Aspects of random maps

Lecture notes of the 2014 Saint-Flour Probability Summer School

preliminary draft

Grégory Miermont[1]

October 9, 2014

[1]ÉNS de Lyon & Institut universitaire de France

Foreword

Random surfaces in physics. The main question motivating this course is the following

> **Question 1:** *what does a uniform random metric on the* 2*-sphere look like?*

This question has its roots in the so-called "quantum gravity" theory from physics[1]. We refer to [5] for an introduction to these aspects. A basic object that arises in this theory is the so-called partition function[2]

$$\int_{\mathcal{R}(M)/\mathrm{Diff}^+(M)} \mathcal{D}[g] \exp(-S(g)) \,, \tag{1}$$

where $\mathcal{R}(M)$ is the space of Riemannian metrics on the 2-dimensional differentiable manifold M, $\mathrm{Diff}^+(M)$ is the set of diffeomorphisms of M acting on $\mathcal{R}(M)$ by pullback[3], $\mathcal{D}g$ is a putative "uniform measure" on $\mathcal{R}(M)$, invariant under the action of $\mathrm{Diff}^+(M)$, and $\mathcal{D}[g]$ is the induced measure on the quotient. Finally, $S(g)$ is an action functional defined in terms of the volume and curvature of the metric, as well as potential "external fields" interacting with the metric: in the simplest, pure form of 2-dimensional gravity, $S(g)$ is simply $\Lambda \cdot \mathrm{vol}_g(M)$ for some constant $\Lambda > 0$. One sees that the mathematically problematic object here is the measure $\mathcal{D}g$, since there is no "uniform"

[1]Bear in mind that the "2-dimensional random metrics" that arise in this context are but a toy model for fluctuations at small scales of the physical 4-dimensional space-time!

[2]In fact, the quantum theory leads one to consider integrals of $\exp(\mathrm{i}S(g))$ rather than $\exp(-S(g))$. The mathematical objects that are involved in these two models are very different, although a formal passage from one to the other, called the Wick rotation in physics, seems relatively common in physics. As far as we know, there is still no rigorous "quantum probabilistic" theory of quantum gravity.

[3]One should see h^*g as a reparametrization of the Riemannian manifold (M, g), since it is isometric to (M, h^*g), hence these two metrics should really be considered the same.

measure on the infinite dimensional space $\mathcal{R}(M)$. Hence, one is naturally lead to Question 1.

At this point, however, it may seem to the reader that there is not much to expect from such a question or from (1). However, a similar situation arises if one asks the following

Question 2: *what does a uniform random path in $\mathbb{R}^d$ look like?*

Even though this question is also ill-posed, there should not be too much debate that the natural answer to it is "Brownian motion". There are several reasons to give a special role to Wiener's measure (the law of Brownian motion), but one of them that is particularly intuitive is that it arises as the *universal scaling limit* of natural discrete analogs of the initial question. Indeed, it is well-known that Brownian motion is the limit in distribution of a uniform random lattice path of length n in $\mathbb{Z}^d$, that is a uniform element of the finite set

$$\{(\omega_0, \omega_1, \dots, \omega_n) \in \mathbb{Z}^d : \omega_0 = 0, |\omega_{i+1} - \omega_i| = 1, 0 \leq i \leq n-1\},$$

where $|\cdot|$ denotes the Euclidean norm in $\mathbb{R}^d$. Of course, this involves a proper rescaling of the path, by a factor n in time and $\sqrt{n}$ in space, hence the name "scaling limit". Moreover, Brownian motion has a universal character: if one replaces lattice paths by random walks (under a natural condition of centering and isotropy), then Brownian motion still arises as the scaling limit.

Hence, one way to understand Question 1 is to find a natural discrete analog of a metric on the sphere[4], study the scaling limit of this model, and accept this limit as the right answer if it universally arises whatever the discrete model is, at least within a reasonable class.

[4]Note also that, rather than asking for an abstract intrinsically defined random metric on $\mathbb{S}^2$ as in Question 1, a natural way to "generalize Brownian motion to two dimensions" would be to ask the following

Question 3: *what does a uniform random 2-sphere immersed in $\mathbb{R}^d$ look like?*

(indeed, loosely speaking, Brownian motion can be seen as a line, for which there is only one intrinsic metric up to reparametrization, immersed in $\mathbb{R}^d$!) The problem considered in Question 1 is called "pure gravity" in physics, describing the fluctuations of a universe with no matter, while immersing random surfaces in $\mathbb{R}^d$ can be understood as defining d "matter fields" on these abstract manifolds. These aspects have been considered in physics, for instance Chapter 3 in [5] but to our knowledge they are far from being well-understood mathematically.

Natural candidates for the discretized version of Question 1 are *maps*, which are graphs that are properly embedded in a 2-dimensional surface. The latter are seen up to “reparametrization”, that is, up to the natural action of homeomorphisms of the underlying surface. In this sense, maps are combinatorial, discrete objects, and they provide a way to endow the surface with a discrete “geometrization”. It is then natural to consider a uniform element in the set of maps on the sphere with n edges (or in another natural class, like triangulations with n triangles), let $n \to \infty$, and renormalize properly the map in order to observe a non-trivial scaling limit.

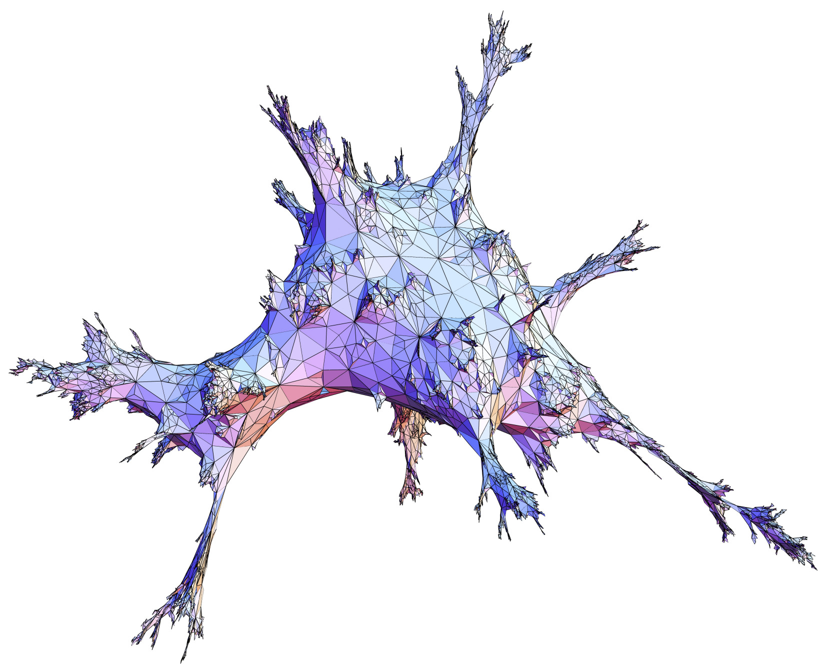

Figure 1: A large random map, simulation by N. Curien

Making sense of the program sketched above will take most of the efforts in this course. One of the difficulties that occurs in this context is that the objects that are involved are rather irregular objects. In the same way as Brownian motion has an arguably very “rough” structure, the random surfaces that arise as limits of random maps are not smooth, Riemannian manifolds or Riemann surfaces. Hence, in order to deal with the geometric aspects of these random objects, one first has to give up a large part of the geometric and analytic arsenal that is usually available, and find what relevant quantities still make sense and are susceptible to pass to appropriate limits.

The first attempts in approaching Question 1 by discretization methods were done by theoretical physicists, starting with Weingarten in the early 1980's, and followed by David, Kazakov, Kostov, Migdal, and many others, see [5] for references. Ambjørn and Watabiki [6] were notably able to compute the so-called *two-point function or pure gravity*, namely the asymptotic probability that two uniformly chosen points in a random map stand at a given graph distance. Their approach was motivated by the rich enumerative theory for maps that was available in the years 1990, after the works by Tutte [79] and his followers, see e.g. [13, 24], and the connection between map enumeration and matrix integrals unveiled by t'Hooft [78], and Brézin, Parisi, Itzykson and Zuber [27].

2002: A Random Metric Space Odyssey. Rigorous derivations of scaling limits for distance functionals in random maps, however, came only in 2002, when the seminal work by Chassaing and Schaeffer [31] was posted in preprint form.[5] They identified the asymptotic distribution of the radius and profile of distances in random quadrangulations. This was made possible by the extensive use of a remarkable bijection between random maps and labeled trees, called the *Cori-Vauquelin-Schaeffer bijection* (CVS bijection), and that we will present in Section 2.3 after recalling some of the basic aspects of maps in Chapter 1 and their enumeration theory at the beginning of Chapter 2. In Chapter 3, we will review the theory of scaling limits of the random labeled trees that appear in the CVS bijection, the theory of which has been developed thoroughly in the years 1990-2000 around ideas of Aldous and Le Gall, see in particular the surveys [51, 53] and references therein. In Chapter 4 we will describe the results by Chassaing and Schaeffer, and start the exploration of the global scaling limit of uniform random quadrangula-

[5]This discussion about "rigorous derivations" calls for a remark. It was considered for a long time that Ambjørn and Watabiki's 1995 computations of the two-point function and its scaling limit were not rigorous. However, it was observed very recently by Timothy Budd (personal communication) that they did derive the exact distance function for a specific model of planar maps, namely uniform trivalent maps with independent exponential edge-lengths, rather than without these edge-lengths as they implicitly claimed in their paper. In retrospect, this sheds a very interesting light on this 20-years old result: from a mathematical point of view, the only problem was that the objects were not identified correctly.

tions. Very important milestones in this topic have been discovered by Le Gall: in particular, he identified the topology of this scaling limit as being that of $\mathbb{S}^2$ in the works [54, 59] (the second with Paulin). We will devote Chapter 5 to the proof of this result. Le Gall also obtained a remarkable description of a family of geodesics in the limiting space [55]. In Chapter 6, we will present an approach of this type of results introduced in [67]. This is based on a generalization of the CVS bijection that allowed Bouttier and Guitter [26] to obtain the three-point function of uniform random quadrangulations. In Chapter 7, we will complete the proof of convergence of rescaled quadrangulations to a limiting space called the Brownian map, a term coined by Marckert and Mokkadem [63]. This requires a rather detailed study of the geodesics of the limiting space that was performed in independent works by Le Gall [56] and the author [68]. Finally, in Chapter 8, we will present an argument due to Le Gall [56], that allows to generalize the convergence result to many other families of maps, showing that the Brownian map is "universal".

