x \mid \mathcal{F}) \right].$$

Exercise 31. Check that the last two definitions agree when $\mathcal{F} = \sigma(Y)$.

Proposition 32. *On a probability space supporting the random variable X we have*

1. *if $\mathcal{G} \subset \mathcal{F}$ are two σ-fields then $\mathcal{H}(X \mid \mathcal{G}) \geq \mathcal{H}(X \mid \mathcal{F})$ with equality if and only if the conditional laws of X given $\mathcal{F}$ or $\mathcal{G}$ are the same. In particular, $\mathcal{H}(X \mid \mathcal{F}) = 0$ if and only if X is independent of $\mathcal{F}$.*

2. *if $\mathcal{F}_n$ is a decreasing sequence of σ-fields then $\mathcal{H}(X \mid \mathcal{F}_n) \uparrow \mathcal{F}(X \mid \bigcap_{n \geq 0} \mathcal{F}_n)$.*

Proof. Using Jensen's inequality for the conditional expectation (remark that $\phi(x) = -x \log x$ is concave) and the tower property we have

$$\begin{aligned}
\mathcal{H}(X \mid \mathcal{F}) = \sum_x \mathbb{E}\left[\phi\left(\mathbb{E}[1_{X=x} \mid \mathcal{F}]\right)\right] &\underset{\text{tower}}{=} \sum_x \mathbb{E}\left[\mathbb{E}\left[\phi\left(\mathbb{E}[1_{X=x} \mid \mathcal{F}]\right) \mid \mathcal{G}\right]\right] \\
&\underset{\text{Jensen}}{\leq} \sum_x \mathbb{E}\left[\phi\left(\mathbb{E}\left[\mathbb{E}[1_{X=x} \mid \mathcal{F}] \mid \mathcal{G}\right]\right)\right] \\
&\underset{\text{tower}}{\leq} \sum_x \mathbb{E}\left[\phi\left(\mathbb{E}\left[1_{X=x} \mid \mathcal{G}\right]\right)\right] = \mathcal{H}(X \mid \mathcal{G}).
\end{aligned}$$

The above inequality is sharp if and only if Jensen's conditional inequality is sharp, i.e. if for every x almost surely we have $\mathbb{E}\left[1_{X=x} \mid \mathcal{G}\right] = \mathbb{E}\left[1_{X=x} \mid \mathcal{F}\right]$. The desired statements easily follows from that.
The second point follows from the fact that $\mathbb{E}[1_{X=x} \mid \mathcal{F}_n]$ is a positive retrograde martingale and so converges almost surely towards $\mathbb{E}[1_{X=x} \mid \cap \mathcal{F}_n]$. By Fatou we thus get $\mathcal{H}(X \mid \cap \mathcal{F}_n) \leq \lim \mathcal{H}(X \mid \mathcal{F}_n)$ but the other inequality is granted by the first point. □

4.2 Tail and invariant σ-fields, Poisson boundary

Let g be a connected graph with positive conductances on every edge. If a origin vertex $o \in \mathrm{V}(\mathsf{g})$ is distinguished, we can consider the random walk $(X_i)_{i \geq 0}$ on g starting from o which we see here as a Markov chain on the vertices of $\mathfrak{g}$. The random walk is thus a random variable taking values in the space Paths of infinite paths of neighboring vertices in $\mathfrak{g}$ starting from ρ on which the shift operation θ acts naturally $\theta((x_i)_{i\geq 0}) = (x_i)_{i\geq 1}$. For each $n \geq 0$ we define $\mathcal{F}^n = \sigma(X_0, ..., X_n)$ as well as $\mathcal{F}^\infty = \sigma(X_i : i \geq 0)$. We then consider two types of σ-fields containing "limit information"

Definition 13. *The* tail *σ-field $\mathcal{T}$ is $\mathcal{T} = \bigcap_{n\geq 0} \sigma(X_i : i \geq n)$.*
The invariant *σ-field $\mathcal{I}$ is made of all the events A such that if $\omega \in A$ then $\theta(\omega) \in A$ as well.*

Clearly $\mathcal{I} \subset \mathcal{F}$. Here is a trivial case where $\mathcal{I} \neq \mathcal{F}$: consider the graph $a - b$ made of a single edge of conductance 1. Then the invariant sigma field is always trivial whereas the tail sigma field contains events of the type "the simple random walk crosses from a to b at every large enough *even* times". Unless for this parity reasons the two σ-fields $\mathcal{I}$ and $\mathcal{T}$ coincide in general (for exemple for the lazy simple random i.e. when the walk has at each step a probability 1/2 to stay put), see [37, Theorem 14.18 and Section 14.6].

There is also a third way of encapsulating limit behavior of a random walk:

Definition 14. *A function $f : \mathrm{V}(\mathfrak{g}) \to \mathbb{R}$ is* harmonic *if for any $x \in \mathrm{V}(\mathfrak{g})$ we have $\mathbb{E}_x[f(X_1)] = f(x)$ where X_1 is a one-step random walk (on the vertices of the graph) started from x (in words $f(x)$ is the mean of its neighbors).*

In general the invariant σ-field can be interpreted in terms of bounded harmonic functions:

Theorem 33 (***Description of the invariant σ-field in terms of Poisson boundary***)

We have a linear map between the space $L^\infty(\mathsf{Paths}, \mathcal{I}, \mathbb{P}_o)$ and the space of bounded harmonic functions :

$$\begin{array}{rcl} L^\infty(\mathsf{Paths}, \mathcal{I}, \mathbb{P}_o) & \to & \textit{Bounded Harmonic Functions} \\ f \in & \mapsto & \left(h_f : x \mapsto \mathbb{E}_x[f(X_0, X_1, ..., X_n, ...)]\right) \\ f_h(X_0, X_1, ..., X_n, ...) \overset{a.s.}{=} \lim\limits_{n\to\infty} h(X_n) & \leftarrow & h \end{array}$$

Proof. Let us go from the left to the right. If f is a bounded invariant function then by invariance and the Markov property at time 1 we have

$$h_f(x) = \mathbb{E}_x[f(X_0, X_1, ..., X_n, ...)] = \mathbb{E}_x[f(X_1, ..., X_n, ...)] = \mathbb{E}_x[h_f(X_1)].$$

Hence h_f is indeed a bounded harmonic function. It is easy to see that h_f in fact does not depend on the equivalence class of f. For the reverse mapping, notice that if h is a bounded harmonic function then $h(X_n)$ is a bounded martingale and so converges almost surely so we can define f_h. To see that the two mappings are inverse one another remark that

$$f(X_0, ..., X_n, ...) \overset{p.s.}{=} \lim_{n\to\infty} \mathbb{E}[f(X_0, ...) \mid \mathcal{F}_n]$$

by the martingale convergence theorem. Yet the right hand side of the last display can be written by the Markov property as

$$\mathbb{E}_x[f(X_0,\dots) \mid \mathcal{F}_n] = \mathbb{E}_{X_n}[f(\tilde{X}_0,\dots)] = h_f(X_n),$$

and so $f_{h_f} = f$ almost surely as desired. □

We deduce in particular from the last theorem that the following definition makes sense:

Definition 15. *The graph* g *has the Liouville property (in short* g *is Liouville) if the only bounded harmonic functions it carries are constant, or equivalently if the invariant* σ*-field* $\mathcal{I}$ *is trivial.*

Examples: All the recurrent graphs are Liouville. The graphs $\mathbb{Z}^d$ for $d \geq 1$ are Liouville, more generally the Cayley graphs of abelian groups are Liouville. The full k-ary tree with $k \geq 3$ are non-Liouville.

4.3 Entropy on stationary random graphs

4.3.1 The mean entropy

Let $\vec{G}$ be a stationary random graph with a random walk on it. In this section it is more convient to see the walk as a sequence of vertices $(X_0, X_1, \dots)$ starting at the extremity of the root edge (so that X_1 is not necessarily deterministic). In particular, the stationarity of the underlying graph shows that $(G; X_0, X_1, \dots)$ has the same law as $(G; X_1, X_2, \dots)$.

Proposition 34. *With the above notation, for any* $1 \leq i_1 < i_2 < \dots < i_\ell$ *we have*

$$\mathbb{E}[\mathcal{H}(X_{i_1}, X_{i_2}, \dots, X_{i_\ell})] = \mathbb{E}[\mathcal{H}(X_{i_1})] + \mathbb{E}[\mathcal{H}(X_{i_2-i_1}, X_{i_3-i_1}, \dots, X_{i_\ell - i_1})].$$

Proof. If (g, ρ) is a pointed graph we denote by $p^I_{(\mathrm{g},\rho)}$ the law of $(X_i : i \in I)$ where $(X_n)_{n\geq 0}$ is the random walk on g started from $X_0 = \rho$. By definition of the entropy we thus have

$$\begin{aligned}
\mathbb{E}[\mathcal{H}(X_{i_1}, X_{i_2}, \dots, X_{i_\ell})] &= \mathbb{E}\left[-\log p^{i_1,\dots,i_\ell}_{(G,X_0)}(X_{i_1}, X_{i_2}, \dots, X_{i_\ell})\right] \\
&\underset{\text{Markov}}{=} \mathbb{E}\left[-\log\left(p^{i_1}_{(G,X_0)}(X_{i_1}) p^{i_2-i_1,\dots,i_\ell-i_1}_{(G,X_{i_1})}(X_{i_2},\dots,X_{i_\ell})\right)\right] \\
&= \mathbb{E}\left[-\log p^{i_1}_{(G,X_0)}(X_{i_1})\right] + \mathbb{E}\left[p^{i_2-i_1,\dots,i_\ell-i_1}_{(G,X_{i_1})}(X_{i_2},\dots,X_{i_\ell})\right] \\
&\underset{\text{stat.}}{=} \mathbb{E}\left[-\log p^{i_1}_{(G,X_0)}(X_{i_1})\right] + \mathbb{E}\left[p^{i_2-i_1,\dots,i_\ell-i_1}_{(G,X_0)}(\tilde{X}_{i_2-i_1},\dots,\tilde{X}_{i_\ell-i_1})\right] \\
&= \mathbb{E}[\mathcal{H}(X_{i_1})] + \mathbb{E}[\mathcal{H}(X_{i_2-i_1}, X_{i_3-i_1}, \dots, X_{i_\ell-i_1})].
\end{aligned}$$

□

Exercise 32. Let $(X_n)_{n\geq 0}$ be an irreducible Markov chain on a finite state $\mathcal{E}$ with stationary distribution $(\mu_i : i \in \mathcal{E})$ and transition probabilities $q_{i,j}$. Show that regardless of the starting point of the chain we have

$$\frac{1}{n}\mathcal{H}(X_1,\dots,X_n) \xrightarrow[n\to\infty]{} \sum_{i\in\mathcal{E}} \mu_i \sum_{j\in\mathcal{E}} -q_{i,j}\log q_{i,j}.$$

(Hint: Show that when the Markov chain starts from stationarity we actually have equality for all $n \geq 1$ in the above display).

Let us derive a few useful consequences of the last proposition. To simplify notation we write $h_I = \mathbb{E}[\mathcal{H}(X_i : i \in I)]$ for a finite subset $I \subset \mathbb{N}$ and write $h_n = \mathbb{E}[\mathcal{H}(X_n)]$. First for any $n, m \geq 0$ we have $\mathcal{H}(X_n, X_{n+m}) \geq \mathcal{H}(X_{n+m})$ and so after taking expectation and applying Proposition 34 we deduce that

$$h_{n+m} \leq h_n + h_m.$$

The sequence $(h_n)_{n \geq 0}$ is then sub additive and so by Fekete's lemma h_n/n converges as $n \to \infty$, and this limit (equal to the infimum of all h_n/n) is denoted by $h \in [0, \infty]$ and called the *mean entropy* of $\vec{G}$. In the following we will always assume that h_1 is finite (this is a very mild assumption and follows for example from the hypothesis $\mathbb{E}[\log(\deg(\vec{e})_*)] < \infty$).

Exercise 33. Let $\vec{G}$ be an ergodic stationary random graph such that $\mathbb{E}[\log(\deg(\vec{e})_*)] < \infty$. Show (using Kingman sub additive theorem) that $-\log p^n_{(G,X_0)}(X_n)$ converges almost surely and in L^1 towards $h \in [0, \infty)$ as $n \to \infty$.

Next, if one applies repeatedly Proposition 34 to h_I for $I = \{0 < i_1 < \ldots < i_\ell\}$ one deduces that letting $i_0 = 0$ we have

$$h_I = \sum_{j=1}^{\ell} h_{i_j - i_{j-1}}.$$

In particular, if one considers conditional entropy, for any $1 \leq k \leq m \leq n$ we have

$$\begin{aligned}
\mathbb{E}[\mathcal{H}(X_1, \ldots, X_k \mid X_m, \ldots, X_n)] &= \mathbb{E}[\mathcal{H}(X_1, \ldots, X_k, X_m, \ldots, X_n) - \mathcal{H}(X_m, \ldots, X_n)] \\
&= h_{\{1,2,\ldots,k,m,m+1,\ldots,n\}} - h_{\{m,m+1,\ldots,n\}} \\
&= kh_1 + h_{m-k} + (n-m)h_1 - (h_m + (n-m)h_1) \\
&= kh_1 + h_{m-k} - h_m.
\end{aligned}$$

Theorem 35 (***Tail σ-field and zero entropy***)

Let $\vec{G}$ be a stationary random graph such that $h_1 < \infty$. Then the mean entropy h is zero if and only if the tail σ-algebra $\mathcal{T}$ relative to the random walk on $\vec{G}$ is almost surely trivial (in particular G is almost surely Liouville).