The topics presented in these notes leave aside many interesting aspects of random maps, including the scaling limits of random maps on surfaces that are more general than the sphere [29, 17], the scaling limits of maps with large faces [57] and their relations to statistical physics models on random maps [23], as well as all the fascinating aspects of local limits of random maps [8, 49, 48, 30, 64, 36, 14, 46, 9] or the links between scaling and local limits, as developed in [35].

Liouville quantum gravity Let us end this introduction by mentioning that there exists a completely different approach to Question 1, that came roughly simultaneously to the first works based on discretization. This approach, due to Polyakov [75, 47], consists in arguing that the random metric one is looking for should formally be the of the form $e^{\gamma h}|\mathrm{d}z|^2$, where $\gamma > 0$ and h is a Gaussian random field, which in its simplest form is the so-called Gaussian free field. Mathematically however, the latter is a random distribution rather than a pointwise defined function, and its "exponential" is a very ill-defined object. In the recent years, the mathematical grounds of this line of research are starting to being explored by Duplantier and Sheffield [38], starting with understanding that the appropriate notion of measure associated to the ill-defined Riemannian metric above should be understood as one instance of the theory of Kahane's multiplicative chaoses. The connection with the metric aspects and the Brownian map now seems to be accessible via new conformally invariant growth processes called Quantum Loewner Evolutions (QLE) [70] introduced by Miller and Sheffield. Indeed, they announce in this paper the construction of a random metric space out of the Gaussian free field that is (a version of) the Brownian map. Let us mention also a promising line of research consisting in considering Brownian motion in Liouville quantum gravity, which is an instance of a singular time-change of Brownian motion, explored by Garban, Rhodes and Vargas, Berestycki, and others [40, 41, 15, 61, 16, 7].

Trying to bridge more concretely this purely "continuum, random complex-analytic" theory with the discrete theory described in these notes is one of the most exciting challenges in random maps theory. Let us mention that recently, Curien [34] has argued that an exploration of local limits of random maps using Schramm-Loewner evolution with parameter 6 can give precious information on the conformal structure of random planar maps, and a link between such explorations and QLE would be a wonderful achievement. However, this is mostly *terra incognita* so far.

Acknowledgements. Parts of these lecture notes are an updated version of a course given at a Clay Mathematical Institute Summer School in 2010 in Buzios, Brazil, jointly with Jean-François Le Gall. Thanks are due to him for allowing me to use this material, and for many stimulating discussions around random maps. Thanks to the regular participants who have given life to the *journées cartes* that have been held since about 2005, at first in Orsay, and which have now taken a nomadic form. These *journées* have been a constant source of motivation, exchange and inspiration over the years. Thanks to the organizers of the Summer School for continuing this beautiful tradition of the probabilistic community. Saint-Flour, *proba pro nobis*!

Contents

Chapter 1

Background on maps

As a reflection of the fact that maps arise in many different domains of mathematics, there are many equivalent definitions for maps. In most of this course, the maps we will consider will be plane maps, i.e. maps defined on the 2-sphere $\mathbb{S}^2$, however, for the purposes of this chapter it is better to consider maps defined on general surfaces.

1.1 Maps as embedded graphs

1.1.1 Embedded graphs

In these notes, the word "graph" will refer to an unoriented multigraph, formally consisting of a set of vertices V, a set of edges E, and an incidence relation, which is a subset I of $V \times E$ such that for every $e \in E$, the set $\{v \in V : (v, e) \in I\}$ has cardinality 1 or 2 (in the first case we say that e is a self-loop). We usually denote a graph by $G = (V, E)$, without explicit mention of the incidence relation I.

The description of maps that is arguably the most intuitive is in terms of graphs embedded in surfaces. The drawback is that it turns out to be quite elaborate mathematically, and for this reason, let us stress at this point that we will sometimes be sketchy in this section with some definitions. A careful treatment of the theory of embedded graphs is done in the book [71].

Let S be a topological surface that is oriented, compact, connected and without boundary. The fundamental theorem of classification of surfaces asserts that S is homeomorphic to one of the surfaces $\mathbb{S}^2 = \mathbb{T}_0, \mathbb{T}_1, \mathbb{T}_2, \ldots$, where for every $g > 0$, the surface $\mathbb{T}_g$ is the connected sum of g copies of the 2-dimensional torus $\mathbb{T}_1$. The number g is called the *genus* of S.

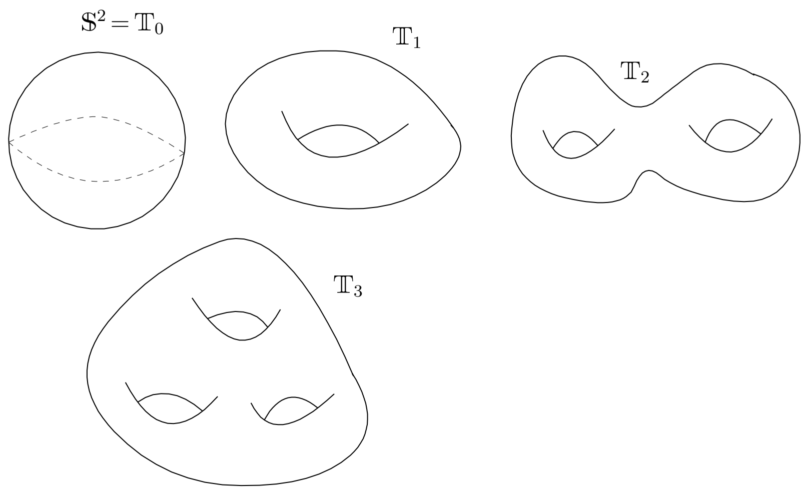

Figure 1.1: Topological surfaces

An *oriented edge* in S is a continuous mapping $e : [0,1] \to S$ such that

- either e is injective
- or the restriction of e to $[0,1)$ is injective and $e(0) = e(1)$.

In the second case, we say that e is a *loop*. We always consider edges up to reparametrization by an increasing homeomorphism $[0,1] \to [0,1]$, and the quantities associated with an edge e are usually invariant under such reparametrization.

The origin of e is $e^- = e(0)$, the target of e is $e^+ = e(1)$, the two together form the extremities of e. The reversal of e is the edge $\overline{e} : t \mapsto e(1-t)$.

An *edge* on S is a pair $\mathbf{e} = \{e, \overline{e}\}$ where e is an oriented edge. The relative interior of an edge $\mathbf{e} = \{e, \overline{e}\}$ is the set $\mathrm{int}(\mathbf{e}) = e((0,1))$, and its extremities are the extremities of e. If E is a set of edges, we let $\vec{E}$ be the associated set of oriented edges, so that $\#\vec{E} = 2\#E$.

Definition 1.1.1. *An embedded graph in S is a multigraph $G = (V(G), E(G))$ such that*

- *$V(G)$ is a finite subset of S,*

- *$E(G)$ is a finite set of edges on S,*
- *for every $\mathbf{e} \in E(G)$, the vertices incident to $\mathbf{e}$ in G are its extremities,*
- *for every $\mathbf{e} \in E(G)$, the relative interior* $\mathrm{int}(\mathbf{e})$ *does not intersect $V(G)$ nor the images of the edges of $E(G)$ distinct from* $\mathbf{e}$.

For simplicity, we usually denote $V(G), E(G)$ by V, E. The *support* of an embedded graph G is the set

$$\mathrm{supp}(G) = V \cup \bigcup_{\mathbf{e} \in E} \mathrm{int}(\mathbf{e}),$$

and a *face* of G is a connected component of $S \setminus \mathrm{supp}(G)$. We let $F(G)$ be the set of faces of G, or simply F if G is clearly given by the context.

Definition 1.1.2. *A map on S is an embedded graph whose faces are all homeomorphic to the unit disk of $\mathbb{R}^2$. A map on $S = \mathbb{S}^2$ is called a plane map.*

A rooted map is a map with a distinguished oriented edge, formally a pair (G, e_*), where G is a map and $e_* \in \vec{E}(G)$.

A map is necessarily a connected graph. In fact, in the case where $S = \mathbb{S}^2$, maps are exactly the connected embedded graphs, but for higher genera, the graph that consists of a single vertex $V = \{v\}$ and no edges $E = \varnothing$ is not a map, because $S \setminus \{v\}$ is not simply connected.

We see that the notion of map does not only depend on the underlying graph structure, that is the graph isomorphism class of G, but also on the way that the graph is embedded in S (and in particular, of S itself). Figure 1.2 shows how the complete graph K_4 can be embedded in $\mathbb{S}^2$ or $\mathbb{T}_1$ to form two different maps.

Let $G = (V, E)$ be a map and $e \in \vec{E}$. Since S is an oriented surface, it makes sense to consider the face of G that lies to the left of e, when one crosses the edge e in its natural orientation from e^- to e^+. We let $f_e \in F(G)$ be this face. Note that it might well happen that $f_{\overline{e}} = f_e$, in which case e is called an *isthmus*. The oriented edges $e_1, e_2, \ldots, e_d$ bounding a face f are naturally arranged in a counterclockwise cyclic order around f, and with every $i \in \{1, \ldots, d\}$ there corresponds a *corner*, namely, an angular sector included in f and bounded by the edges e_{i-1} and e_i, where by convention $e_0 = e_d$ (to be mathematically precise, one should rather speak of germs of

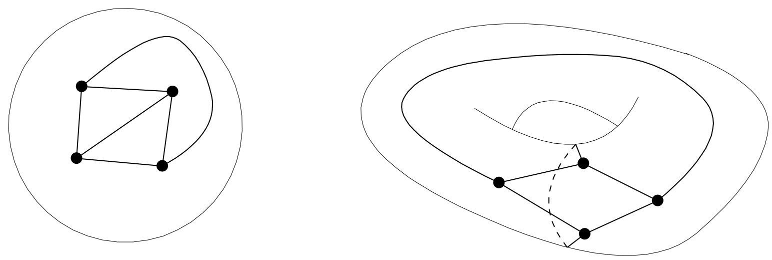

Figure 1.2: Two maps with underlying graph K_4.

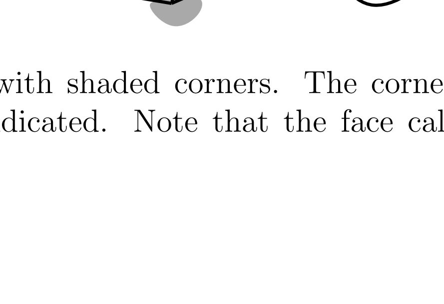

Figure 1.3: A map with shaded corners. The corner corresponding to an oriented edge e is indicated. Note that the face called f is incident to 4 corners.

such sectors). It will be useful in the sequel to identify oriented edges in a map and their incident corners.

The degree of a vertex $v \in V$ is defined by $\deg(v) = \#\{e \in \vec{E} : e^- = v\}$, which in words is the number of edges that are incident to v, counting self-loops twice. The degree of a face $f \in F$ is $\deg(f) = \#\{e \in \vec{E} : f_e = f\}$, and this is the number of corners incident to the face f: if we understand f as being the interior of a topological polygon, then its degree is just the number of edges of this polygon. Note that this is also the number of edges of the map that lie in the closure of f (we say that those edges are incident to f), but with the convention that the isthmuses are counted twice.

1.1.2 Duality and Euler's formula

There is a natural notion of duality for maps. Let $G = (V, E)$ be a map. Let v_f be a point inside each face $f \in F(G)$. For every $\{e, \overline{e}\} \in E$, draw a "dual" edge from v_{f_e} to $v_{f_{\overline{e}}}$ that intersects $\mathrm{int}(e)$ at a single point, and does not intersect $\mathrm{supp}(G)$ otherwise. It is possible to do this in a way such that the dual edges do not cross, so that the graph with vertex set $V^* = \{v_f : f \in F(G)\}$ and edge-set E^* equal to the set of dual edges is an embedded graph G^*, and in fact, a map. See Figure 1.4 for an example, and observe how the roles of self-loops and isthmuses are exchanged by this operation. One can observe also that the roles and degrees of vertices and faces are also exchanged by duality: one has $\deg(v_f) = \deg(f)$ with the above notation, and conversely, with every vertex $v \in V$ one can associate a unique face $f_v \in F(G^*)$ such that $v \in f_v$, and with $\deg(f_v) = \deg(v)$. Proving all of the previous assertions would require a considerable effort, but these will become rather transparent when we introduce the algebraic description of maps.

A very important property of maps is given by Euler's formula.

Theorem 1.1.1. *Let $G = (V, E)$ be a map on the surface $\mathbb{T}_g$. Then*

$$\#V - \#E + \#F = 2 - 2g.$$

The number $\chi = 2 - 2g$ is called the Euler characteristic of $\mathbb{T}_g$.

Example It is not possible to embed the complete graph K_5 in $\mathbb{S}^2$. To see this, note that $\#V(K_5) = 5$ and $\#E(K_5) = 10$. Assume that G is a map

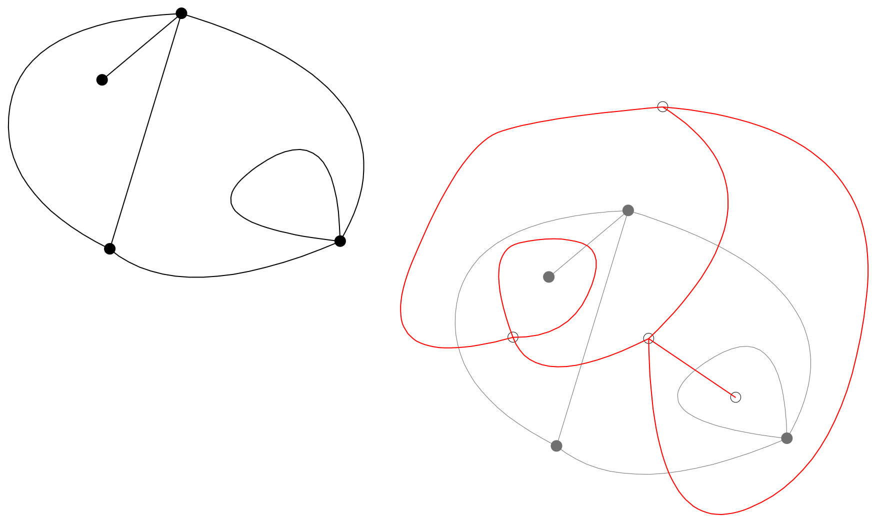

Figure 1.4: Duality

that is isomorphic to K_5 as a graph. By Euler's formula for $g = 0$, it holds that

$$\#F = 10 - 5 + 2 = 7.$$

Since K_5 has no self-loops or multiple edges, all faces in G must have degree at least 3. Thus

$$20 = 2\#E = \#\vec{E} = \sum_{f \in F} \deg(f) \geq 3\#F.$$

Hence $\#F < 7$, a contradiction.

For a similar reason, it is not possible to embed the complete bipartite graph $K_{3,3}$ in $\mathbb{S}^2$.

1.1.3 Isomorphisms

The next important notion for maps is that of isomorphisms. Let G, G' be two maps on surfaces S, S'. The two maps are called isomorphic if there exists an orientation-preserving homeomorphism $h : S \to S'$ such that $V(G') = h(V(G))$ and $\vec{E}(G') = \{h \circ e : e \in \vec{E}(G)\}$. The mapping h is called a map isomorphism. Note that it is in particular a graph isomorphism, in the sense that it preserves the incidence relations between vertices and edges, but it is more than that since it also preserves the incidence relations between edges and faces. If G, G' are maps that are rooted at e_*, e'_*, we say that h is a root-preserving isomorphism if it is an isomorphism such that $h \circ e_* = e'_*$. We say that the rooted maps (G, e_*) and (G', e'_*) are isomorphic if there exists a root-preserving isomorphism sending one to the other.

Intuitively, two maps are isomorphic if one can "distort" the first one into the other, although this is more subtle than that. It is not always the case, at least in genus $g \geq 1$, that one can join two isomorphic maps on the same surface by homotopy. The classical example is given at the bottom part of Figure 1.5. Here, the homeomorphism h is a so-called Dehn twist, obtained by cutting the torus along a generating circle, rotate one side by a full turn, and glue back. If one sees $\mathbb{T}_1$ as $(\mathbb{R}/\mathbb{Z})^2$, such a mapping is given by $(x, y) \mapsto (x, x + y)$.

From now on, we will almost always identify isomorphic (rooted) maps. We will usually adopt bold notation $\mathbf{m}, \mathbf{t}, \ldots$ to denote isomorphism classes of maps. Also, most of the maps we will be interested in will be rooted, although we will usually not specify the root in the definition.

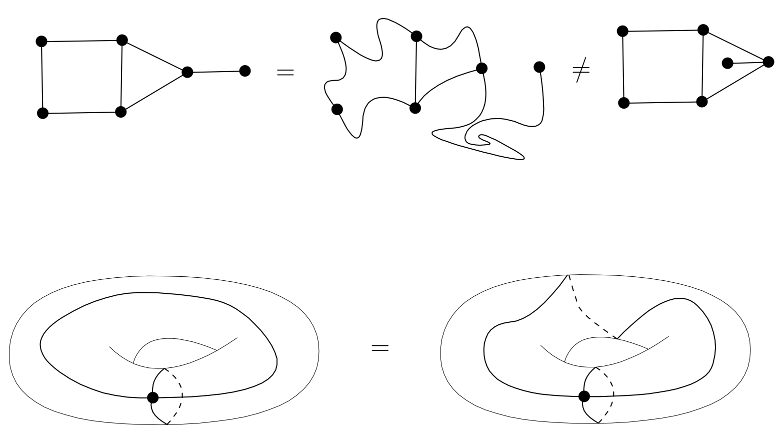

Figure 1.5: Isomorphisms

Finally, a map automorphism is an isomorphism from a map onto itself, that is, a symmetry of the map. The reason for considering rooted maps is the following.