Proof. We examine the expected conditional entropy of $\mathbb{E}[\mathcal{H}(X_1 \mid X_m, \ldots, X_n)]$ when $1 < m < n$. By the above calculation this is equal to $h_1 + h_{m-1} - h_m$ and in particular does not depend upon n (this is also true without taking expectation and follows from Markov property at time n). If we let $n \to \infty$ and then $m \to \infty$, we get by monotonicity of conditional entropy (Proposition 32) and monotone convergence that

$$\begin{aligned}
h_1 + h_{m-1} - h_m &= \mathbb{E}[\mathcal{H}(X_1 \mid \sigma(X_m, \ldots)] \\
&\underset{m\to\infty}{\nearrow} \mathbb{E}[\mathcal{H}(X_1 \mid \mathcal{T})],
\end{aligned}$$

in particular $h_m - h_{m-1}$ is decreasing and it must converge to h by Cesaro's summation. We finally deduce that $\mathbb{E}[\mathcal{H}(X_1 \mid \mathcal{T})] = h_1 - h$. Similarly we get $\mathbb{E}[\mathcal{H}(X_1, X_2, ..., X_k \mid \mathcal{T})] = k(h_1 - h)$. If $h = 0$, it means that $\mathcal{H}(X_1, X_2, ..., X_k \mid \mathcal{T}) = \mathcal{H}(X_1, X_2, ..., X_k)$ almost surely and so by Exercise 1 that $\mathcal{F}_k$ is independent of $\mathcal{T}$, and so $\mathcal{T}$ is trivial. Reciprocally, if $\mathcal{T}$ is trivial then $\mathcal{H}(X_1 \mid \mathcal{T}) = 0$ almost surely and in expectation (recall that $h_1 < \infty$). □

4.3.2 Volume and speed

Using the inequality $\mathcal{H}(X) \le \log \#\mathrm{Supp}(X)$ we deduce that $h_n \le \mathbb{E}[\log \#[(G, X_0)]_n]$ and in particular if G has subexponential growth in the sense that $n^{-1}\mathbb{E}[\log \#[(G, X_0)]_n] \to 0$ then its mean entropy is automatically 0 and in particular it is Liouville (this gives another proof of the fact that $\mathbb{Z}^d$ is Liouville for any $d \ge 1$). Actually, even if the volume of balls growth exponentially with the radius the graph may still be Liouville as soon as the speed of the random walk is zero:

Corollary 36. *Let $\vec{G}$ be a stationary and ergodic random graph of degree almost surely bounded by $M > 0$. Conditionally $\vec{G}$ let $(X_n)_{n\ge 0}$ be the vertices visited by a simple random walk on G starting from the origin. We denote* the speed *of the random walk by s and* the exponential volume growth *of G by v, namely*

$$\begin{aligned} s &= \limsup_{n\to\infty} n^{-1}\mathbb{E}\left[\mathrm{d_{gr}}(X_0, X_n)\right], \\ v &= \limsup_{n\to\infty} n^{-1}\mathbb{E}\left[\log(\#[\vec{G}]_n)\right]. \end{aligned}$$

Then the mean entropy h of $\vec{G}$ satisfies $h \le vs$.

Remark 2. We even have the inequality $h \ge s^2/2$, but the proof of it is based on a deep inequality called the Carne-Varopoulos inequality and so we do not present the proof here.

Proof. Fix $\varepsilon > 0$. To simplify notation, we write B_s for the ball of radius $(s+\varepsilon)n$ around the origin in the graph G and B_s^c for the complement of the ball of radius n in B_s. We decompose the entropy $\mathcal{H}(X_n)$ as follows

$$\begin{aligned} \mathcal{H}(X_n) &= \mathbb{E}[-\log p^n_{(G,X_0)}(X_n)] \\ &= \mathbb{E}\left[-\log p^n_{(G,X_0)}(X_n)1_{X_n\in B_s}\right] + \mathbb{E}\left[-\log p^n_{(G,X_0)}(X_n)1_{X_n\in B_s^c}\right] \\ &\le \mathbb{E}\left[-\log p^n_{(G,X_0)}(X_n)1_{X_n\in B_s}\right] + n\log(M)\mathbb{P}\left(X_n \in B_s^c\right). \end{aligned}$$

Because $p^n_{(\mathrm{g},\rho)}(x) \le M^{-n}$ in any graph with degrees bounded by M. On the other hand, by ergodicity it follows that $n^{-1}\mathrm{d_{gr}}(X_0, X_n) \to s$ almost surely and hence $\mathbb{P}(X_n \in B_s) \to 1$. It follows from the last display that

$$\mathbb{E}[\mathcal{H}_n] \le \mathbb{E}\left[-\log p^n_{(G,X_0)}(X_n) \mid X_n \in B_s\right](1 + o(1)) + o(n) \le \mathbb{E}[\log(B_s)](1 + o(1)),$$

and the latter is eventually less than sv as desired. □

Bibliographical references. The concept of entropy of a random variable has been introduced by Shannon in the last century in order to theorize channel coding. It has since then been used in many context. The use of the entropy to characterize the behavior of random walks on regular graphs has its origin in the work of Avez [9] and has been later developed by Kaimanovich [31, 29, 30]. The extension to stationary random graphs is taken from [11], see also [20]. Basics on entropy of random variables and tail behavior of random walk can be found in many textbooks but is here largely inspired by [37].

Chapter 5

Stationary random trees

In this chapter we focus on the case when the underlying graph is a stationary/stationary reversible/unimodular random tree and use this geometric structure to our advantage. We start by giving an example of stationary and reversible random tree based on the standard Galton–Watson trees. We also prensent the local limit of large critical Galton–Watson trees known as Kesten's tree.

5.1 Augmented Galton–Watson trees

We now introduce a very important example of stationary and reversible random graph based on Galton–Watson trees. As it will turn out, Galton–Watson trees are generally not stationary and reversible random graphs and we need to trick a little their distributions to gain stationarity (and reversibility) along the simple random walk.

5.1.1 Definition

Recall from Section 1.3 the definition of a Galton–Watson plane tree. Clearly a plane tree τ can be seen as a pointed graph by forgetting the order and distinguishing the ancestor vertex. It can also be seen as a rooted graph by distinguishing the edge $\varnothing \to 1$ provided of course that τ has more than two vertices. We denote $\pi_{\circlearrowleft}(\tau)$ the obtained rooted graph.

It should be clear that the origin vertex $\varnothing$ of a Galton–Watson tree has stochastically fewer neighbors than the other vertices since it has no ancestor. To cope with this problem we introduce the augmented p-Galton–Watson tree (AGW) obtained by grafting two independent p-GW trees at each extremity of the oriented root edge $(\varnothing, 1)$, see figure below. Equivalently, the offspring distribution of the root vertex is changed from $(p_k)_{k\geq 0}$ to $(0, p_0, p_1, ...)$ whereas the offspring distribution of the other vertices is still p.

Theorem 37 (*Augmented Galton–Watson tree*)

The image by $\pi_{\circlearrowleft}$ of an augmented p*-AGW tree is a stationary (and reversible) random graph.*

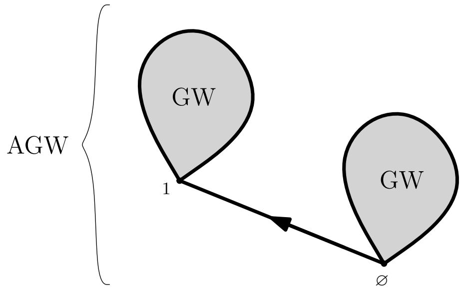

In particular we can get from an AGW tree a unimodular random graph after forgetting the orientation, distinguishing the origin vertex and biasing by $(\deg(\varnothing))^{-1}$.

Proof. We will directly show that the rooted plane AGW tree (with the plane orientation) is itself stationary and reversible. By forgetting the plane ordering we will have the desired statement. Fix $k, \ell \geq 0$ and measurable subsets of trees $A_1, ..., A_k, B_1, ..., B_\ell$. We form the event $\mathcal{E}$ represented on the figure on the right. By the definition of the augmented Galton–Watson probability we get that

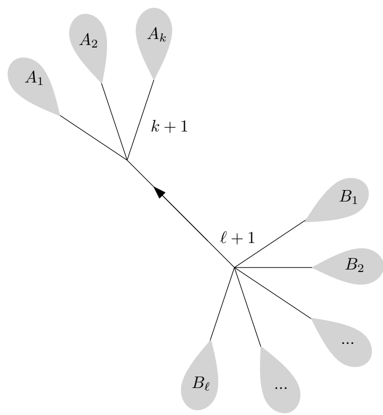

$$\mathbb{P}_{\mathrm{AGW}}(\mathcal{E}) = p_k p_\ell \prod_{i=1}^{k} \mathbb{P}_{\mathrm{GW}}(A_i) \prod_{i=1}^{\ell} \mathbb{P}_{\mathrm{GW}}(B_i).$$

We then compute the probability that the tree re-rooted on the first edge (after the root edge) taken by a simple random walk falls in $\mathcal{E}$ and we find:

$$\underbrace{p_\ell \frac{\ell}{\ell+1} \prod_{i=1}^{\ell} \mathbb{P}_{\mathrm{GW}}(B_i) p_k \prod_{i=1}^{k} \mathbb{P}_{\mathrm{GW}}(A_i)}_{\text{if } \vec{E}_1 \neq \overleftarrow{E}_0} + \underbrace{p_\ell \frac{1}{\ell+1} p_k \prod_{i=1}^{\ell} \mathbb{P}_{\mathrm{GW}}(B_i) \prod_{i=1}^{k} \mathbb{P}_{\mathrm{GW}}(A_i)}_{\text{if } \vec{E}_1 = \overleftarrow{E}_0}.$$

Hence the probability of $\mathcal{E}$ is the same for the original plane rooted tree as for the one obtained by rooting after one step of random walk. Events of the type of $\mathcal{E}$ form a π-system and they generate the full σ-algebra. We conclude by an application of the monotone class theorem. □

Exercise 34. Suppose $\mathbb{E}[\mathrm{p}] > 1$. Show that an augmented Galton–Watson tree conditioned to be infinite is a stationary and reversible random graph. Compute the degree distribution of its origin vertex.

Example : Geometric Galton–Watson. Let $\mathrm{p}^{(\xi)}$ be the geometric offspring distribution with success parameter $\xi \in (0,1)$, that is

$$p_k^{(\xi)} = \xi(1-\xi)^k.$$

An easy computation shows that $\mathrm{p}^{(\xi)}$ indeed is a probability distribution whose generating function is given by $\mathrm{F}_{\mathrm{p}^{(\xi)}}(z) = \xi(1-(1-\xi)z)^{-1}$. In particular, the mean of $\mathrm{p}^{(\xi)}$ is $(1-\xi)/\xi$ and

$\mathrm{p}^{(\xi)}$ is supercritical as soon as $\xi < 1/2$ and the extinction probability is the smallest solution of $z = \mathrm{F}_{\mathrm{p}^{(\xi)}}(z)$ in $[0,1]$ and thus equal to

$$\mathrm{q}_{\mathrm{die}}^{(\xi)} = \frac{\xi}{1-\xi}, \qquad \text{for } \xi \in (0,1/2].$$

We denote in the following T_ξ a $\mathrm{p}^{(\xi)}$-Galton–Watson tree. We already convinced ourselves that T_ξ is not a stationary and reversible random graphs and constructed the associated augmented Galton–Watson tree denoted by $\hat{T}_\xi$ which is stationary and reversible. In the case of the geometric distribution, this augmented Galton–Watson has a particularly simple interpretation:

Proposition 38. *For $\xi \in (0,1/2)$, the $\mathrm{p}^{(\xi)}$-Galton–Watson tree T_ξ* conditioned to be infinite *is a stationary and reversible random graph (after taking the image by $\pi_{\circlearrowleft}$). Furthermore this law coincides with the law of the augmented $\mathrm{p}^{(\xi)}$-Galton–Watson tree $\hat{T}_\xi$ conditioned to be infinite.*

Proof. We will show that T_ξ conditioned to be infinite is equal in law to $\hat{T}_\xi$ conditioned to be infinite. Note that since $\hat{T}_\xi$ is (in its random rooted graph version) a stationary and reversible random graph, the property passes to the conditioning and we get the statement of the proposition. Let us compute first the degree distribution of the ancestor in T_ξ, we have

$$\mathbb{P}(c_\varnothing(T_\xi) = k_0 \mid \#T_\xi = \infty) = \frac{\xi(1-\xi)^{k_0}\left(1-(\mathrm{q}_{\mathrm{die}}^{(\xi)})^{k_0}\right)}{1-\mathrm{q}_{\mathrm{die}}^{(\xi)}} = \frac{\xi(1-\xi)\left((1-\xi)^{k_0}-\xi^{k_0}\right)}{1-2\xi}.$$

Furthermore, conditionally on $c_\varnothing(T_\xi) = k_0$ and T_ξ being infinite, the k_0 trees above $1,2,...,k_0$ have the law of k_0 independent $\mathrm{p}^{(\xi)}$-Galton–Watson trees conditioned on having at least one tree being infinite. The same is true in $\hat{T}_\xi$: conditionally on $\hat{T}_\xi$ being infinite and on $c_\varnothing(\hat{T}_\xi) = k_0$ the k_0 subtrees above $1,2,...,k_0$ have the above law and we have

$$\mathbb{P}(c_\varnothing(\hat{T}_\xi) = k_0 \mid \#\hat{T}_\xi = \infty) = \frac{\xi(1-\xi)^{k_0-1}\left(1-(\mathrm{q}_{\mathrm{die}}^{(\xi)})^{k_0}\right)}{1-(\mathrm{q}_{\mathrm{die}}^{(\xi)})^2} = \frac{\xi(1-\xi)\left((1-\xi)^{k_0}-\xi^{k_0}\right)}{1-2\xi}.$$

Comparing the last two displays, we realized that T_ξ and $\hat{T}_\xi$ have the same law once conditioned on being infinite. This completes the proof. □

5.1.2 Computation of the speed

In this section we study the behavior of the simple random walk on supercritical augmented Galton–Watson trees. For simplicity we focus on the case when $p_0 = 0$ i.e. when all individuals have at least one child and $p_1 \neq 1$. In particular, the random trees considered are all infinite. We begin with a $0-1$ law for transience of Galton–Watson trees.

Proposition 39. *Let $\vec{T} = (T, \vec{E})$ be (the random rooted graph obtained from) a Galton–Watson tree with offspring distribution $(p_k)_{k\geq 0}$ with $p_1 \neq 1$ and $p_0 = 0$. Then $\mathbb{P}(T$ is transient$) \in \{0,1\}$ and the same is true for the associated augmented Galton–Watson tree.*

Proof. Obviously if the offspring distribution is subcritical or critical, the random tree T is almost surely finite and thus recurrent. We can thus focus on the supercritical case. A tree is transient if and only one of its subtrees above the ancestor is transient. Hence the probability q that T is recurrent satisfies $q = \sum_{k \geq 1} p_k q^k = \mathrm{F_p}(q)$ whose unique solutions are 0 or 1. The same property holds for the augmented Galton–Watson tree since the latter is composed of the grafting of two Galton–Watson trees through a single edge. □

Theorem 40 (*Speed in AGW*)

Let $\vec{T}$ be (the random rooted graph obtained from) an augmented Galton–Watson tree with offspring distribution $(p_k)_{k \geq 0}$ satisfying $p_0 = 0$ and $p_1 \neq 1$. We denote by $(X_i)_{i \geq 0}$ the vertices visited by the simple random walk on T then

- $\mathrm{d_{gr}}(X_0, X_i) \to \infty$ *almost surely (T is almost surely transient),*
- *we have*
$$\lim_{n \to \infty} n^{-1} \cdot \mathbb{E}[\mathrm{d_{gr}}(X_0, X_n)] = \sum_{k \geq 1} p_k \frac{k-1}{k+1}.$$

Remark 3. A similar statement holds for offspring distribution with $p_0 \neq 0$ but when $\sum_k k p_k > 1$ and the AGW is conditioned to be infinite. Also, the convergence in mean of $n^{-1}\,\mathrm{d_{gr}}(X_0, X_n)$ actually holds in a almost sure sense, a fact that follows from (3.5) once we know that an augmented Galton–Watson tree is ergodic (see [37, Section 16.3] for a proof).