Proposition 1.1.2. *An automorphism of a map that fixes one oriented edge fixes all the edges of the map.*

Intuitively, the reason for this is that an automorphism fixing an oriented edge e must send every oriented edge with origin $v = e^-$ to another such edge, while preserving the cyclic order of these edges. Hence, every edge with origin e^- must be fixed. For the same reason, if e is fixed, then all the edges incident to e^+ must be fixed. By connectedness of the graph, all edges must be fixed.

1.2 Algebraic description of maps

Rather than viewing a map as a graph embedded in a surface, one can also see it as a gluing of polygons. Namely, since cutting a surface S along the edges of a map on S produces a finite number of topological polygons (the faces), one can adopt the opposite viewpoint by considering a finite collection of polygons and glue the edges of these polygons in pairs in order to create a topological surface. One can also decide that the boundaries of the polygons are oriented in counterclockwise order, so that the polygons lie to the left of their oriented edges.

Around each face of $\mathbf{m}$, the incident oriented edges form a cycle, and the collection of these cycles is a permutation φ of $\vec{E}$. This permutation has $\#F$ cycles, with sizes equal to the degrees of the corresponding faces. Since the oriented edges are glued in pairs, namely, e and $\overline{e}$ are glued together with reversed orientation, one can define a second permutation α of $\vec{E}$ whose cycles are $(e, \overline{e})$ for $e \in \vec{E}$. The permutation α is an involution without fixed point. Note that the permutation $\sigma = \alpha\varphi^{-1}$ acts on $\vec{E}$ in the following way. Starting from $e \in \vec{E}$, one goes in clockwise order along f_e, obtaining an oriented edge e', and then one takes the reversal of this edge. Clearly, $\alpha\varphi^{-1}(e)$ is the oriented edge with origin e^- that appears just after e in counterclockwise order around e^-. From this observation, we deduce that σ is a permutation of $\vec{E}$ whose cycles are formed by the oriented edges originating from the vertices of $\mathbf{m}$, arranged in counterclockwise order around these vertices.

Note that if e, e' are two elements of $\vec{E}$, then there exists a word in the letters $\{\sigma, \alpha, \varphi\}$ that sends e to e', due to the connectedness of the map. Otherwise said, the group generated by σ, α, φ acts transitively on $\vec{E}$.

Remarkably enough, a map (considered up to isomorphism) is entirely determined by the permutations σ, α, φ, or by any two of them due to the

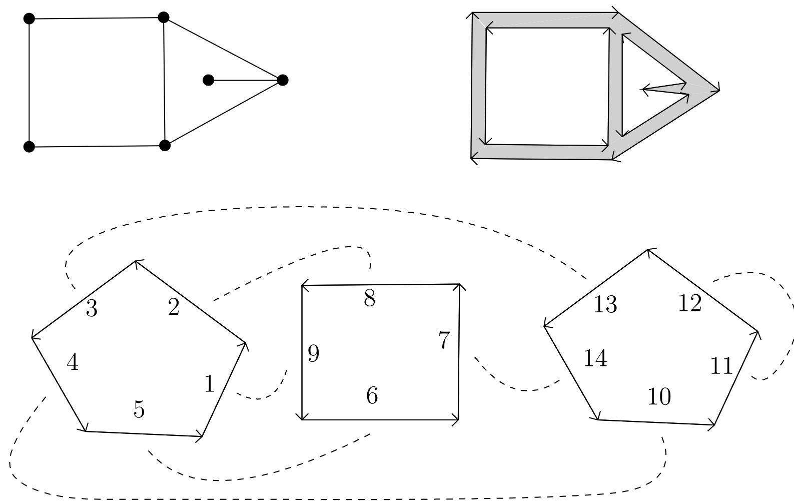

Figure 1.6: Maps are obtained by gluing polygons

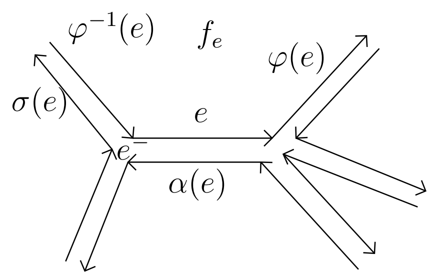

Figure 1.7: The permutations σ, α, φ

identity

$$\varphi\alpha\sigma = 1.$$

In particular, once we know these permutations we can completely forget that they act on a set of oriented edges on a surface, but rather choose to let them act on any finite label set X, and a canonical choice is to take the integers $\{1, 2, \ldots, 2n\}$, where n is the number of edges of the map. For instance, in Figure 1.6, we have

$$\varphi = (1, 2, 3, 4, 5)(6, 7, 8, 9)(10, 11, 12, 13, 14),$$

$$\alpha = (1, 9)(2, 8)(3, 13)(4, 10)(5, 6)(7, 14)(11, 12),$$
$$\sigma = (1, 6)(2, 9)(3, 8, 14)(4, 13, 11)(5, 10, 7)(12).$$

Definition 1.2.1. *Let X be a finite set of even cardinality. A fatgraph structure on X is a triple $(\sigma, \alpha, \varphi)$ of permutations of X such that $\varphi\alpha\sigma = 1$, α is an involution without fixed points, and the subgroup $\langle \sigma, \alpha, \varphi \rangle$ of the permutation group of X generated by σ, α, φ acts transitively on X.*

Two fatgraph structures $(\sigma, \alpha, \varphi)$ on X and $(\sigma', \alpha', \varphi')$ on X' are isomorphic if there exists a bijection $\pi : X \to X'$ such that

$$\pi\sigma = \sigma'\pi, \qquad \pi\alpha = \alpha'\pi, \qquad \pi\varphi = \varphi'\pi.$$

Intuitively, two fatgraphs are isomorphic if they are the same "up to relabeling".

Theorem 1.2.1. *The set of maps considered up to map isomorphisms is canonically identified with the set of fatgraph structures on finite sets considered up to fatgraph isomorphisms.*

The canonical identification is the one we described earlier, which associates with every map $\mathbf{m}$ a fatgraph structure on $\vec{E}(\mathbf{m})$. For a proof of this result (and a fully rigorous description of the canonical identification), see Section 3.2 of [71]. Mathematically, fatgraph structures are of course much more elementary than maps. Another nice aspect of fatgraphs is that certain algebraic properties of maps become rather transparent in this language.

For instance, the duality for maps has a very simple interpretation in terms of fatgraphs. Clearly, $(\sigma, \alpha, \varphi)$ is a fatgraph structure if and only if $(\varphi, \alpha, \alpha\sigma\alpha)$ is a fatgraph structure, and the reader will easily be convinced that the second corresponds to the map that is dual to the one associated with the first.

Consider also Proposition 1.1.2. In terms of a fatgraph $(\sigma, \alpha, \varphi)$ on a set X, an automorphism is just a permutation of X that commutes with σ, α, φ, and hence with every element of $H = \langle \sigma, \alpha, \varphi \rangle$. Now let π be an automorphism of a fatgraph that fixes an element x of X. Since H acts transitively on X, for every $y \in X$, there exists an element $\rho \in H$ such that $\rho(x) = y$. But then,

$$\pi(y) = \pi(\rho(x)) = \rho(\pi(x)) = \rho(x) = y.$$

Hence, π is the identity of X. This is the content of Proposition 1.1.2.

Note that in the context of fatgraph structures, the analog of a rooted map is a fatgraph structure on X together with a distinguished element $x_* \in X$. Two such rooted fatgraph are of course isomorphic if there exists a fatgraph isomorphism that preserves the respective distinguished elements.

1.3 Plane trees

A graph-theoretic tree is a finite graph $G = (V, E)$ that satisfies any of the following equivalent conditions:

- G is connected and has no cycle
- G is connected and has $\#V = \#E + 1$
- G has no cycle and has $\#V = \#E + 1$
- For every $u, v \in V$, there exists a unique simple path between u and v.

A plane tree is a map which, as a graph, is a tree. Note that, due to the absence of cycles, a tree $\mathbf{t}$ has a contractible support, and therefore this implies that the surface S on which $\mathbf{t}$ lives is $\mathbb{S}^2$. Moreover, by the Jordan curve theorem, the absence of cycles is clearly equivalent to the fact that $\mathbf{t}$ has a unique face.

Proposition 1.3.1. *A plane tree is a map with one face.*

In terms of the fatgraph representation, a tree is a gluing of a $2n$-gon by identifying its sides in pairs, in order to produce a surface homeomorphic to a sphere. Equivalently, if we label the edges of the $2n$-gon with the integers $1, 2, \ldots, 2n$, the identifications must produce a *non-crossing matching* of $\{1, 2, \ldots, 2n\}$.

In purely algebraic terms, a tree is a factorization of the cyclic permutation $(1, 2, \ldots, 2n)$ as a product $\sigma^{-1}\alpha$, where α is an involution without fixed points, and σ has $n+1$ cycles. Note that the transitivity hypothesis is automatically satisfied in this case.