Proof. We denote by $X_0, X_1, \ldots$ the vertices visited iteratively by the simple random walk with the convention that $(X_0, X_1) = \vec{E}_0$ is the root edge of $\vec{T}$. For $n \geq 0$ we put $H_n = \mathrm{d_{gr}}(X_0, X_n)$ and want to prove that $n^{-1} \cdot H_n \to s$ in mean. To do so we first associate a martingale to H_n in a standard way. Indeed, conditionally on $\vec{T}$ we denote by $\mathcal{F}_n$ the filtration generated by $\{X_0, \ldots, X_n\}$ and put
$$M_n = H_n - \sum_{k=1}^{n} \mathrm{E}[H_k - H_{k-1} \mid \mathcal{F}_{k-1}].$$
(We use E for expectation over the random walk and $\mathbb{E}$ for averaging over the underlying random graph). By construction, conditionally on $\vec{T}$, the process (M_n) is a martingale. Taking the average over the random walk and the underlying tree we get that
$$\mathbb{E}[\mathrm{E}[\mathrm{d_{gr}}(X_0, X_n)]] = \underbrace{\mathbb{E}[\mathrm{E}[M_n]]}_{=0} + \sum_{k=1}^{n} \mathbb{E}[\mathrm{E}[H_k - H_{k-1} \mid \mathcal{F}_{k-1}]]. \tag{5.1}$$

It turns out that, conditionally on $\vec{T}$, since the walk displaces on a *tree* we can explicitly express the conditional expectation $\mathrm{E}[H_{n+1} - H_n \mid \mathcal{F}_n]$: Indeed, if the walk is sitting on a vertex $X_n \neq X_0$ of degree k, there are $k-1$ edges leading to vertices further apart from X_0 by one unit and a unique edge yielding to a vertex closer to X_0 by one unit, so

$$\begin{aligned} \mathrm{E}[H_{n+1} - H_n \mid \mathcal{F}_n] &= \mathbb{1}_{X_n \neq X_0} \frac{\deg(X_n) - 2}{\deg(X_n)} + \mathbb{1}_{X_n = X_0} \\ &= \frac{\deg(X_n) - 2}{\deg(X_n)} + \frac{2}{\deg(X_n)} \mathbb{1}_{X_0 = X_n}. \end{aligned} \tag{5.2}$$

Since the underlying random graph is stationary (we do not use reversibility in this proof) we get that for all $n \geq 0$

$$\mathbb{E}\left[\mathrm{E}\left[\frac{\deg(X_n) - 2}{\deg(X_n)}\right]\right] = \sum_{k \geq 1} p_{k-1} \frac{k-2}{k} = \sum_{k \geq 0} p_k \frac{k-1}{k+1} =: s.$$

Combining this calculation with (5.2) and (5.1) we deduce that $\mathbb{E}[\mathrm{E}[H_n]] \geq n \cdot s$. From this we first deduce that $\vec{T}$ is almost surely transient. Indeed, we know from Proposition 39 that $\vec{T}$ is either almost surely transient or almost surely recurrent. If it is almost surely recurrent, using (3.7) we get that range of the walk is sublinear in mean in particular we would have

$$\lim_{n \to \infty} n^{-1} \, \mathbb{E}[\mathrm{E}[H_n]] = 0,$$

which is incompatible with $\mathbb{E}[\mathrm{E}[H_n]] \geq n \cdot s$. Consequently we almost surely have $H_n \to \infty$ and bootstrapping in (5.2) we deduce that the probability that $X_0 = X_n$ tends to zero as $n \to \infty$ and we get the desired convergence

$$\frac{\mathbb{E}[\mathrm{E}[H_n]]}{n} \to s.$$

□

Exercise 35. Show that the speed of the random walk on a Galton–Watson tree with offspring distribution satisfying $p_0 = 0$ and $p_1 \neq 1$ exists and is the same as in the associate augmented Galton–Watson tree.

Exercise 36. Do you have any easier way to prove that supercritical Galton–Watson trees with $p_0 = 0$ are almost surely transient?

Remark that since the function $x \mapsto \frac{x-1}{x+1}$ is concave we have $\sum_{k \geq 0} p_k \frac{k-1}{k+1} \leq \frac{m-1}{m+1}$ where m is the mean of the offspring distribution. Hence, the speed on a Galton–Watson tree is lower than the speed on a regular tree of the same growth. Do you have a heuristic explanation for this fact?

5.2 General results

We now extend some of the above results to the case of general unimodular/stationary reversible random trees. We also discuss the status of Open Question 5 in the case of random trees.

5.2.1 Geometric properties

Theorem 41 (***Recurrence, zero speed, mean degree, ends...***)

Let $(T, \vec{E})$ be a stationary random tree a.s. infinite whose origin vertex is denoted by ρ. Then the following four assertions are equivalent

(i) $\mathbb{E}[\deg(\rho)^{-1}] = 1/2$,

(ii) the simple random walk a.s. has zero speed on T,

(iii) T is almost surely recurrent,

(iv) T almost surely has 1 or 2 ends.

In particular, such stationary random trees are also reversible by Theorem 30.

Proof. With the same notation as in the proof of Theorem 25 or 40 we have

$$\mathrm{E}[\mathrm{d_{gr}}(X_0, X_{n+1}) - \mathrm{d_{gr}}(X_0, X_n) \mid \mathcal{F}_n] = \frac{\deg(X_n) - 2}{\deg(X_n)} + \frac{2}{\deg(X_n)} 1_{X_0 = X_n}.$$

Recall also that by stationarity we have

$$\mathbb{E}\left[\mathrm{E}\left[\frac{\deg(X_n) - 2}{\deg(X_n)}\right]\right] = \mathbb{E}\left[\frac{\deg(\rho) - 2}{\deg(\rho)}\right].$$

On the other hand, since the invariant measure of the walk on T is proportional to the degree we have $n^{-1} \sum_{i=0}^{n-1} 1_{X_0 = X_i} \to 0$ in P-probability. Hence we deduce that

$$\text{speed} = \lim_{n\to\infty} n^{-1} \mathbb{E}[\mathrm{E}[\mathrm{d_{gr}}(X_0, X_n)]] = \frac{\mathbb{E}[\deg(\rho) - 2]}{\mathbb{E}[\deg(\rho)]},$$

and so (i) is equivalent to (ii). Also $(iv) \Rightarrow (iii) \Rightarrow (ii)$ is easy to prove. Indeed, if (iv) holds then T is made of a single infinite or bi-infinite line on which finite trees are grafted and as far as recurrence is concerned, a simple random walk on such a tree behaves like the simple random walk on $\mathbb{N}$ or $\mathbb{Z}$ which in both cases is recurrent. If we suppose (iii) that T is almost surely recurrent then the range of the walk is zero by (3.7) and so the speed of the SRW on T is zero and so (ii) holds. Let us now focus on $(ii) \Rightarrow (iv)$. We can suppose that T is ergodic (this is a standard procedure in ergodic theory: if a property A holds almost surely for the stationary measure then it holds also almost surely for almost all –with respect to the Choquet integral– ergodic components of the measure). In order to prove that T has almost surely 1 or 2 ends, we will use the tree $(\tilde{T}, \tilde{\rho})$ obtained by biasing T by the inverse of the origin degree. If we can prove that T is reversible on top of being stationary then $\tilde{T}$ will be unimodular by Proposition 19. Since $\mathbb{E}[\deg(\tilde{\rho})] = 2$ we could apply Theorem 13 and deduce that $\tilde{T}$ has one or two ends almost surely. The same obviously would hold for T and the proof would be complete. It thus remains to show that T is reversible. To do this, we proceed as in the proof of Theorem 30 and suppose by contradiction that it is not reversible. Then there exists $\alpha \neq \beta$ and a bounded positive mesurable function F such that by ergodicity we have

$$\frac{1}{n}\sum_{i=0}^{n-1} F(T,\vec{E}_i) \xrightarrow[n\to\infty]{a.s.} \alpha \quad \text{and} \quad \frac{1}{n}\sum_{i=0}^{n-1} F(T,\overleftarrow{E}_i) \xrightarrow[n\to\infty]{a.s.} \beta.$$

Now since we are on a tree, the set of edges visited by the walk from $\vec{E}_0$ to $\vec{E}_n$ are almost the same as those visited from $\overleftarrow{E}_n$ to $\overleftarrow{E}_0$. Those edges which differ are the edges which lie on the geodesic path from X_0 to X_n, but since T has zero speed by assumption there are only $o(n)$ such edges. Since F is bounded, we can use this fact together with the last display to deduce that $\alpha = \beta$. Absurd! □

Exercise 37. Prove directly that a recurrent stationary random tree must have 1 or 2 ends without passing through its unimodular version and requiring the mass-transport principle used in Theorem 13.

Exercise 38. Give an example of a recurrent (equivalently zero speed by the above theorem) stationary (and reversible by Theorem 30) random tree T such that the exponential volume growth is strictly positive

$$\liminf_{n\to\infty} n^{-1}\log(\#[T]_n) > 0.$$

5.2.2 Approximations by finite random graphs

Theorem 42 (***Soficity for unimodular random trees***)

Any unimodular random tree (T,ρ) is the local limit in distribution of a sequence of uniformly pointed random graphs.

We do not give the proof and refer to [19, 24, 13, 18]. Notice that in the above statement nothing is said about the geometry of the finite sequence approximating (T,ρ) and they may very well not be trees. Indeed, if (T_n,ρ_n) is a sequence of random finite uniformly pointed trees which converges locally in distribution towards an infinite random pointed tree (T_∞,ρ_∞), then using Theorem 13 as well as Fatou's lemma, we deduce that

$$\mathbb{E}[\deg(\rho)] \le 2.$$

Using Theorem 13 again it follows that T_∞ has at most two ends. Reciprocally, those random trees are approximated by finite random trees: Consider inside T_∞ a bond Bernoulli percolation of parameter $p \in (0,1)$ and denote by (τ_p,ρ) the cluster of the origin, pointed at the origin. We have seen that (τ_p,ρ) is indeed a unimodular random graph. It thus suffices to see that it is almost surely finite in order to apply Proposition 7. But recall that since $\mathbb{E}[\deg(\rho_\infty)] \le 2$, it follows from Theorem 13 that T_∞ almost surely has $0,1$ or 2 ends. It is thus clear that any percolation with parameter $p<1$ on such a tree creates only finite clusters a.s. Translating this in the realm of stationary random trees we have proved:

Corollary 43. *The class of stationary trees considered in Theorem 41 is exactly those random trees which are local limit of finite stationary random trees.*

Remark 4. In order to compare with the forthcoming theorem of Benjamini & Schramm on recurrence of local limits of planar graphs let us draw the easy corollary of the above discussion: If (T_n, ρ_n) is a sequence of finite unimodular random trees converging locally in law towards (T, ρ) then T is almost surely recurrent.

5.2.3 A few constructions

A stationary tree not reversible. We begin with the construction of a random tree which is stationary but not reversible (and hence not unimodular once biaised by the inverse of the root degree). As explained in Theorem 30 it must be transient. We start with the fixed tree T_0 below obtained where each vertex at generation equal to 0 mod 3 has 1 child, each vertex at generation equal to 1 mod 3 has 2 children and each vertex at generation equal to 2 mod 3 has three children. The generations are indexed by $\mathbb{Z}$. It is an easy exercice to see that the simple

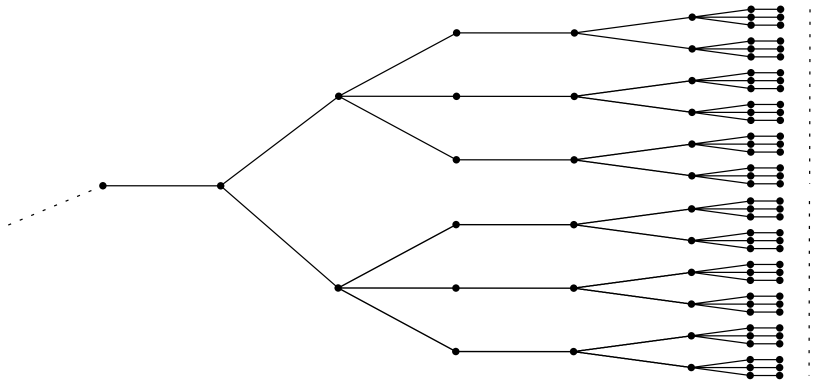

Figure 5.1: The $1-2-3$ tree T_0.

random walk on T_0 is transient. We claim that there is a random choice of the root edge so that $(T_0, \vec{E})$ is stationary. One abstract way to see this is to consider a subsequential limit of the probability measure on $\mathcal{G}^{\rightarrow}$ given by

$$\mu_n = \frac{1}{n} \sum_{i=0}^{n-1} \delta_{(T_0, (\vec{E}_{i+k})_{k \geq 0})},$$

where, as usual, we write $(\vec{E}_i)_{i \geq 0}$ for the oriented edges visited by the random walk on T_0. The fact that this sequence of probability measure is tight follows from the fact that the degrees in T_0 are all bounded by 4. Clearly if θ is the shift operator on $\mathcal{G}^{\rightarrow}$ we have

$$\mathrm{d_{LP}}(\mu_n \circ \theta^{-1}, \mu_n) \leq \frac{2}{n},$$

where $\mathrm{d_{LP}}$ is the Lévy Prokhorov distance. It follows that if μ is any subsequential limit of $(\mu_n)_{n \geq 0}$ we must have $\mu \circ \theta^{-1} = \mu$. In order words the random rooted graph associated to μ is stationary, which is what we desired since it corresponds to the tree T_0 with a random choice of root edge (they are only finitely many choices of root edge which give different rooted trees).

Exercise 39. Give explicitly a choice of a root edge on T_0 to turn it into a stationary random graph.

We claim that the above random rooting of T_0 turns it into a stationary but not reversible random graph. Indeed, since T_0 is transient the walk escapes to infinity and during its run the walker will successively visit vertices of degree $2, 3, 4$ more often in this order than in the order $4, 3, 2$ breaking the reversibility. We leave the reader turn this heuristic explanation into a proof.

Arbitrary degree distribution. The augmented Galton–Watson tree construction enables us to produce stationary random trees with an arbitrary degree distribution of the origin. Using the inverse biasing, one can get a unimodular random tree with an arbitrary degree distribution provided that it has a finite mean.

5.3 Local limit of critical Galton–Watson trees

In this chapter $\mathrm{p} = (p_0, p_1, \ldots)$ is a critical offspring distribution $\sum_{k \geq 0} k p_k = 1$ such that $p_1 \neq 1$. We denote by T a p-Galton–Watson tree which is thus an almost surely finite plane tree. We are interested in the local convergence of large critical Galton–Watson trees towards an infinite random tree called "the Galton–Watson tree T conditioned to be infinite". Of course since T is almost surely finite the last conditioning is degenerate and necessitates some care to be properly defined.

5.3.1 Kesten's tree

In this section we introduce an infinite random tree which will be shown later to be the local limit of critical Galton–Watson tree conditioned to be large in some sense.

Martingale construction

Let T be a critical Galton–Watson tree. For $n \geq 0$, recall that $[T]_n$ is the tree obtained by restricting to the first n generations of T and denote by $(\mathcal{F}_n)$ the filtration generated by these variables. We write $\#\partial[T]_n$ for the number of vertices at generation n of the tree T. Then by the branching structure of Galton–Watson trees and the criticality of p we get that

$$\#\partial[T]_n \text{ is an } \mathcal{F}_n\text{-martingale.}$$

Obviously this is not a uniformly integrable martingale since it starts from $\#\partial[T]_0 = 1$ and eventually yields the value 0 because the tree is almost surely finite (recall that $p_1 \neq 1$). We will use this martingale as a Radon–Nikodym derivative to create a new probability distribution on plane trees.

More precisely for each $n \geq 0$ define (the law of) a random plane tree τ_n by putting for any measurable non-negative function f on the space of plane trees

$$\mathbb{E}[f(\tau_n)] = \mathbb{E}[\#\partial[T]_n f([T]_n)].$$

Since $\#\partial[T]_n$ is a non-negative random variable of expectation 1, the last display makes sense and indeed defines (the law of) a random plane tree. Now, the martingale property of $\#\partial[T]_n$ translates into a coherence relation between the τ_n in the sense that $[\tau_n]_m = \tau_m$ in distribution for any $0 \le m \le n$:

$$\mathbb{E}[f([\tau_n]_m)] \underset{\text{def.}}{=} \mathbb{E}[\#\partial[T]_n f([[T]_n]_m)] = \mathbb{E}[\#\partial[T]_n f([T]_m)] \underset{\text{mart.}}{=} \mathbb{E}[\#\partial[T]_m f([T]_m)] = \mathbb{E}[f(\tau_m)].$$

This enables us to define a random tree τ such that $[\tau]_n = \tau_n$ in distribution for each n (to be precise, the coherence relation shows that τ_n converges locally in distribution toward a random plane tree τ; one then check that $[\tau]_n = \tau_n$ in distribution for every n). Notice that τ is almost surely an infinite random tree since almost surely τ_n has descendant at the generation n. The random tree τ will be later called *the* p-*Galton–Watson tree conditioned to be infinite* or *Kesten's infinite* p-*Galton–Watson tree*. Its law is characterized by the fact that for any plane tree t_0 of height n we have

$$\mathbb{P}([\tau]_n = \mathsf{t}_0) = \mathbb{P}([T]_n = \mathsf{t}_0) \cdot \#\partial[\mathsf{t}_0]_n.$$

The above construction is quite general (see Exercise 22) and shows that the law of τ is obtained by biasing the standard Galton–Watson measure with a non-negative martingale. However, the above description is not very practical and we shall give another construction of τ in the next section.

Description of τ as a Galton–Watson tree with immigration

Recall that p is the offspring distribution of T which is supposed to be critical. We denote by $\overline{\mathrm{p}}$ the size biaised distribution obtained by putting for $k \geq 0$

$$\overline{p}_k = k p_k,$$

which is indeed a probability distribution thanks to the criticality of p. Notice that $\overline{\mathrm{p}}$ is supported by (strictly) positive integers. We now construct a random infinite tree $\tilde{\tau}$ which is the genealogical tree made of two sorts of particles : standard and mutant particles. Initially there is only one mutant particle. All particles reproduce independently of each other, and standard particles produce a random number of standard particles distributed as p. Mutant particles however, reproduce according to $\overline{\mathrm{p}}$. Among the descendant of a mutant particle, a uniform child is picked (independently of the past) and is declared "mutant" whereas the other children are standard particles. Clearly, in the construction of the plane tree $\tilde{\tau}$ there is a distinguished infinite ray corresponding to the genealogical line of the mutant particles (there is exactly one mutant particle at each level of the tree). We call this distinguished line the "spine" of $\tilde{\tau}$. By the above description, all the trees hanging off this spine are independent p-Galton–Watson trees and so are all a.s. finite.

Proposition 44. *The random infinite tree $\tilde{\tau}$ has the same law as τ.*

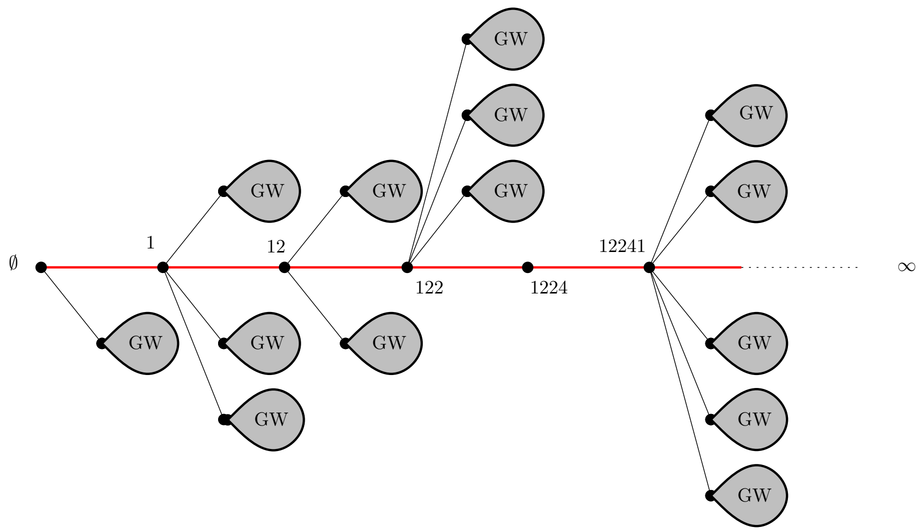

Figure 5.2: The construction of the tree $\tilde{\tau}$ from a spine of mutant particles (in red) and p-Galton–Watson trees produced by the standard children of the mutants.

Proof. Let us fix a plane tree t_0 with height n. We need to show that

$$\mathbb{P}([\tilde{\tau}]_n = \mathrm{t}_0) = \mathbb{P}([T]_n = \mathrm{t}_0) \cdot \#\partial[\mathrm{t}_0]_n.$$

If $x \in \partial[\mathrm{t}_0]_n$ is a vertex at height n, we say that we have $\tilde{\tau} = \mathrm{t}_0^{(x)}$ if in the above construction of $\tilde{\tau}$ the spine at height n exactly goes through the vertex x and the tree $\tilde{\tau}$ agrees with t_0 on the first n levels. Using the construction of $\tilde{\tau}$ we see that

$$\mathbb{P}([\tilde{\tau}]_n = \mathrm{t}_0^{(x)}) = \prod_{y \in \mathrm{t}_0 \backslash \partial[\mathrm{t}_0]_n \backslash [[\varnothing, x]]} p_{c_y} \prod_{z \in [[\varnothing, x]] \backslash \{x\}} \overline{p}_{c_z} \frac{1}{c_z},$$

where c_u is the number of children of the vertex u and $[[a, b]]$ is the geodesic line in the tree t_0 between vertices a and b. Using the definition of $\overline{\mathrm{p}}$, we see that the last display is equal to $\prod_{x \in \mathrm{t}_0 \backslash \partial[\mathrm{t}_0]_n} p_{c_x}$ which is exactly $\mathbb{P}([T]_n = \mathrm{t}_0)$. The desired result then follows by summing over all points $x \in \partial[\mathrm{t}_0]_n$. □

Exercise 40. Suppose that p has a second moment and denote $\sum p_k k(k-1) = \sigma^2$ its variance. We denote by T the p-Galton–Watson tree and by τ Kesten's infinite p-Galton–Watson tree. Recall that for any $n \geq 0$ we have $\mathbb{E}[\#\partial[T]_n] = 1$ by the martingale property.

1. Show that $\mathbb{E}[\#\partial[\tau]_n] = \mathbb{E}[\left(\#\partial[T]_n\right)^2]$.

2. Using the representation of τ given in Proposition 44 show that

$$\mathbb{E}[\#\partial[\tau]_n] = n \sum_{k=1}^{\infty} \overline{p}_k (k-1) = n\sigma^2 + 1.$$

3. Do you see another proof of the fact $\mathbb{E}[\left(\#\partial[T]_n\right)^2] = n\sigma^2 + 1$?

5.3.2 Local convergence

In this section, we show that the random tree τ naturally appears as the local limit of Galton–Watson trees conditioned to be large in some sense.

5.3.3 Trees with one spine

We denote by $\mathcal{T}_f$ the set of all finite plane trees and by $\mathcal{T}_1$ the set of all infinite plane trees with only one end (i.e. a unique infinite path starting from the origin of the tree). Note that $\tau \in \mathcal{T}_1$ almost surely. This is clear from the second construction of τ (Proposition 44) since all the trees grafted to the spine are a.s. finite. This remark is important since it restricts our state space a lot and thus to check convergence in distribution it is sufficient to check convergence on a much smaller set of events.

If t, s are plane trees and x is a leaf of t (that is a vertex with no child), we denote by $\mathsf{t} \circledast (\mathsf{s}, x)$ the tree obtained by grafting s on the vertex x of t, or formally the set $\{u \in \mathsf{t}\} \cup \{xv : v \in \mathsf{s}\}$. We also introduce the set

$$\mathcal{T}(\mathsf{t}, x) = \{\mathsf{t} \circledast (\mathsf{s}, x) : \mathsf{s} \text{ plane tree}\}.$$

These sets are nice since they generate the Borel σ-field on $\mathcal{T}_f \cup \mathcal{T}_1$.

Proposition 45. *Let θ_n for $n \geq 1$ and θ be random variables taking values in $\mathcal{T}_f \cup \mathcal{T}_1$ almost surely. For θ_n to converge in distribution for the local distance towards θ, it is sufficient to prove that for any $\mathsf{t} \in \mathcal{T}_f$ and any leaf x of t we have*

$$\lim_{n\to\infty} \mathbb{P}(\theta_n \in \mathcal{T}(\mathsf{t}, x)) = \mathbb{P}(\theta \in \mathcal{T}(\mathsf{t}, x)) \quad \text{and} \quad \lim_{n\to\infty} \mathbb{P}(\theta_n = \mathsf{t}) = \mathbb{P}(\theta = \mathsf{t}).$$

Proof. The family of all events of the form $\{\theta \in \mathcal{T}(\mathsf{t}, x)\}$ or $\{\theta = \mathsf{t}\}$ is a π-system and generate the Borel σ-field for the local topology on $\mathcal{T}_f \cup \mathcal{T}_1$. Hence by the monotone class theorem, the knowledge of a probability distribution on this class determines it completely. To see that these events form a convergence determining class, it is sufficient to check that every open set for the local topology on $\mathcal{T}_f \cup \mathcal{T}_1$ is obtained as a finite or countable union of those elements (see [15, Theorem 2.2]). This is easily checked. □

5.3.4 Conditioning a tree to be large

There are different ways to say that a tree is large : either by considering its number of vertices, its height, its number of leaves, or more exotic thoughts. Following [2], we unify these notions as follows. Let A be an integer-valued function which is finite on the set of all finite plane trees and which satisfies an "asymptotic additive property": for any finite tree t with a leaf x, as soon as $A(\mathsf{t} \circledast (\mathsf{s}, x))$ is large enough we have

$$A(\mathsf{t} \circledast (\mathsf{s}, x)) = A(\mathsf{s}) + D(\mathsf{t}, x), \tag{5.3}$$

where $D(\mathsf{t}, x)$ is some function of t and x. Let us right away give examples of such functions:

- $A(\mathsf{t})$ is the size (number of vertices) of t, in this case $D(\mathsf{t},x) = \#\mathsf{t} - 1$,

- $A(\mathsf{t})$ is the height of t, in this case $D(\mathsf{t},x) = \mathrm{Height}(x)$,

- $A(\mathsf{t})$ is the number of leaves of t, in this case $D(\mathsf{t},x) = \#\mathrm{Leaves}(\mathsf{t}) - 1$.

We also denote $\mathcal{A}_n$ the set of all trees in $\mathcal{T}_f \cup \mathcal{T}_1$ such that $A(\mathsf{t}) \in [n, n+n_0)$ where $n_0 \in \{1,2,3,\ldots\} \cup \{\infty\}$ is fixed. Usually we think of $n_0 = 1$ or $n_0 = \infty$.

Theorem 46

Let T_n be the random plane tree T conditioned on the event $T \in \mathcal{A}_n$ (we restrict our attention to the values of n such that the latter event has positive probability) then as soon as

$$\lim_{n\to\infty} \frac{\mathbb{P}(T \in \mathcal{A}_{n+1})}{\mathbb{P}(T \in \mathcal{A}_n)} = 1,$$

we have $T_n \to \tau$ in distribution for the local distance as $n \to \infty$.