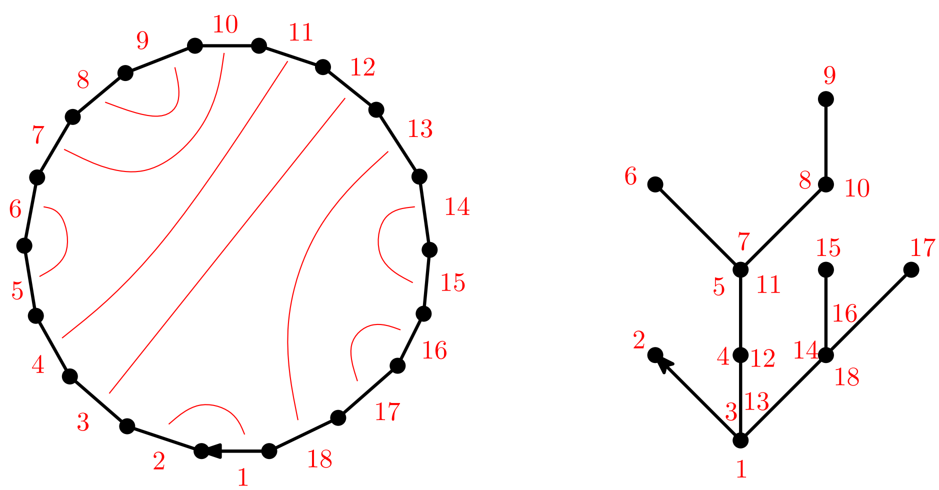

$\sigma = (1,3,13)(2)(4,12)(5,7,11)(6)(8,10)(9)(14,16,18)(15)(17)$

$\alpha = (1,2)(3,12)(4,11)(5,6)(7,10)(8,9)(13,18)(14,15)(16,17)$

Figure 1.8: Gluing the edges of a polygon along a non-crossing matching produces a tree. The fatgraph representation is displayed. Note that on the left picture, integers from 1 to 18 represent half-edges, while on the right picture, we used integers to represent the corresponding corners.

There are many other possible natural encodings for trees. The so-called Ulam-Harris encoding of rooted plane trees is arguably the simplest mathematically, and it stresses the fact that rooted trees are the mathematical object corresponding to genealogical trees. It identifies the set of vertices of a rooted tree $\mathbf{t}$ with a subset of the set

$$\mathcal{U} = \bigcup_{n \geq 0} \mathbb{N}^n$$

of integer words, where $\mathbb{N}^0 = \{\varnothing\}$ consists only on the empty word. The root vertex is the word $\varnothing$, and the neighbors of the latter are labeled $1, 2, \dots, k_\varnothing$ from left to right (the root is then the oriented edge from $\varnothing$ to 1). More generally, a non-root vertex $u = u_1 \dots u_n$ of the tree is a neighbor to its parent $\neg u = u_1 \dots, u_{n-1}$, located "below" u, and its k_u children labeled $u1, \dots, uk_u$ from left to right "above" u. We will not use much this encoding, however, we will keep the notation $\neg u$ for the parent of a non-root vertex u in a plane tree $\mathbf{t}$ (the neighbor of u that lies closest to the root vertex).

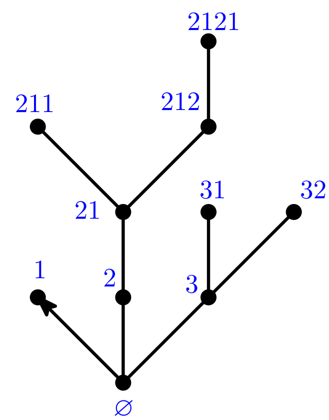

Figure 1.9: The Ulam-Harris encoding of the rooted tree of Figure 1.8

One that will be of particular importance to us is the encoding by *contour functions*. Namely, let $e_0, e_1, \dots, e_{2n-1}$ be the sequence of the oriented edges bounding the unique face of a tree $\mathbf{t}$, starting with the root edge e_0, and where $n = \#E(\mathbf{t})$. We call this the contour exploration of $\mathbf{t}$. Then let $u_i = e_i^-$ denote the i-th visited vertex in the contour exploration, and set

$$C_{\mathbf{t}}(i) = d_{\mathbf{t}}(u_0, u_i), \quad 0 \leq i < 2n,$$

be the height of u_i (distance to the root vertex). By convention, let $e_{2n} = e_0, u_{2n} = u_0$, and $C_{\mathbf{t}}(2n) = 0$. It is natural to extend $C_{\mathbf{t}}$ by linear interpolation between integer times: for $0 \leq s < 2n$, we let

$$C_{\mathbf{t}}(s) = (1 - \{s\})C_{\mathbf{t}}(\lfloor s \rfloor) + \{s\}C_{\mathbf{t}}(\lfloor s \rfloor + 1),$$

where $\{s\} = s - \lfloor s \rfloor$ is the fractional part of s.

The contour $C_{\mathbf{t}}$ is then a non-negative path of length $2n$, starting and ending at 0, with slope ± 1 between integer times. We call such a path a *discrete excursion* of length $2n$. Conversely, it is not difficult to see that any discrete excursion is the contour process of a unique rooted plane tree $\mathbf{t}$ with n edges. More specifically, the trivial path with length 0 corresponds to the vertex map, and inductively, if C is a discrete excursion of positive length, then we let $C^{(1)}, \ldots, C^{(k)}$ be the excursions of C above level 1, and the tree encoded by C is obtained by taking the rooted trees $\mathbf{t}_{(1)}, \ldots, \mathbf{t}_{(k)}$ encoded by $C^{(1)}, \ldots, C^{(k)}$ and linking their root vertices in this cyclic order to a same vertex v by k edges, and finally rooting the map at the edge going from v to the root of $\mathbf{t}_{(1)}$.

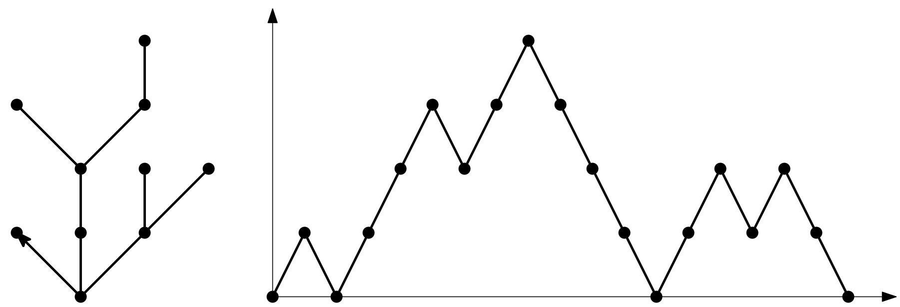

Figure 1.10: The contour process of the same rooted tree

Notes for Chapter 1

We have given two equivalent points of view on maps: a "topological" point of view, and an "algebraic" point of view. See [71] for a rather complete treatment of these questions, as well as for many developments on actual plane embeddings of graphs, embeddability problems, the 4-color theorem, and other topics.

There exist other ways to look at maps, and in particular, a "geometric" point of view that is better suited to the links between maps and complex structures (Riemann surfaces). This point of view introduces maps as coverings of the sphere $\mathbb{S}^2 = \hat{\mathbb{C}}$ that are ramified at exactly three points (say $0, 1, \infty$). One of the appeals of this approach is that it allows to select canonically a distinguished embedding of every map in $\mathbb{R}^2$, called its "true shape" or "dessin d'enfant". We refer the reader to Chapters 1 and 2 of [50] for an introduction to this fascinating point of view.

Chapter 2

A quick introduction to the enumeration of maps

The enumerative theory of maps has a rich history, which is one of the reasons of the success of the theory of random maps. We will quickly mention some cornerstones of this theory, focusing mostly on the recent bijective methods.

2.1 Tutte's quadratic method

The enumeration of maps starts with the work of Tutte, who approached the problem of enumeration of plane maps by solving the recursive equations satisfied by the associated generating functions. Suppose that one wants to find the cardinality of the set $\mathcal{M}_n$ of rooted plane maps with n edges. This can be achieved by finding an explicit form for the formal power series

$$M(x) = \sum_{\mathbf{m} \in \mathcal{M}} x^{\# E(\mathbf{m})} = \sum_{n \geq 0} \# \mathcal{M}_n x^n \in \mathbb{Q}[[x]],$$

where $\mathcal{M} = \bigcup_{n \geq 0} \mathcal{M}_n$. In order to do this, one should first try to find a functional equation for M. To this end, it is useful to introduce a second "catalytic" variable y that counts the degree of the root face, and set

$$M(x, y) = \sum_{\mathbf{m} \in \mathcal{M}} x^{\# E(\mathbf{m})} y^{\deg(f_*)},$$

where $f_* = f_{e_*}$ denotes the root face of $\mathbf{m}$. Note in particular that $M(x) = M(x, 1)$, and that the coefficient of y^m in this formal power series is the power series $[y^m]M(x,y) \in \mathbb{Q}[[x]]$ counting maps with root face degree equal to m, with a weight x per edge.

If a map $\mathbf{m}$ is not reduced to the vertex map, we can decompose it by removing its root edge. If the latter is an isthmus, it separates the map into two new maps $\mathbf{m}_1, \mathbf{m}_2$ that can be canonically rooted in such a way that the degrees of their root faces add up to $\deg(f_*) - 2$ (the isthmus of $\mathbf{m}$ counts twice in the computation of $\deg(f_*)$), and with a total of $\#E(\mathbf{m}) - 1$ edges. If e_* is not an isthmus, then $\mathbf{m}$ with e_* removed is a new map that can be canonically rooted in such a way that the root face is $f_{e_*} \cup f_{\overline{e}_*} \cup \mathrm{int}(e_*)$. See the next picture.