Proof. Since τ is almost surely infinite and $\mathbb{P}(T_n = \mathsf{t}) \to 0$ for any finite plane tree t, using Proposition 45 it is sufficient to check that for any finite tree t with a leaf x we have

$$\mathbb{P}(T_n \in \mathcal{T}(\mathsf{t},x)) \xrightarrow[n\to\infty]{} \mathbb{P}(\tau \in \mathcal{T}(\mathsf{t},x)) = \prod_{u \in \mathsf{t}\setminus\{x\}} p_{c_u},$$

where the last equality has been shown in the proof of Proposition 44. On the event $T \in \mathcal{T}(\mathsf{t},x)$ we denote by s the tree grafted on x. Using the assumption (5.3) we made on the "size" function A we can write for all n large enough

$$\begin{aligned}
\mathbb{P}(T_n \in \mathcal{T}(\mathsf{t},x)) &= \frac{1}{\mathbb{P}(T \in \mathcal{A}_n)} \mathbb{P}(T \in \mathcal{T}(\mathsf{t},x) \text{ and } A(T) \in [n, n+n_0)) \\
&\underset{\text{for large } n}{=} \frac{1}{\mathbb{P}(T \in \mathcal{A}_n)} \mathbb{P}(T \in \mathcal{T}(\mathsf{t},x) \text{ and } \mathsf{s} \in \mathcal{A}_{n-D(\mathsf{t},x)}) \\
&\underset{\text{branching}}{=} \frac{\mathbb{P}(T \in \mathcal{A}_{n-D(\mathsf{t},x)})}{\mathbb{P}(T \in \mathcal{A}_n)} \mathbb{P}(T \in \mathcal{T}(\mathsf{t},x)).
\end{aligned}$$

Since $D(\mathsf{t},x)$ is a fixed number, by our assumption on $\mathbb{P}(T \in \mathcal{A}_n)$ the fraction in the last display tends to 1 as $n \to \infty$. The second term is easily seen to be equal to $\prod_{u\in\mathsf{t}\setminus\{x\}} p_{c_u}$ as desired. □

5.3.5 Applications

Conditioning at large heights We consider the size function $A(\mathsf{t})$ to be the height, i.e. the maximal generation attained by the tree t. Clearly this function satisfies (5.3). So in order to apply the last result, one needs to verify that $\mathbb{P}(T \in \mathcal{A}_{n+1})/\mathbb{P}(T \in \mathcal{A}_n) \to 1$. We first treat the case when $n_0 = \infty$. In this case $T \in \mathcal{A}_n$ if and only if T is not extinct at generation n. However, the extinction probability for a Galton–Watson tree is known to obey

$$\mathbb{P}(\partial[T]_n \neq \varnothing) = 1 - \mathrm{F}_{\mathrm{p}}^{(n)}(0),$$

where $\mathrm{F_p}(z) = \sum_{k \geq 0} z^k p_k$ is the generating function of the offspring distribution and $\mathrm{F_p}^{(n)}$ is its n-fold composition. Recall that when p is critical (and $p_1 \neq 1$) we have $\mathrm{F_p}^{(n)}(0) \to 1$ (the Galton–Watson tree almost surely dies out). Since $\mathrm{F_p}^{(n)}(0)$ is a sequence defined by iterations of $\mathrm{F_p}$ we have

$$\frac{\mathbb{P}(\mathcal{A}_{n+1})}{\mathbb{P}(\mathcal{A}_n)} = \frac{\left(1 - \mathrm{F_p}^{(n+1)}(0)\right)}{\left(1 - \mathrm{F_p}^{(n)}(0)\right)} \to \mathrm{F_p}'(1) = 1.$$

Now let us treat the case $n_0 = 1$ meaning that $T \in \mathcal{A}_n$ is the height of T is exactly n. Using the last arguments in this case we have

$$\begin{aligned}\frac{\mathbb{P}(T \in \mathcal{A}_{n+1})}{\mathbb{P}(T \in \mathcal{A}_n)} &= \frac{\mathbb{P}(\mathrm{Height}(T) \geq n+1) - \mathbb{P}(\mathrm{Height}(T) \geq n+2)}{\mathbb{P}(\mathrm{Height}(T) \geq n) - \mathbb{P}(\mathrm{Height}(T) \geq n+1)} \\ &= \frac{\left(1 - \mathrm{F_p}^{(n+1)}(0)\right) - \left(1 - \mathrm{F_p}^{(n+2)}(0)\right)}{\left(1 - \mathrm{F_p}^{(n)}(0)\right) - \left(1 - \mathrm{F_p}^{(n+1)}(0)\right)} \to 1.\end{aligned}$$

We can thus apply Theorem 46 in both cases $n_0 = 1$ or $n_0 = \infty$ and get that Galton–Watson tree conditioned on having an extinction after a large height or at an exact large height converge towards the infinite Galton–Watson tree conditioned to survive.

Conditioning at large size If the function $A(t)$ is the number of vertices it is a little more difficult to prove the required condition on $\mathbb{P}(T \in \mathcal{A}_n)$ demanded by Theorem 46. However when the offspring distribution has a finite variance σ^2 we can use the connection of Galton–Watson trees with random walk and the local central limit theorem to deduce that the conditions of Theorem **??** are always satisfied.

Bibliographical notes. Augmented Galton–Watson trees have been introduced in the beautiful paper [36] of Lyons, Pemantle and Peres where they proved Theorem 40. The proof presented here is taken from [37, Chapter 17]. Proposition 38 is taken from [6] where the supercritical geometric GW tree conditioned to survive arise as the local limit of unicellular maps in high genus. Theorem 41 and its proof seems to be new in the literature. The general introduction of Kesten's infinite random tree is usually dated back to [32] where Kesten used it as a model of random graph where the simple random walk has an anomalous diffusive behavior. In the particular case of Poisson Galton–Watson tree this fact has been known earlier [25]. The construction of Kesten's infinite tree as a change of measure is mainly due to works of Lyons, Pemantle and Peres, see [37]. The approach to the convergence given here (one of the simplest as of today) is taken from [2] to which we refer for further details and references.

Chapter 6

Unimodular random triangulations

In this chapter we study unimodular random planar triangulation (in particular we will prove the Benjamini–Schramm's theorem stating that any local limit of uniformly pointed planar graphs with bounded degree is recurrent). This chapter will also serve us as a good pretext to introduce the basics on random maps and random triangulations as well as describing the marvelous theory of circle packings.

6.1 Planar maps

6.1.1 Generalities

A planar graph is a locally finite (multiple) graph which can be drawn on the plane (or equivalently on the sphere) in such a way that the edges are non-crossing except at the vertices. Such a drawing is called a proper embedding. Notice that a planar graph may have several topologically different proper embeddings and the definition only tells us the existence of such. In particular, the notion of face of the graph is subject to vary with the embedding.

Definition 16. *A finite planar map is a finite connected planar graph properly embedded in the plane (or on the sphere) viewed up to homeomorphisms that preserve the orientation.*

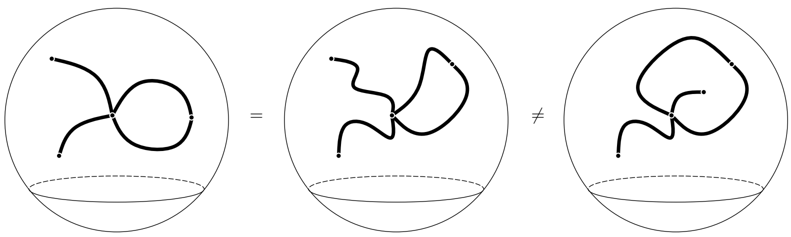

Figure 6.1: The same underlying planar graph can yield different planar maps.

In other words, a finite planar map is an equivalence class of embeddings of a finite planar graph. This will allow us to properly define the number of faces of the map and the incidence relations between vertices, edges and faces. The degree of a face is the number of edges incident to this face with the convention that when an edge is lying completely inside a face it is counted twice in the degree. There is an analogous definition of a finite map drawn on the torus or more generally on a compact (orientable) surface of genus $g \geq 0$, but since we will restrict ourselves to the planar case we sometimes drop the adjective planar and speak of a map instead of a planar map. May the reader forgive this. Definition 16 may seem hard to manipulate at first glance, luckily it admits several equivalent points of view:

- a finite planar map can be seen as a topological gluing of finitely many polygons (the faces of the map) along their edges so that the manifold produced this way is a topological sphere,
- a finite planar map can also be seen as a finite graph with a system of coherent orientations around each vertex of the graph which correspond to the cyclic ordering of the edges when going clockwise around a vertex in the map.

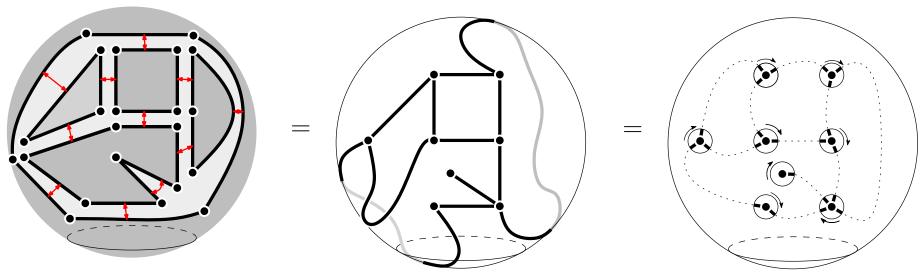

Figure 6.2: A planar map seen as a gluing of polygons, notice that two edges of a same polygon could be folded to give a single edge in the map.

Using the last definition it should be clear that the number of planar maps with a given number of edges is finite. Also, this definition is practical since we can make sense of *infinite* maps just as being infinite graphs given with a system of coherent (i.e. giving rise to a planar structure) of cyclic orientations of the edges around each vertex. In the following we will denote by $\mathcal{M}^\bullet$ (resp. $\vec{\mathcal{M}}$) the set of all pointed (resp. rooted) finite of infinite maps. If $\mathsf{m} \in \mathcal{M}^\bullet$ or $\mathsf{m} \in \vec{\mathcal{M}}$ we can make sense of the restriction of radius r, denoted by $[\mathsf{m}]_r$ as the pointed or rooted planar map obtained by keeping only vertices and edges within distance r from the origin. This yields the introduction of the the local topology on $\mathcal{M}^\bullet$ and $\vec{\mathcal{M}}$. A *simple* map is a map in which multiple edges or loops are forbidden.

Remark 5. Although the above definition of infinite planar maps seems inoffensive, it is not clear at all that these objects can be represented as equivalence classes of planar graphs drawn on the plane. This is true however in the case of one-ended infinite maps (that is when the underlying graph is one-ended, or equivalently simply-connected). We call later these maps “infinite maps of the plane” since they can be properly drawn on $\mathbb{R}^2$ such that there is no accumulation of edges or vertices inside $\mathbb{R}^2$.

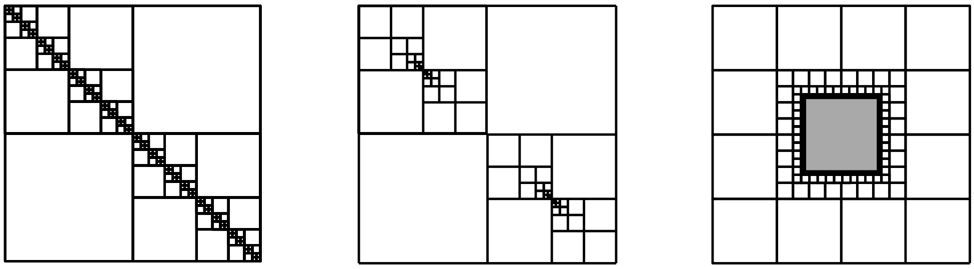

Figure 6.3: Three examples of infinite maps, the left-most one has infinitely many ends, the center one has two ends whereas the right-most one has only one end (the centered region in gray is not a face). In particular, the right-most map can be drawn on the plane $\mathbb{R}^2$ without accumulation points for the edges.

If m is a finite of infinite planar map such that all the faces of m are of finite degree, one can define the dual map $\mathsf{m}^\dagger$ obtained informally speaking by placing inside each face of m a vertex of $\mathsf{m}^\dagger$ and linking two vertices of $\mathsf{m}^\dagger$ by an edge if the corresponding faces in m share an edge. The duality mapping is clearly an involution on the set of all planar maps with finite face degrees. If the map m is rooted then $\mathsf{m}^\dagger$ inherits its dual edge as root edge.

A famous theorem on planar map which looks childish is the 4-colors theorem which proves that 4 colors suffice to color any planar map such that any pair of adjacent faces (i.e. sharing an edge) have different colors. The proof is extremely difficult and requires the help of a computer to check numerous cases, but a version with 5 colors is much easier to do.

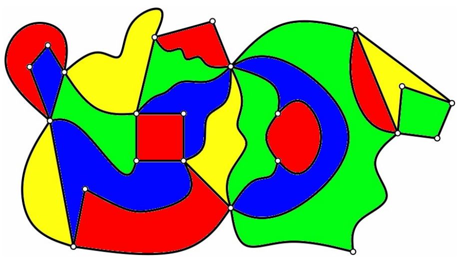

We can easily extend the notion of uniformly rooted or pointed maps, unimodular random map, stationary and reversible random maps.

Exercise 41. Prove that if $(M, \vec{E})$ is a uniformly rooted random map, then $(M^\dagger, \vec{E}^\dagger)$ is also uniformly rooted. $(*)$ Prove that if $(M, \vec{E})$ is a stationary and reversible random map, then $(M^\dagger, \vec{E}^\dagger)$ is also stationary and reversible (use unimodularity). $(**)$ Is it true without assuming reversibility?

6.1.2 Euler's formula

The first non-trivial result about planar maps is the famous Euler relation which links the number of faces, of edges and of vertices of any finite planar map. We denote the set of vertices, edges and faces of a map m by $\mathrm{V}(\mathsf{m}), \mathrm{E}(\mathsf{m})$ and $\mathrm{F}(\mathsf{m})$ respectively.

Theorem 47

For any finite planar map m *we have*

$$\#\mathrm{V}(\mathsf{m}) + \#\mathrm{F}(\mathsf{m}) - \#\mathrm{E}(\mathsf{m}) \quad = \quad 2. \tag{6.1}$$

Proof. The proof is done by induction on the number of edges. The only map with 0 edge has 1 vertex and 1 face so that (6.1) is true. Suppose now that $\mathrm{E}(\mathsf{m}) \geq 1$ and erase an arbitrary edge of m, then two cases may happen:

- either the new map m' is still connected and so applying the induction hypothesis we have $\#\mathrm{V}(\mathsf{m}')+\#\mathrm{F}(\mathsf{m}')-\#\mathrm{E}(\mathsf{m}') = 2$. Also we have $\#\mathrm{V}(\mathsf{m}) = \#\mathrm{V}(\mathsf{m}')$ and $\#\mathrm{E}(\mathsf{m}) = \#\mathrm{E}(\mathsf{m}')+1$ and a careful inspection shows that $\#\mathrm{F}(\mathsf{m}) = \#\mathrm{F}(\mathsf{m}') + 1$. Gather-up the pieces we find that m obeys (6.1).

- or the removal of the edge breaks m into two connected maps m_1 and m_2. Applying (6.1) to each block we find that $\#\mathrm{V}(\mathsf{m}_1)+\#\mathrm{F}(\mathsf{m}_1)-\#\mathrm{E}(\mathsf{m}_1) = 2$ as well as $\#\mathrm{V}(\mathsf{m}_2)+\#\mathrm{F}(\mathsf{m}_2)-\#\mathrm{E}(\mathsf{m}_2) = 2$. Also, we have $\#\mathrm{V}(\mathsf{m}) = \#\mathrm{V}(\mathsf{m}_1)+\#\mathrm{V}(\mathsf{m}_2)$ and $\#\mathrm{E}(\mathsf{m}) = \#\mathrm{E}(\mathsf{m}_1)+\#\mathrm{E}(\mathsf{m}_2)+1$ and another careful inspection shows that $\#\mathrm{F}(\mathsf{m}) = \#\mathrm{F}(\mathsf{m}_1) + \#\mathrm{F}(\mathsf{m}_2) - 1$, the minus 1 terms stems from the fact that the external face of m_1 and m_2 is counted twice otherwise. Putting everything together we indeed verify (6.1).