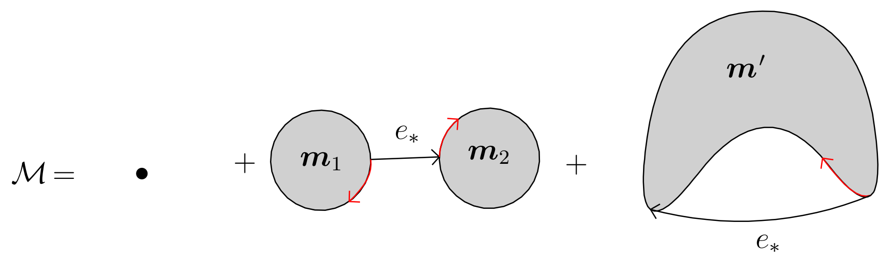

Figure 2.1: The recursive decomposition of a rooted map

This translates into the following equation for $M(x,y)$:

$$\begin{aligned} M(x,y) &= 1 + xy^2 M(x,y)^2 + x \sum_{m \geq 0} ([y^m]M(x,y)(\, y + y^2 + \cdots + y^{m+1}) \\ &= 1 + xy^2 M(x,y)^2 + xy \frac{M(x) - yM(x,y)}{1-y}. \end{aligned}$$

The first term 1 accounts for the fact that the vertex map has no edge, and root face degree 0. The next term accounts for the choice of $\mathbf{m}_1$ and $\mathbf{m}_2$, which participates by the term $M(x,y)^2$, and the fact that the resulting map obtained by bridging the two maps has one more edge, and root face degree increased by 2, hence the factor xy^2. The last term is a bit more complicated: $[y^m]M(x,y)$ counts maps $\mathbf{m}'$ with root face degree m, as we saw earlier (and we sum over all possible values of m which corresponds to taking the union over all possibilities). The factor x accounts for the root edge that one adds (in such a way that the root vertex of $\mathbf{m}$ and $\mathbf{m}'$ coincide) to obtain $\mathbf{m}$, and the factor $(y + y^2 + \cdots + y^{m+1})$ corresponds to the possible degrees of the resulting root face of $\mathbf{m}$, depending on where the target of the root edge is located on the root face of $\mathbf{m}'$. If e_* is a self-loop then the degree of the root face is $m+1$, if it parallels the root edge of $\mathbf{m}'$ then the degree of the root face is m, and so on.

This equation is a quadratic equation of the form

$$(a_1 M(x,y) + a_2)^2 = a_{3,}$$

where a_1, a_2, a_3 are formal power series in the variables x, y, that tacitly depend on the unknown series $M(x)$. The idea of Tutte is to introduce a parameterization of y in terms of x, $y = \alpha(x)$ where α is a formal power series, along which $a_3(x, \alpha(x)) = 0$. If this is the case, then since a_3 is a square, we will have not only $a_3 = 0$ but also

$$\frac{\partial a_3}{\partial y} = 0,$$

which gives two equations in the unknown series $\alpha(x)$ and $M(x)$. The rest of the story is rather technical, so we will not develop it here, but one can already note that at this point, the catalytic variable y has disappeared, and the problem boils down to computing $M(x)$ rather than $M(x,y)$, which in any case was our initial concern. This is one of the little miracles of this method. It ultimately allows to find an implicit parameterization of $M(x)$ of the following form. Let $\theta = \theta(x)$ be the formal power series such that

$$\theta(x) = \frac{x}{1 - 3\theta(x)},$$

then it holds that

$$M(x) = \frac{1 - 4\theta(x)}{(1 - 3\theta(x))^2}.$$

The so-called *Lagrange inversion formula* then allows to compute explicitly the coefficients of M, and one finds

$$\#\mathcal{M}_n = [x^n]M(x) = \frac{2}{n+2} \cdot \frac{3^n}{n+1} \binom{2n}{n}. \tag{2.1}$$

The quadratic method and its generalizations are discussed in depth in [39, 24, 42] for instance. These are powerful and systematic enumeration methods, with the drawback that the implementation of the method itself is a kind of "blackbox" giving no real insight into the specific structure of the objects that are involved. Furthermore, it does not provide an explanation of the reason why a formula like (2.1) is so strikingly simple. This is one of the reasons why one can desire *a posteriori* to get bijective explanations for such formulas, as we will do shortly.

2.2 Matrix integrals

The easiest connection between Gaussian random variables and combinatorics arises when one computes the moments of a standard Gaussian random variable $h \sim \mathcal{N}(0,1)$. Indeed

$$E[h^k] = \#\,\mathrm{Matchings}(k)$$

where Matchings(k) is the number of perfect matchings on a set with k elements, or equivalently of permutations of $\mathfrak{S}_k$ that are involutions without fixed points. In particular, we see that Matchings(k) is empty when k is odd, which corresponds to the fact that $E[h^k] = 0$ since the Gaussian distribution is symmetric. When $k = 2p$ is even, this number is

$$E[h^{2p}] = (2p-1)(2p-3)\cdots 3\cdot 1 = (2p-1)!!$$

Using the fatgraph representation we can easily associate a map with a matching α, viewed as an involution without fixed points in $\mathfrak{S}_{2n}$. Let $\varphi = (1, 2, \ldots, 2n)$ be the cycle of length $2n$. Then clearly $(\alpha\varphi^{-1}, \alpha, \varphi)$ is a fatgraph structure. The map it corresponds to consists in gluing in pairs the edges of a $2n$-gon, labeled as $1, 2, \ldots, 2n$ in counterclockwise order, according to α. The gluing should be made in such a way that the resulting surface is oriented. Yet otherwise said, a matching is canonically associated with a map with one face, and which is also rooted (the role of the oriented edge labeled 1 is distinguished, and if one knows where this oriented edge is, one can recover all the other labels by going around the unique face of the map).

By duality, we can also view matchings as maps with only one vertex.

Note that the association between a matching and a map does not give a straightforward way to see the genus of the resulting map, other than by computing the number of cycles of σ and applying Euler's formula.

It was observed by physicists that higher-dimensional Gaussian calculus allows one to enumerate maps by keeping track of the genus. Let $N \geq 1$ be an integer, and $\mathcal{H}_N$ be the space of $N \times N$ Hermitian matrices. This is a vector space over $\mathbb{R}$ with dimension N^2, since an element $H = (h_{ij})_{1 \le i,j \le n}$ of $\mathcal{H}_N$ is determined by the real diagonal entries $h_{ii}, 1 \le i \le n$ and by the real and imaginary parts of the upper-diagonal $\Re(h_{ij}), \Im(h_{ij}), 1 \le i < j \le n$. It is naturally endowed with the scalar product $\langle A, B \rangle = \mathrm{Tr}(AB^*) = \mathrm{Tr}(AB)$, which corresponds to the Euclidean scalar product up to the identification of $\mathcal{H}_N$ with $\mathbb{R}^{N^2}$ discussed above. In particular

$$\|H\|^2 = \mathrm{Tr}(H^2) = \sum_{i=1}^{n} h_{ii}^2 + 2 \sum_{1 \le i < j \le n} (\Re(h_{ij})^2 + \Im(h_{ij})^2).$$

There is a natural Gaussian measure associated with this Euclidean structure, given by

$$\gamma_N(\mathrm{d}H) = \frac{1}{Z_N} \exp(-\mathrm{Tr}(H^2)/2)\mathrm{d}H,$$

where $\mathrm{d}H = \prod_{i=1}^{n} \mathrm{d}h_{ii} \prod_{1 \le i < j \le n} \mathrm{d}\Re(h_{ij})\mathrm{d}\Im(h_{ij})$, and where

$$Z_N = 2^{N/2} \pi^{N^2/2}.$$

Under this measure, the entries $h_{ii}, 1 \le i \le n$ and $\Re(h_{ij}), \Im(h_{ij}), 1 \le i < j \le N$ are all independent real centered Gaussian random variables, with $h_{ii}, 1 \le i \le n$ having variance 1 and $\Re(h_{ij}), \Im(h_{ij}), 1 \le i < j \le N$ having variance 1/2.

The analog of moments in this setting are averages of the form

$$E\left[\prod_{i=1}^{n} \mathrm{Tr}(H^{k_i})\right]$$

where $k_1, \dots, k_n$ are integers. Suppose for instance that $n = 1$ and that k_1 is even, for this reason, we write it $k_1 = 2k$. Then

$$E[\mathrm{Tr}(H^{2k})] = E\left[\sum_{i_1, i_2, \dots, i_{2k}} h_{i_1 i_2} h_{i_2 i_3} \dots h_{i_{2k-1} i_{2k}} h_{i_{2k} i_1}\right].$$

At this point, one can take the sum out of the expectation by linearity, and use the well-known Wick formula.

Lemma 2.2.1 (Wick formula)**.** *Let $(X_i, i \in I)$ be a centered complex Gaussian vector, where I is a finite index-set with say $\#I = 2k$. Then*

$$E\left[\prod_{i\in I} X_i\right] = \sum_{\alpha=(i_1j_1)\dots(i_kj_k)\in\mathrm{Matchings}(I)} \prod_{r=1}^{k} E[X_{i_r}X_{j_r}],$$

where Matchings(I) *is the set of perfect matchings of I.*

In our context, we obtain that $E[\mathrm{Tr}(H^{2k})]$ is a sum of terms of the form $E[h_{i_1i_2}h_{i_2i_3}\dots h_{i_{2k-1}i_{2k}}h_{i_{2k}i_1}]$, for indices $i_1, i_2, \dots, i_{2k}$ in $\{1, 2, \dots, N\}$. Applying the Wick formula for the index set $I = \{(i_1, i_2), (i_2, i_3), \dots, (i_{2k}, i_1)\}$, we see that the expectation in question is a sum over all perfect matchings of I of a product of terms $E[h_{ij}h_{kl}]$, where the indices (i, j) and (k, l) are matched. The crucial observation is that many of these terms are in fact equal to 0, due to the easily checked fact that