$\square$

Remark 6. As we already noticed, the notion of face is not well-defined for planar graphs as it may depend on its planar embedding. However, we see from Euler's formula that the number of faces does not depend on the embedding but only on the underlying graph structure.

Exercise 42. Using Euler's formula show that the complete graph K_5 on 5 vertices (with an edge between any pair of vertices) and the graph $K_{3,3}$ made of 3 black vertices and 3 white vertices such that there is an edge between any pair of black and white vertices are not planar graphs.

A well-known application of Euler's formula is the classification of all regular polyhedrons or Platonic solids. Indeed, a regular polyhedron can be seen as a finite map such that the degree of the vertices and faces are constant.

Exercise 43 (regular polyhedrons). Show that there are only 5 regular polyhedrons such that the degree of the vertices and the degree of the faces satisfy

Name	vertex degree	face degree
Tetrahedron	3	3
Cube	3	4
Octahedron	4	3
Dodecahedron	3	5
Icosahedron	5	3

6.1.3 Mean degree in uniformly pointed planar maps

In the following we will use a lot *triangulations*. A triangulation is a planar map whose faces have all degree three (they are triangles). Beware, since we allow multiple edges and loops, a triangle can be folded on itself and look weird at first glance, see Fig. 6.4. When multiple edges or loops are forbidden we speak of *simple* triangulations.

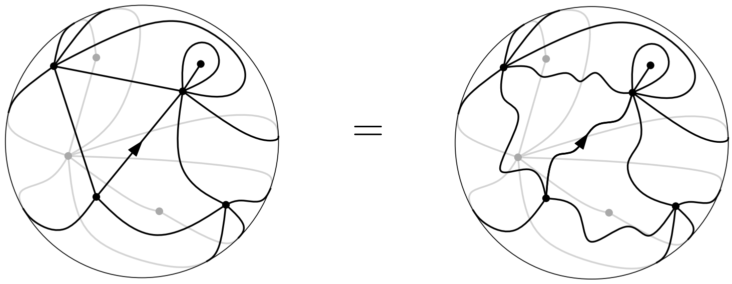

Figure 6.4: A finite triangulation of the sphere. Notice the triangle which is folded on itself and looks like a loop with an inner edge : this is indeed a triangle, since the enclosed face has degree 3!

If m is a finite planar map with no face of degree 1 or 2 then we must have $3\#\mathrm{F}(\mathsf{m}) \leq 2\#\mathrm{E}(\mathsf{m})$ since any edge is counted twice in the total sum of the degrees of the faces. So after applying Euler's formula we get that

$$\begin{aligned} \#\mathrm{V}(\mathsf{m}) + \#\mathrm{F}(\mathsf{m}) - \#\mathrm{E}(\mathsf{m}) &= 2 \\ 3\#\mathrm{V}(\mathsf{m}) + 2\#\mathrm{E}(\mathsf{m}) - 3\#\mathrm{E}(\mathsf{m}) &\geq 6 \\ 6\left(1 - \frac{2}{\#\mathrm{V}(\mathsf{m})}\right) &\geq \frac{2\#\mathrm{E}(\mathsf{m})}{\#\mathrm{V}(\mathsf{m})} = \text{mean degree}. \end{aligned} \tag{6.2}$$

We will remember this formula as the fact that the mean degree in a finite planar graph without loops or faces of degree 2 is always less than or equal to 6. In particular it is the case for simple maps (except the map made of a single edge). If m is a triangulation then we even have equality in (6.2). Let us deduce a corollary on limits of uniformly pointed simple maps.

Proposition 48. *Let $M_n^\bullet = (M_n, \rho_n)$ be a sequence of random uniformly pointed finite planar maps without face of degree 1 or 2 converging locally towards some infinite random pointed map $M^\bullet = (M, \rho)$ as $n \to \infty$ then we have*

$$\mathbb{E}[\deg(\rho)] \leq 6.$$

If furthermore $M_n^\bullet$ are triangulations then $\mathbb{E}[\deg(\rho)] = 6$.

Proof. Since $M_n^\bullet = (M_n, \rho_n)$ is uniformly pointed we deduce that

$$\mathbb{E}[\deg(\rho_n)] = \mathbb{E}\left[\frac{1}{\#\mathrm{V}(M_n)} \sum_{x \in \mathrm{V}(M_n)} \deg(x)\right] = \mathbb{E}\left[\frac{2\#\mathrm{E}(M_n)}{\#\mathrm{V}(M_n)}\right] \underset{(6.2)}{\leq} 6\,\mathbb{E}\left[1 - \frac{2}{\#\mathrm{V}(M_n)}\right]. \tag{6.3}$$

By the local convergence of the maps and since the degree of the origin is a continuous function for the local topology, we get that $\deg(\rho_n) \to \deg(\rho)$ in distribution as $n \to \infty$. Thus by the above display and Fatou's lemma we get that

$$\mathbb{E}[\deg(\rho)] \leq 6.$$

In the case of triangulations, there is an equality in (6.3) and so to deduce $\mathbb{E}[\deg(\rho)] = 6$ we must show that $(\deg(\rho_n))_{n\geq 1}$ is a uniformly integrable family of random variables. Or equivalently that for any $\varepsilon > 0$ we can find k_0 such that uniformly in n we have

$$\mathbb{E}[\deg(\rho_n)\mathbb{1}_{\deg(\rho_n)\geq k_0}] = \mathbb{E}\left[\frac{1}{\#\mathrm{V}(M_n)} \sum_{\substack{x\in\mathrm{V}(M_n)\\ \deg(x)\geq k_0}} \deg(x)\right] \leq \varepsilon.$$

This will follow from the local convergence of $M_n^\bullet$. For simplicity we suppose now that the maps M_n are simple. Since $(M_n^\bullet)$ is locally converging in distribution, it is a tight sequence, and so for any $\varepsilon > 0$, there exists k_0 such that the probability that there exists a vertex within distance 1 from the ρ_n having a degree larger than k_0 is smaller than ε uniformly in n,

$$\mathbb{P}(\exists y : \mathrm{d_{gr}}(\rho_n, y) \leq 1 \text{ and } \deg(y) \geq k_0) = \mathbb{E}\left[\frac{1}{\#\mathrm{V}(M_n)} \sum_{x\in\mathrm{V}(M_n)} \mathbb{1}\{\exists y : \mathrm{d_{gr}}(x, y) \leq 1 \text{ and } \deg(y) \geq k_0\}\right] \leq \varepsilon.$$

Now, for any set $\mathcal{S}$ of vertices of degree larger than or equal to k_0 in a simple map m, we consider the set of vertices $\mathcal{S}^{(1)}$ formed by $\mathcal{S}$ as well as its neighbor vertices. Since we assumed that m is simple, the maps formed by the connected components of $\mathcal{S}^{(1)}$ are also simple and have no face of degree 1 or 2. For each connected component C of $\mathcal{S}^{(1)}$ we thus have by (6.2) $\sum_{x\in\mathrm{V}(\mathsf{C})} \deg(x) \leq 6\,\#\mathrm{V}(\mathsf{C})$. Summing-up over the connected components we deduce that

$$\sum_{x\in\mathcal{S}} \deg(x) \leq 6\,\#\mathcal{S}^{(1)}.$$

Combining the last three displays we get that

$$\mathbb{E}[\deg(\rho_n)\mathbb{1}_{\deg(\rho_n)\geq k_0}] \leq 6\,\mathbb{P}(\exists y : \mathrm{d_{gr}}(\rho_n, y) \leq 1 \text{ and } \deg(y) \geq k_0) \leq 6\,\varepsilon.$$

This proves that $(\deg(\rho_n))_{n\geq 1}$ is indeed uniformly integrable (in the case of simple maps) and so converges in expectation as desired. We leave the details of the proof in the triangulation case (the hiccup is that the connected components of $\mathcal{S}^{(1)}$ are not necessarily simple maps!) to the reader. □

Remark 7. We deduce from the above result that any unimodular triangulation of mean degree strictly larger than 6 (see for example the 7-regular triangulation of Fig. 6.10) cannot be obtained as a local limit of uniformly pointed finite planar maps. Notice that by the Aldous–Lyons conjecture 5 they can be obtained (at least their pointed graph structures) as local limit of finite unimodular random graphs but not necessarily planar.

Exercise 44. Let $\vec{T}_n$ be a sequence of uniformly rooted triangulations of the sphere converging locally towards an infinite rooted triangulation $\vec{T}$. Show that $\mathbb{E}[(\deg((\vec{E})_*))^{-1}] = 1/6$.

6.1.4 Curvature and isoperimetric inequality

In this section we consider triangulations only. Recall that if $\mathtt{t}$ is a finite triangulation with $v > 0$ vertices then its mean degree $2e/v = 6 - 12v^{-1}$. In conformal geometry the mean degree represents the "average curvature": if the mean degree is equal to 6 the surface is flat, if it is larger than 7 the surface is negatively curved and if it is smaller than 6, it is positively curved. Everybody knows the standard 6-regular triangulation, which is flat, known as the honey-comb lattice. However, it is easy to see that there exists infinite triangulations whose vertex degrees are bounded from below by 7 say (e.g. the 7-regular triangulation) but they grow very rapidly. This can be encoded in the so-called isoperimetric profile:

Theorem 49 (***Degrees and isoperimetric profile***)

Let $\mathtt{t}$ be a triangulation with a boundary of length p, that is a planar map whose faces are all triangles except for one face, called the external face which is of degree p. We denote by n the number of inner vertices of $\mathtt{t}$

- *If all the inner vertex degrees are larger than or equal to 7 then for some $c > 0$*

$$p \geq c \cdot n.$$

- *If all the inner vertex degrees are larger than or equal to 6 then*

$$p \geq \sqrt{12n}.$$

Proof. Let $\mathtt{t}$ be a triangulation with n inner vertices and whose minimal inner vertex degree is $d \in \{6, 7\}$. We may choose $\mathtt{t}$ so that p is the smallest possible. Notice that this forces the boundary $\partial\mathtt{t}$ to be a simple cycle since otherwise if there are pinch points by a simple surgical operation we can glue two edges and diminish the perimeter while keeping the number of inner vertices and their degrees unchanged. We write f for the number of faces of $\mathtt{t}$ and e its number of edges. Counting the edges from the face point of view gives $2e = 3(f-1) + p$ and Euler's formula writes $(n+p) + f - e = 2$ which once combined give

$$3n + 2p = 3 + e. \tag{6.4}$$

On the other hand counting edges from the vertex point of view yields

$$2e = \sum_{u \in \mathsf{Vertices}(\mathtt{t})} \deg(u). \tag{6.5}$$

From the above display we deduce that $2e \geq dn$ where d is the minimal inner vertex degree and combining this with (6.4) already yields the first point of the theorem.

In the case $d = 6$ we must do better and we adapt here the proof of Angel, Benjamini and Horesh communicated to us by Itai Benjamini. We introduce Σ the edges incident to both a vertex on the boundary $\partial\mathrm{t}$ of t and an inner vertex of t and Δ the edges linking two vertices of $\partial\mathrm{t}$. Coming back to (6.5) more carefully we get $2e = 6n + 2p + \#\Sigma + 2\#\Delta$ which together with (6.4) yields

$$\#\Sigma + 2\#\Delta = 2p - 6. \tag{6.6}$$

We already deduce that $p \geq 4$ unless t is made of a single triangle. We now assume that the triangulation has been chosen so that the ratio $c = p^2/n$ is the smallest possible among all triangulations with boundary so that $n \leq N$ where N is fixed (if there are several choices, pick one with minimal p). Let us examine a bit more the structure of such a minimal triangulation. Recall that the boundary is necessarily simple, and let us now rule-out the possibility of anon-boundary edge $\partial\mathrm{t}$ linking two vertices of $\partial\mathrm{t}$. Indeed, if there was such an edge it would split the map into two triangulations with boundary of perimeter $p_1 + 1$ and $p_2 + 1$ with n_1 and n_2 inner vertices respectively such that $n_1 + n_2 = n$ and $p_1 + p_2 = p$.

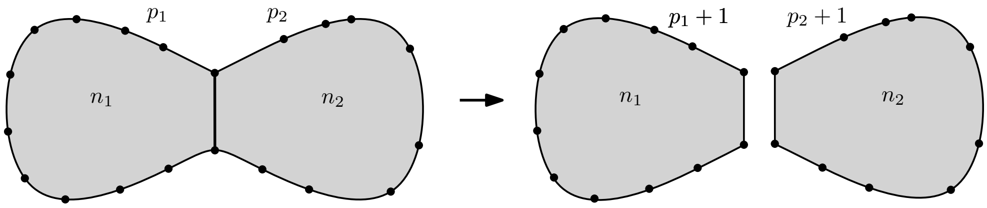

Figure 6.5: One cannot split t into two parts by the minimality assumption.

Notice then that necessarily $p > p_1 \geq 2$ and $p > p_2 \geq 2$ and $(p_1 - 1)(p_2 - 1) > 1$ otherwise t is made of two triangles glued together and n would be equal to 0. By our minimality assumption we must have $(p_1 + 1)^2 \geq cn_1 + 1$ and $(p_2 + 1)^2 \geq cn_2 + 1$. Using the fact that $(p_1 - 1)(p_2 - 1) > 1$ it follows that

$$(p_1 + p_2)^2 > (p_1 + 1)^2 + (p_2 + 1)^2 - 2 \geq c(n_1 + n_2),$$

which is absurd. Hence we have with the above notation $\Delta = \varnothing$. Now if we consider the set of all inner vertices adjacent to the boundary of t, by our deduction on t they form a connected subset which encloses a triangulation t' with a boundary of perimeter p'. It is easy to see on the figure that $\#\Sigma = p + p'$ and so using (6.6) we have that $p \geq p' + 6$. Using our minimality assumption again we deduce that $(p')^2 \geq cv' + 1$ where v' is the number of inner vertices of t'. Since we obviously have $v \leq v' + p'$ (with equality if t' has a simple boundary) we deduce that

$$c(v' + p') \geq cv = p^2 \geq (p' + 6)^2 \geq (p')^2 + 12p' \geq cv' + 12p',$$

and this can only work if $c \geq 12$. As N was arbitrary, this proves the second statement of the theorem. □

Remark 8. The isoperimetric constant $\sqrt{12}$ is achieved in the case of balls of large radius in the standard infinite 6-regular triangulation.

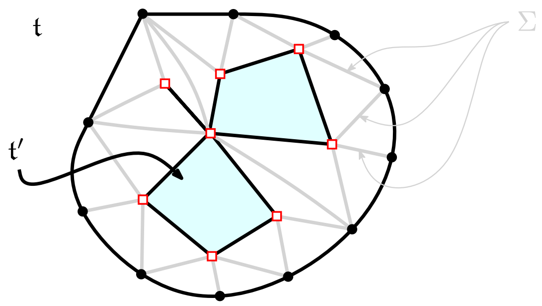

Figure 6.6: Induction hypothesis: passing from $\mathfrak{t}$ to $\mathfrak{t}'$.

Exercise 45. Show that there is not infinite triangulation of the plane whose vertex degrees are bounded by 5. Is there an infinite triangulation of the plane with degrees only in $\{5, 6\}$ and having infinitely many vertices of degree 5?

6.2 Circle packings

A planar map does not a priori have any canonical representation in the plane (or the sphere) since even in the finite case, it is given as an equivalence class of embeddings. Still one can ask if we can make sense of a "faithful" representation of a map. In particular, one can wonder whether a finite map can always be drawn with straight lines on the plane. Surprisingly, once the trivial cases have been excluded the answer is yes!

6.2.1 Finite Circle Packings

For now on we focus on the case of simple maps where multiple edges and loops have been forbidden (since in any representation with straight lines, the latter are squashed). We say that a simple map $\mathfrak{m}$ is represented by a circle packing if there is a collection $(\mathcal{C}_v : v \in \mathrm{V}(\mathfrak{m}))$ of non overlapping disks in the plane $\mathbb{R}^2$ such that $\mathcal{C}_v$ is tangent to $\mathcal{C}_u$ if and only if u and v are neighbors in $\mathfrak{m}$. Recall that the completed plane $\hat{\mathbb{C}} = \mathbb{R}^2 \cup \{\infty\}$ can be identified with the Riemann sphere $\mathbb{S}^2$ by the stereographic projection from the north pole. This projection transforms circles and lines in $\hat{\mathbb{C}}$ into circles on the Riemann sphere. Recall also that the Möbius group

$$\mathcal{M} = \left\{ z \in \hat{\mathbb{C}} \mapsto \frac{az+b}{cz+d} \right\}$$

acts triply transitively on the Riemann sphere (i.e. we can map any triplet of points to any other triplet of points) and preserves circles.

Theorem 50 (***Finite circle packing theorem***)

Any finite simple map m *admits a circle packing representation on the Riemann sphere. Furthermore if* m *is a simple triangulation then the circle packing is unique up to a Möbius transformation.*

Remark 9. It follows from the circle packing theorem that any simple planar graph can be drawn on the plane with straight lines. This fact is not obvious at all and is known as Fáry's theorem (1948).

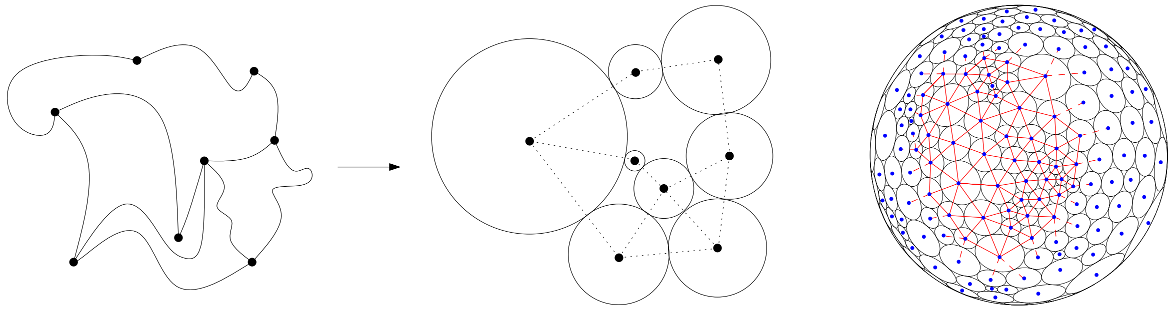

Figure 6.7: On the left, a finite circle packing of a planar map. On the right a circle packing of a triangulation seen on the Riemann sphere.

Sketch of proof. First, it is easy to see that it suffices to prove the theorem for simple triangulations because we can embed any simple planar map inside a simple triangulation by further triangulating inside each face. Fix a triangulation t and pick a face $f \in \mathrm{F}(\mathsf{t})$ that we will see as the exterior face. We will prove that we can construct a circle packing of t such that the three circles corresponding to this outer face are three mutually tangent circles of radius 1, or equivalently that the three vertices of the triangles form an equilateral triangle. The rest of the circles are in-between these three circles. We start with the uniqueness statement.

Uniqueness. Since the Möbius group of the Riemann sphere acts triply transitively we can transform any circle packing into a packing of the above form (with the marked face forming an equilateral triangle and the rest of the vertices inside). Imagine that we are given two packings $\mathcal{P}$ and $\mathcal{P}'$ of the above form, in particular the three exterior circles are of radius 1. We then choose an interior vertex v of the triangulation such that the ratio of the corresponding circles in the packing is maximal i.e.

$$\lambda(v) = \frac{r_{\mathcal{P}}(v)}{r_{\mathcal{P}'}(v)} \quad \text{is maximal.}$$

We then examine the structure of the packing around this circle in $\mathcal{P}$. By dividing all the distances by $\lambda(v)$ we end up with a circle of radius $\mathrm{r}_{\mathcal{P}'}(v)$ and such that all the neighboring circles have a radius which is less than the corresponding radius in $\mathcal{P}'$. By an obvious monotonicity property of the angles around a circle we deduce that these new radii must coincide with those in $\mathcal{P}'$ i.e.

$$\lambda(u) = \frac{r_{\mathcal{P}}(u)}{r_{\mathcal{P}'}(u)} = \lambda(v),$$

for all u neighbors of v (for otherwise if the inequality were strict, the neighboring circles would not surround the circle associated to v). Since the graph is connected we deduce step by step that $\lambda(\cdot)$ is constant and must be equal to 1 by the assumption on the exterior circles. Hence $\mathcal{P} = \mathcal{P}'$.

Existence. We will not prove the existence but just describe the algorithm that can be used (even in practice!) to construct the packing. The idea is to first find all the radii of the circles. Once these radii are found, one can reconstruct the packing step by step by starting from the external face and deploying the circles one by one around the circles already explored (notice that given the radii and the combinatorial layout we can determine the angles, and for this we crucially use the fact that the underlying map is a triangulation). To find the radii we start with an arbitrary assignment of radii to the vertices of the triangulations except the three vertices of the marked face which have their radii fixed for ever to 1. We then examine all the internal vertices in a cyclic order and repeat forever the following adjustment : see Fig. 6.8.

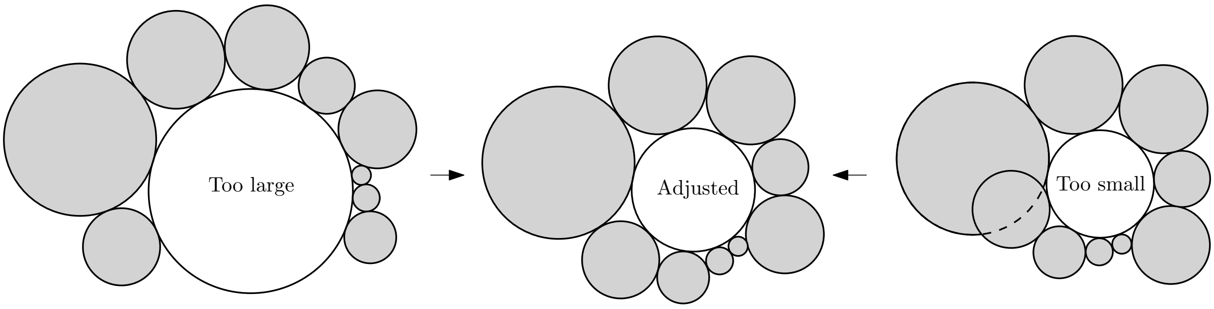

Figure 6.8: Adjustment rule : For an internal vertex v with radius r_v, we examine the radii of the neighbors u of v. Using these radii, one can see whether or not placing the circles of the corresponding radii around a circle of radius r_v would close exactly. For most of the time it will not. But by a simple monotonicity property, one can always update the radius r_v so that the latter property holds true.

Repeatedly applying this updating rule, it can be proved (but it is not easy) that this algorithm indeed converges towards the unique fixed point for the right values of the radii for the circles (with the outer three triangles normalized) and that these values give rise to a non-degenerate (all the radii are positive) circle packing for the triangulation $\mathfrak{t}$. □

Although we will not need it, we cannot resist stating the link between circle packings and the Riemann mapping theorem. Imagine that we have a circle packing of a region Ω, with the hexagonal packing say, and that the same combinatorial triangulation structure is circle-packed in the disk (we can always do so by the finite circle packing theorem). We suppose also that the center $z \in \Omega$ is mapped to $0 \in \mathbb{D}$. Then as the maximal radius of the circle goes to 0, the mapping induced by the circle packings approximates a conformal function bijection $\Omega \to \mathbb{D}$ such that z is mapped to 0.

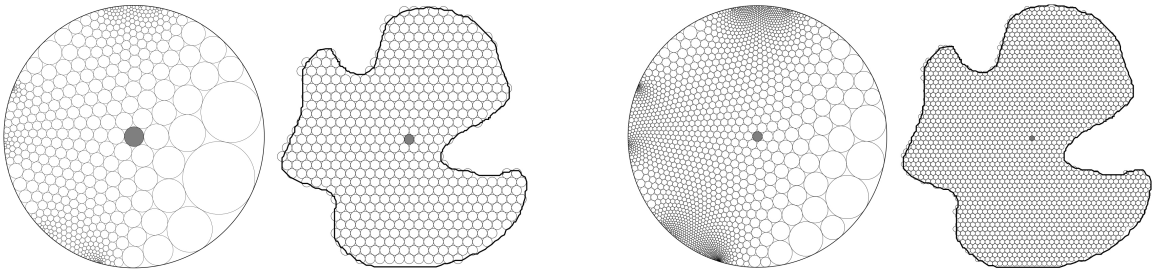

Figure 6.9: Thurston conjecture (Rodin–Sullivan/Schramm theorem) : Circle packings can be used to approximate conformal mappings. Images of Kenneth Stephenson.

6.2.2 Infinite Circle Packings

Recall the definition of infinite planar maps as oriented planar graphs and in particular of one-ended infinite triangulation also called infinite triangulation of the plane. If $\mathcal{P}$ is an infinite circle packing in the plane the *carrier* of $\mathcal{P}$ is the subset of the plane made of the union of all the circles as well as the interstices between them, see Fig. 6.10 for a carrier equal to the unit disk $\mathbb{D}$ (left) or the plane (right).

Theorem 51 (*Infinite circle packing theorem*)

Let $\mathtt{t}$ *be a infinite simple triangulation of the plane (i.e. 1-ended). Then we have one of the mutually excluding alternatives:*

- *Parabolic case: either there is a circle packing whose carrier is* $\mathbb{R}^2$ *representing* $\mathtt{t}$,
- *Hyperbolic case: or there is a circle packing whose carrier is* $\mathbb{D}$ *representing* $\mathtt{t}$.

In the first case, the packing is unique up to rotation, translation and dilation, whereas in the second case it is unique up to Möbius transformation preserving the unit disk. Furthermore, if the vertex degrees of $\mathtt{t}$ *are bounded, the above dichotomy corresponds to the case when* $\mathtt{t}$ *is recurrent (packing in* $\mathbb{R}^2$*) or transient (packing in* $\mathbb{D}$*).*

The above theorem can be seen as a discrete counter-part to the dichotomy for simple connected Riemann surfaces homeomorphic to the disk: either such a surface is conformally equivalent to the disk (and Brownian motion on the surface is transient) or it is conformally equivalent to the plane (and Brownian motion on the surface is recurrent). We will not give its proof which is based on the notion of discrete (vertex) extremal length, [28].

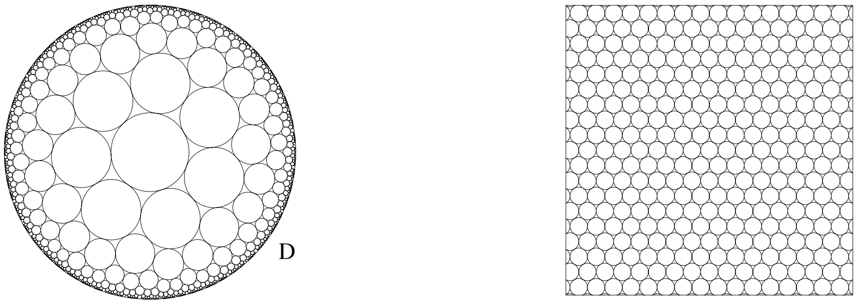

Figure 6.10: Illustration of Theorem 51: the 7-regular infinite triangulation is circle-packed in the disk (and is transient) whereas the 6-regular infinite triangulation is circle-packed in the plane (and is recurrent). Images of Kenneth Stephenson.

6.3 Applications

In this section we mix the mass-transport principle and the circle packing to get powerful results on unimodular planar maps and in particular unimodular planar triangulations. In particular we will prove the Benjamini–Schramm theorem [14] in the case of triangulations.

6.3.1 Mean degree in unimodular random triangulation

Theorem 52

Let (T, ρ) be a unimodular infinite random simple triangulation which is almost surely one-ended. Then we have $\mathbb{E}[\deg(\rho)] \geq 6$. Furthermore if $\mathbb{E}[\deg(\rho)] = 6$ then T and T is almost surely parabolic.

Proof. By Theorem 51, the triangulation T admits a circle packing either in the plane or in the disk (this could be random). In both cases, we consider a mass-transport function $f(\mathrm{t}, x, y)$ which transports from x to y inside the triangulation t the two angles formed at x by the two sides of the edge x–y. More precisely, if the circle packing is in the plane we transport the Euclidean angle and if the circle packing is in the hyperbolic disk then we transport the hyperbolic angle (this angle is the local Euclidean angle formed by the hyperbolic lines joining the hyperbolic centers of the circles[1]).