$$E[h_{ij}h_{kl}] = \delta_{il}\delta_{jk}.$$

Assuming for instance that (i_1, i_2) is matched with (i_4, i_5), means that $i_1 = i_5$ and $i_2 = i_4$. Consider a $2k$-gon with vertices labeled $i_1, i_2, \dots, i_{2k}$ in counterclockwise order, corresponding to incident edges $i_1i_2, \dots, i_{2k-1}i_{2k}, i_{2k}i_1$. We interpret the presence of a factor $E[h_{i_ai_{a+1}}h_{i_bi_{b+1}}]$ (where the addition in indices is taken modulo $2k$) by the fact that the edges i_ai_{a+1} and i_bi_{b+1} are glued together with reversed orientation, corresponding to $i_a = i_{b+1}$ and $i_b = i_{a+1}$. Hence, every term of the sum can be represented as a map with one face of degree $2k$. It remains to sum over the different indices $i_1, i_2, \dots, i_{2k}$, taking into account that several of them are constrained to be equal. A moment's thought shows that the number of degrees of freedom (i.e. of free indices) is equal to the number of vertices of the resulting map. Since each index can take N values, we obtain that

$$E[\mathrm{Tr}(H^{2k})] = \sum_{\mathbf{m}\in\mathcal{M}(2k)} N^{V(\mathbf{m})} = N^{k+1}\sum_{g\geq 0} \frac{\#\mathcal{M}^{(g)}(2k)}{N^{2g}},$$

where $\mathcal{M}(2k)$ is the set of rooted maps (on arbitrary surfaces) with exactly one face of degree $2k$, and $\mathcal{M}^{(g)}(2k)$ is the subset formed by those maps which have genus g. Here, we have used Euler's formula and the fact that maps in $\mathcal{M}(2k)$ have k edges and one face. In this way, we see that the squared inverse dimension N^{-2} serves as a variable counting the genus.

A similar reasoning shows that

$$E\left[\prod_{i=1}^{n}\mathrm{Tr}(H^{k_i})\right] = \sum_{\mathbf{m}\in\mathcal{M}(k_1,\ldots,k_n)} N^{V(\mathbf{m})} = N^{k-n+2}\sum_{g\geq 0}\frac{\#\mathcal{M}^{(g)}(k_1,\ldots,k_n)}{N^{2g}},$$

where $\mathcal{M}(k_1,\ldots,k_n)$ is the set of maps[1] with n labeled faces $f_1,\ldots,f_n$ of respective degrees $k_1,\ldots,k_n$, each face coming with a distinguished incident oriented edge, and $\mathcal{M}^{(g)}(k_1,\ldots,k_n)$ is the subset of those maps with genus g. We also let $k=(k_1+\ldots+k_n)/2$ be the number of edges of maps in $\mathcal{M}(k_1,\ldots,k_n)$. See for instance [73] for a beautiful use of this identity in the context of spectral asymptotics of random matrices and the relation with the statistics of random permutations.

We refer the reader to the excellent reference [81] for a nice introduction to map enumeration *via* matrix integrals.

2.3 The Cori-Vauquelin-Schaeffer bijection

Motivated by the very simple form of the formula (2.1) enumerating $\mathcal{M}_n$, Cori and Vauquelin [33] gave in 1981 a bijective approach to this formula. These approaches reached their full power with the work of Schaeffer starting in his 1998 thesis [76]. We now describe the bijective approach in the case of quadrangulations.

Quadrangulations

A map $\mathbf{q}$ is a quadrangulation if all its faces are of degree 4. We let $\mathcal{Q}_n$ be the set of all (rooted) quadrangulations with n faces. Quadrangulations are a very natural family of maps to consider, in virtue of the fact that there exists a "trivial" bijection between $\mathcal{M}_n$ and $\mathcal{Q}_n$, which can be described as follows.

Let $\mathbf{m}$ be a map with n edges, and imagine that the vertices of $\mathbf{m}$ are colored in black. We then create a new map by adding inside each face of $\mathbf{m}$ a white vertex, and by joining this white vertex to every corner of the

[1] There is an important caveat here, which is that the maps considered in this situation are defined on arbitrary compact orientable surfaces without boundary, that are not necessarily connected. Hence, the maps in question can also have several connected components.

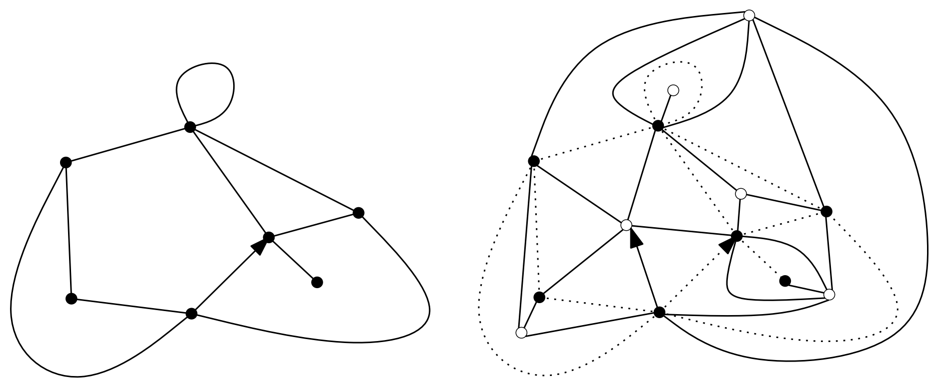

Figure 2.2: The so-called "trivial" bijection

face f it belongs to, by non-intersecting edges inside the face f. In doing so, notice that some black vertices may be joined to the same white vertex with several edges. Lastly, we erase the interiors of the edges of the map $\mathbf{m}$. We end up with a map $\mathbf{q}$, which is a plane quadrangulation with n faces, each face containing exactly one edge of the initial map. We adopt a rooting convention, for instance, we root $\mathbf{q}$ at the first edge coming after e in counterclockwise order around e^-, where e is the root of $\mathbf{m}$.

Notice that $\mathbf{q}$ also comes with a bicoloration of its vertices in black and white, in which two adjacent vertices have different colors. This says that $\mathbf{q}$ is *bipartite*, and as a matter of fact, every (plane!) quadrangulation is bipartite. So this coloring is superfluous: one can recover it by declaring that the black vertices are those at even distance from the root vertex of $\mathbf{q}$, and the white vertices are those at odd distance from the root vertex.

Conversely, starting from a rooted quadrangulation $\mathbf{q}$, we can recover a bipartite coloration as above, by declaring that the vertices at even distance from the root edge are black. Then, we draw the diagonal linking the two black corners incident to every face of $\mathbf{q}$. Finally, we remove the interior of the edges of $\mathbf{q}$ and root the resulting map $\mathbf{m}$ at the first outgoing diagonal from e^- in clockwise order from the root edge e of $\mathbf{q}$. One checks that this is indeed a left- and right-inverse of the previous mapping from $\mathcal{M}_n$ to $\mathcal{Q}_n$. See Fig. 2.2 for an illustration of these bijections.

For the record, we state the following useful fact.

Proposition 2.3.1. *A plane map is bipartite if and only if its faces all have even degree.*

The CVS bijection

Recall that $\mathcal{Q}_n$ is the set of all rooted plane quadrangulations with n faces. A simple application of Euler's formula shows that any element of $\mathcal{Q}_n$ has $2n$ edges ($4n$ oriented edges, 4 for each face) and $n+2$ vertices.

Let $\mathbf{t}$ be a rooted plane tree, with root edge e_0 and root vertex $u_0 = e_0^-$. An *admissible label function* on $\mathbf{t}$ is a function $\ell : V(\mathbf{t}) \to \mathbb{Z}$, such that $\ell(u_0) = 0$ and

$$|\ell(u) - \ell(v)| \leq 1\,, \qquad \text{for every adjacent } u, v \in V(\mathbf{t})\,.$$

Let $\mathbb{T}_n$ be the set of all pairs $(\mathbf{t}, \ell)$, where $\mathbf{t}$ is a rooted plane tree with n edges, and ℓ is an admissible label function.

Let $e_0, e_1, \ldots, e_{2n-1}$ be the contour exploration of the oriented edges of $\mathbf{t}$, and $u_i = e_i^-$, as in Section 1.3. We extend the sequences (e_i) and (u_i) to infinite sequences by $2n$-periodicity. With each oriented edge e_i, we can associate a corner around u_i, as explained in Section 1.1. We will often identify the oriented edge e_i with the associated corner, and we adopt the notation $\ell(e_i) = \ell(u_i)$.

For every $i \geq 0$, we define the *successor* of i by

$$s(i) = \inf\{j > i : \ell(e_j) = \ell(e_i) - 1\}\,,$$

with the convention that $\inf \varnothing = \infty$. Note that $s(i) = \infty$ if and only if $\ell(e_i)$ equals $\min\{\ell(v) : v \in V(\mathbf{t})\}$. This is a simple consequence of the fact that the integer-valued sequence $(\ell(e_i), i \geq 0)$ can decrease only by taking unit steps.