We can (and should!) wonder whether this mass-transport function is well-defined, as this angle could a priori depend on the packing. However, from Theorem 51 we know that the circle packing representing t is unique up to Möbius transformations of the disk or rotation, translation and dilation of the plane depending whether the carrier is $\mathbb{D}$ or $\mathbb{R}^2$. Since all these transformations preserve local angles, the angle of the edge x–y at x is indeed well-defined and it follows that $f(\mathrm{t}, x, y)$ is an honest transport function (the measurability follows from the

[1] we recall that a Euclidean circle in $\mathbb{D}$ can also be seen as a hyperbolic circle with a center which is generally different from the Euclidean center

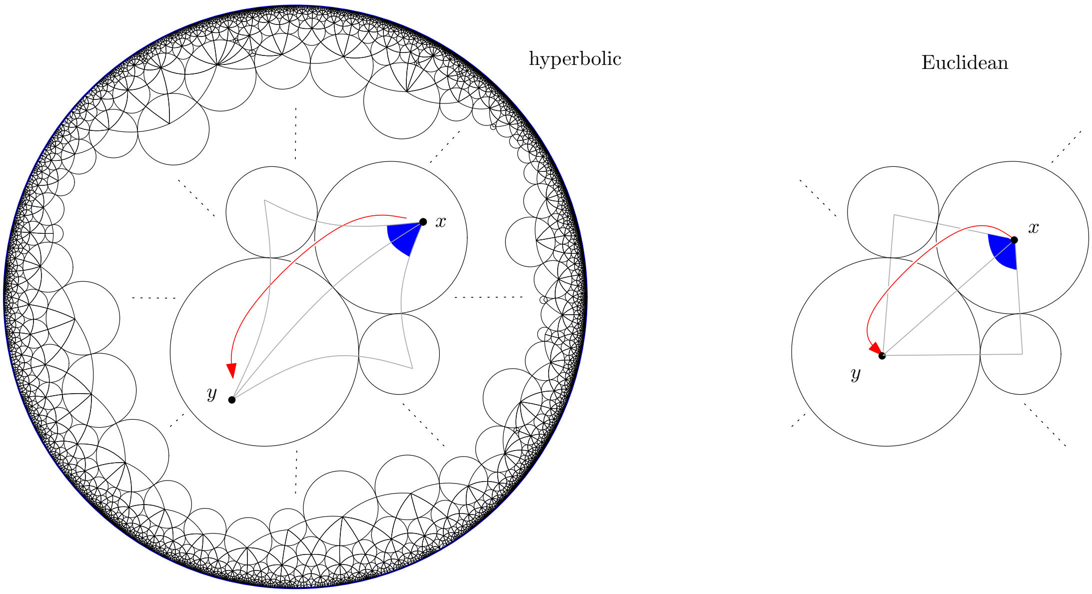

Figure 6.11: Illustration of the mass-transport function in the case when the packing is in the disk or when the packing is in the plane.

measurability of the packing and is admitted). Applying the mass-transport principle (2.1) and recalling that the total angle sum around a vertex is 2π we get that

$$4\pi = \mathbb{E}[\text{total angle sum of triangles incident to } \rho] - 2\pi.$$

Now recall that in Euclidiean geometry, the total sum of the angles of a triangle is equal to π, whereas in hyperbolic geometry, the total sum of the angle of a triangle is strictly less than π (it is equal to π minus the hyperbolic area of the triangle). We thus deduce that $6 \le \mathbb{E}[\deg(\rho)]$ and with equality if and only if the circle packing is almost surely carried by the plane $\mathbb{R}^2$. This proves the result. □

By the above result if $\mathbb{E}[\deg(\rho)] = 6$ and if furthermore T almost surely has bounded degree then T is almost surely recurrent by Theorem 51. To cover the case of triangulations which are not necessarily one-ended we need to use its universal cover:

Proposition 53. *To any (infinite) pointed triangulation* (t, ρ) *we can associate a unique infinite pointed triangulation* $(\tilde{\mathsf{t}}, \tilde{\rho})$ *called the universal covering of* (t, ρ) *such that* $(\tilde{\mathsf{t}}, \tilde{\rho})$ *is one-ended (equivalently simply-connected) and there is a surjection mapping* $\mathfrak{p} : \tilde{\mathsf{t}} \to \mathsf{t}$ *which is locally bijective (and preserves the orientation of edges around each vertex) and sends* $\tilde{\rho}$ *to* ρ.

Proof. The construction of this space is analogous to the construction of the universal covering of a topological space and is left to the imagination of the reader. □

Corollary 54. *If* (T, ρ) *is a unimodular infinite random triangulation (not necessarily one-ended) then* $\mathbb{E}[\deg(\rho)] \ge 6$. *If* $\mathbb{E}[\deg(\rho)] = 6$ *and if* T *has almost surely bounded degree then* T *is almost surely recurrent.*

Proof. We apply Theorem 52 to the universal covering $(\tilde{T}, \tilde{\rho})$ of (T, ρ) which is easily seen to be still unimodular. Since the mean degree of $\tilde{\rho}$ is the same as the mean degree of ρ we deduce the first assertion. For the second assertion, note that since there is a surjection map from $\tilde{T}$ to T, it follows that if $\tilde{T}$ is recurrent then so is T. □

6.3.2 Benjamini–Schramm theorem

Let $(T_n^\bullet)$ be a sequence of finite random uniformly pointed simple triangulations of the sphere which converges locally towards an infinite random pointed triangulation $T^\bullet$ of the plane. We already know from Theorem 14 that the number of ends of T belongs to $\{1, 2, \infty\}$ almost surely and from Proposition 48 that the expected degree of the origin in T is equal to 6.

Theorem 55

Let $(T_n^\bullet)$ be a sequence of finite random uniformly pointed simple triangulations of the sphere which converges locally towards an infinite random pointed triangulation $T^\bullet$ of the plane. We suppose that the vertex degrees in T_n are uniformly bounded by some constant $C > 0$, then $T^\bullet$ is almost surely recurrent.

Proof. Follows by combining Corollary 54 together with Proposition 48. □

The initial theorem of [14] is slightly more general since it allows more general planar maps (still with a uniform bound on the vertex degrees) rather than triangulations. To deduce this more general statement one should use standard monotonicity result on recurrence using the theory of electrical networks [23], since this is not the primary goal of this course we leave it aside.

Bibliographical notes. Planar maps are key objects in combinatorics and nowadays probability, see [33] for a range of applications. The finite circle packing theorem has first been discovered by Koebe and then forgotten to be rediscovered by Thurston as a consequence of a result of Andreev hence the name of the theorem. It has had many applications in the theory of planar maps. See the book of Stephenson [43]. The link between circle packing and Riemann uniformization's theorem was conjectured by Thurston and proved in [41]. The uniqueness part of the proof of the finite circle packing theorem is taken from wikipedia and is due to Schramm. The existence part follows the algorithm proposed by Thurston. The "infinite" circle packing theorem, Theorem 51 is due to He and Schramm [28]. Section 6.3, and in particular the proof of the Benjamini–Schramm theorem [14], is adapted from the recent and beautiful paper of Angel–Hutchroft–Nachmias–Ray [7]. These authors also have an alternative proof of Corollary 54 but rather involved (it uses uniform spanning forests), so Question **??** remains "open". The bounded degree assumption in the Benjamini–Schramm has been an obstacle for several years preventing to apply the theorem to limit of uniform random triangulations. This obstruction has recently been overcome by Gurel-Gurevich and Nachmias [26] hence proving that the so-called Uniform Infinite Planar Triangulation is recurrent.

Bibliography

[1] M. Abért, P. Csikvári, P. E. Frenkel, and G. Kun, *Matchings in Benjamini–Schramm convergent graph sequences*, Trans. Amer. Math. Soc., 368 (2016), pp. 4197–4218.

[2] R. Abraham and J.-F. Delmas, *Local limits of conditioned Galton-Watson trees: the infinite spine case*, Electron. J. Probab., 19 (2014), pp. 1–19.

[3] R. Abraham and J.-F. Delmas, *An introduction to galton-watson trees and their local limits*, arxiv:1506.05571, (2015).

[4] D. Aldous and R. Lyons, *Processes on unimodular random networks*, Electron. J. Probab., 12 (2007), pp. no. 54, 1454–1508 (electronic).

[5] D. Aldous and J. M. Steele, *The objective method: probabilistic combinatorial optimization and local weak convergence*, in Probability on discrete structures, vol. 110 of Encyclopaedia Math. Sci., Springer, Berlin, 2004, pp. 1–72.

[6] O. Angel, G. Chapuy, N. Curien, and G. G. Ray, *The local limit of unicellular maps in high genus*, Electron. Commun. Probab., 18 (2013), pp. 1–8.

[7] O. Angel, T. Hutchcroft, A. Nachmias, and G. Ray, *Unimodular hyperbolic triangulations: Circle packing and random walk*, arXiv:1501.04677.

[8] O. Angel and O. Schramm, *Uniform infinite planar triangulation*, Comm. Math. Phys., 241 (2003), pp. 191–213.

[9] A. Avez, *Théorème de Choquet-Deny pour les groupes à croissance non exponentielle*, C. R. Acad. Sci. Paris Sér. A, 279 (1974), pp. 25–28.

[10] A.-L. Barabási and R. Albert, *Emergence of scaling in random networks*, Science, 286 (1999), pp. 509–512.

[11] I. Benjamini and N. Curien, *Ergodic theory on stationary random graphs*, Electron. J. Probab., 17 (2012), pp. 1–20.

[12] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm, *Group-invariant percolation on graphs*, Geom. Funct. Anal., 9 (1999), pp. 29–66.

[13] I. Benjamini, R. Lyons, and O. Schramm, *Unimodular random trees*, Ergodic Theory Dynam. Systems, (2013).

[14] I. Benjamini and O. Schramm, *Recurrence of distributional limits of finite planar graphs*, Electron. J. Probab., 6 (2001), pp. no. 23, 13 pp. (electronic).

[15] P. Billingsley, *Convergence of probability measures, first edition*, Wiley, New-York, (1968).

[16] B. Bollobás, *A probabilistic proof of an asymptotic formula for the number of labelled regular graphs*, European J. Combin., 1 (1980), pp. 311–316.

[17] C. Bordenave, *Notes on random graphs and combinatorial optimization*, http://www.math.univ-toulouse.fr/ bordenave/coursRG.pdf.

[18] C. Bordenave and P. Caputo, *Large deviations of empirical neighborhood distribution in sparse random graphs*, Probab. Theory Related Fields, 163 (2015), pp. 149–222.

[19] L. Bowen, *Periodicity and circle packings of the hyperbolic plane*, Geom. Dedicata, 102 (2003), pp. 213–236.

[20] M. a. Carrasco Piaggio and P. Lessa, *Equivalence of zero entropy and the Liouville property for stationary random graphs*, Electron. J. Probab., 21 (2016), pp. Paper No. 55, 24.

[21] N. Curien and J.-F. Le Gall, *The Brownian plane*, J. Theoret. Probab., 27 (2014), pp. 1249–1291.

[22] F. M. Dekking, *Branching processes that grow faster than binary splitting*, Amer. Math. Monthly, 98 (1991), pp. 728–731.

[23] P. G. Doyle and J. L. Snell, *Random walks and electric networks*, vol. 22 of Carus Mathematical Monographs, Mathematical Association of America, Washington, DC, 1984.

[24] G. Elek, *On the limit of large girth graph sequences*, Combinatorica, 30 (2010).

[25] G. R. Grimmett, *Random labelled trees and their branching networks*, J. Austral. Math. Soc. Ser. A, 30 (1980/81), pp. 229–237.

[26] O. Gurel-Gurevich and A. Nachmias, *Recurrence of planar graph limits*, Ann. Maths, 177 (2013), pp. 761–781.

[27] O. Häggström, *Infinite clusters in dependent automorphism invariant percolation on trees*, Ann. Probab., 25 (1997), pp. 1423–1436.

[28] Z.-X. He and O. Schramm, *Hyperbolic and parabolic packings*, Discrete Comput. Geom., 14 (1995), pp. 123–149.

[29] V. A. Kaimanovich, *Boundary and entropy of random walks in random environment*, in Probability theory and mathematical statistics, Vol. I (Vilnius, 1989), "Mokslas", Vilnius, 1990, pp. 573–579.

[30] V. A. Kaimanovich and F. Sobieczky, *Stochastic homogenization of horospheric tree products*, in Probabilistic approach to geometry, vol. 57 of Adv. Stud. Pure Math., Math. Soc. Japan, Tokyo, 2010, pp. 199–229.

[31] V. A. Kaimanovich and A. M. Vershik, *Random walks on discrete groups: boundary and entropy*, Ann. Probab., 11 (1983), pp. 457–490.

[32] H. Kesten, *Subdiffusive behavior of random walk on a random cluster*, Ann. Inst. H. Poincaré Probab. Statist., 22 (1986), pp. 425–487.

[33] S. K. Lando and A. Zvonkin, *Graphs on surfaces and their applications*, Springer-Verlag, 2004.

[34] L. Lovasz, *Large networks and graph limits*, http://www.cs.elte.hu/ lovasz/bookxx/hombook-almost.final.pdf.

[35] R. Lyons, *Asymptotic enumeration of spanning trees*, Combin. Probab. Comput., 14 (2005), pp. 491–522.

[36] R. Lyons, R. Pemantle, and Y. Peres, *Ergodic theory on Galton-Watson trees: speed of random walk and dimension of harmonic measure*, Ergodic Theory Dynam. Systems, 15 (1995), pp. 593–619.

[37] R. Lyons and Y. Peres, *Probability on Trees and Networks*, Current version available at http://mypage.iu.edu/ rdlyons/, In preparation.

[38] A. Montanari, *Statistical mechanics and algorithms on sparse and random graphs*, http://web.stanford.edu/ montanar/OTHER/STATMECH/stflour.pdf.

[39] J. Neveu, *Arbres et processus de Galton-Watson*, Ann. Inst. H. Poincaré Probab. Statist., 22 (1986), pp. 199–207.

[40] F. Paulin, *Propriétés asymptotiques des relations d'équivalences mesurées discrètes*, Markov Process. Related Fields, 5 (1999), pp. 163–200.

[41] B. Rodin and D. Sullivan, *The convergence of circle packings to the riemann mapping*, J. Differential Geom., 26 (1987), pp. 349–360.

[42] R. P. Stanley, *Enumerative combinatorics. Vol. 2*, vol. 62 of Cambridge Studies in Advanced Mathematics, Cambridge University Press, Cambridge, 1999. With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin.

[43] K. Stephenson, *Introduction to circle packing*, Cambridge University Press, Cambridge, 2005. The theory of discrete analytic functions.

[44] R. van der Hofstad, *Lecture notes random graphs and complex networks*, In preparation, R. van der Hofstad (May 2013). Random graphs and complex networks. http://www.win.tue.nl/ rhofstad/NotesRGCN.pdf, (2013).