Consider a point v_* in $\mathbb{S}^2$ that does not belong to the support of $\mathbf{t}$, and denote by e_∞ a corner around v_*, i.e. a small neighborhood of v_* with v_* excluded, not intersecting the corners $e_i, i \geq 0$. By convention, we set

$$\ell(v_*) = \ell(e_\infty) = \min\{\ell(u) : u \in V(\mathbf{t})\} - 1.$$

For every $i \geq 0$, the successor of the corner e_i is then defined by

$$s(e_i) = e_{s(i)}\,.$$

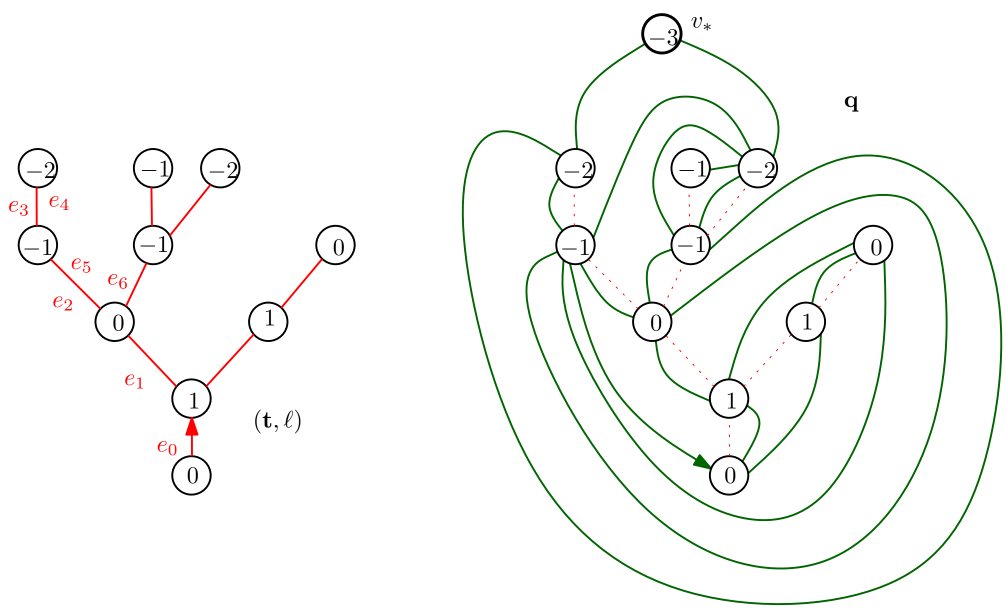

Figure 2.3: Illustration of the Cori-Vauquelin-Schaeffer bijection, in the case $\epsilon = 1$. For instance, e_3 is the successor of e_0, e_2 the successor of e_1, and so on.

The CVS construction consists in drawing, for every $i \in \{0, 1, \ldots, 2n-1\}$, an *arc*, which is an edge from the corner e_i to the corner $s(e_i)$ inside $\mathbb{S}^2 \setminus (\{v_*\} \cup \mathrm{supp}\,(\mathbf{t}))$. See Fig.2.3 for an illustration of the CVS construction.

Lemma 2.3.2. *It is possible to draw the arcs in such a way that the graph with vertex-set $V(\mathbf{t}) \cup \{v_*\}$ and edge-set consisting of the edges of $\mathbf{t}$ and the arcs is an embedded graph.*

Proof. Since $\mathbf{t}$ is a tree, we can see it as a map with a unique face $\mathbb{S}^2 \setminus \mathrm{supp}\,(\mathbf{t})$. The latter can in turn be seen as an open polygon, bounded by the edges $e_0, e_1, \ldots, e_{2n-1}$ in counterclockwise order. Hence, the result will follow if we can show that the arcs do not cross, i.e. that it is not possible to find pairwise distinct corners $e^{(1)}, e^{(2)}, e^{(3)}, e^{(4)}$ that arise in this order in the cyclic order induced by the contour exploration, and such that $e^{(3)} = s(e^{(1)})$ and $e^{(4)} = s(e^{(2)})$.

If this were the case, then we would have $\ell(e^{(2)}) \geq \ell(e^{(1)})$, as otherwise the successor of $e^{(1)}$ would be between $e^{(1)}$ and $e^{(2)}$. Similarly, $\ell(e^{(3)}) \geq \ell(e^{(2)})$. But by definition, $\ell(e^{(3)}) = \ell(e^{(1)}) - 1$, giving $\ell(e^{(2)}) \geq \ell(e^{(3)}) + 1 \geq \ell(e^{(2)}) + 1$, which is a contradiction. □

We call $\mathbf{q}$ the graph with vertex-set $V(\mathbf{t}) \cup \{v_*\}$ and edge-set formed by the arcs, now excluding the (interiors of the) edges of $\mathbf{t}$.

Lemma 2.3.3. *The embedded graph $\mathbf{q}$ is a quadrangulation with n faces.*

Proof. First we check that $\mathbf{q}$ is connected, and hence is a map. But this is obvious since the consecutive successors of any given corner e, given by $e, s(e), s(s(e)), \ldots$, form a finite sequence ending at e_∞. Hence, every vertex in $\mathbf{q}$ can be joined by a chain to v_*, and the graph is connected.

To check that $\mathbf{q}$ is a quadrangulation, let us consider an edge of $\mathbf{t}$, corresponding to two oriented edges $e, \overline{e}$. Let us first assume that $\ell(e^+) = \ell(e^-) - 1$. Then, the successor of e is incident to e^+ and the preceding construction gives an arc starting from e^- (more precisely from the corner associated with e) and ending at e^+. Next, let e' be the corner following $\overline{e}$ in the contour exploration around $\mathbf{t}$. Then $\ell(e') = \ell(e^-) = \ell(\overline{e}) + 1$, giving that $s(\overline{e}) = s(s(e'))$. Indeed, $s(e')$ is the first corner coming after e' in contour order and with label $\ell(e') - 1 = \ell(e) - 1$, while $s(s(e'))$ is the first corner coming after e' with label $\ell(e) - 2$. Therefore, it has to be the first corner coming after $\overline{e}$, with label $\ell(e) - 2 = \ell(\overline{e}) - 1$.

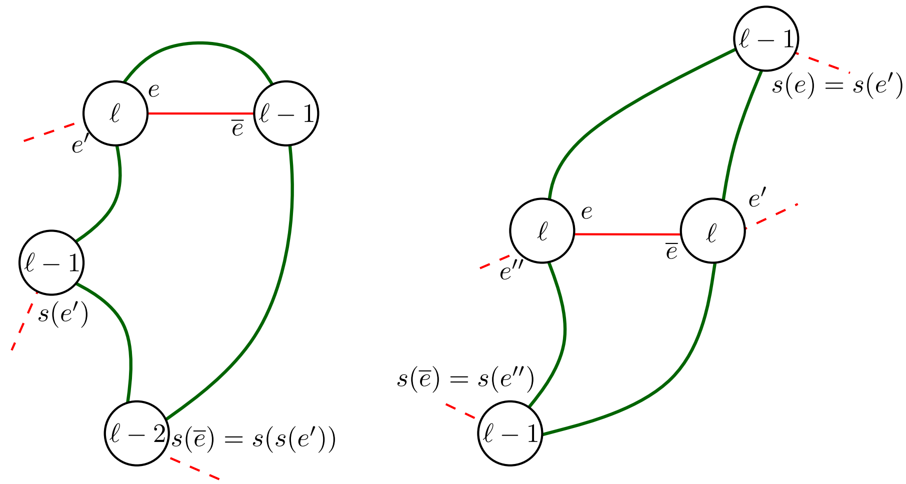

Figure 2.4: Illustration of the proof of Lemma 2.3.3. In this figure, $\ell = \ell(e)$

We deduce that the arcs joining the corners e to $s(e)$, resp. $\overline{e}$ to $s(\overline{e})$, resp. e' to $s(e')$, resp. $s(e')$ to $s(s(e')) = s(\overline{e})$, form a quadrangle, that contains the edge $\{e, \overline{e}\}$, and no other edge of $\mathbf{t}$.

If $\ell(e^+) = \ell(e^-) + 1$, the situation is the same by interchanging the roles of e and $\overline{e}$.

The only case that remains is when $\ell(e^+) = \ell(e^-)$. In this case, if e' and e'' are the corners following e and $\overline{e}$ respectively in the contour exploration of $\mathbf{t}$, then $\ell(e) = \ell(e') = \ell(\overline{e}) = \ell(e'')$, so that $s(e) = s(e')$ on the one hand and $s(\overline{e}) = s(e'')$ on the other hand. We deduce that the edge $\{e, \overline{e}\}$ is the diagonal of a quadrangle formed by the arcs linking e to $s(e)$, e' to $s(e') = s(e)$, $\overline{e}$ to $s(\overline{e})$ and e'' to $s(e'') = s(\overline{e})$. The different cases are summed up in Fig.2.4.

Now, notice that $\mathbf{q}$ has $2n$ edges (one per corner of $\mathbf{t}$) and $n+2$ vertices, so it must have n faces by Euler's formula. So all the faces must be of the form described above. This completes the proof. $\square$

Note that the quadrangulation $\mathbf{q}$ has a distinguished vertex v_*, but for now it is not a rooted quadrangulation. To fix this root, we will need an extra parameter $\epsilon \in \{-1, 1\}$. If $\epsilon = 1$ we let the root edge of $\mathbf{q}$ be the arc linking e_0 with $s(e_0)$, and oriented from $s(e_0)$ to e_0. If $\epsilon = -1$, the root edge is this same arc, but oriented from e_0 to $s(e_0)$.

In this way, we have defined a mapping Φ, from $\mathbb{T}_n \times \{-1, 1\}$ to the set

$\mathcal{Q}_n^\bullet$ of pairs $(\mathbf{q}, v_*)$, where $\mathbf{q} \in \mathcal{Q}_n$ and $v_* \in V(\mathbf{q})$. We call such pairs *pointed quadrangulations*.

Theorem 2.3.4. *For every $n \geq 1$, the mapping Φ is a bijection from $\mathbb{T}_n \times \{-1, 1\}$ onto $\mathcal{Q}_n^\bullet$.*

We omit the proof of this result. See Chassaing and Schaeffer [31, Theorem 4].

Corollary 2.3.5. *We have the following formula for every $n \geq 1$:*

$$\#\mathcal{M}_n = \#\mathcal{Q}_n = \frac{2}{n+2} 3^n \mathrm{Cat}_n \,,$$

where

$$\mathrm{Cat}_n = \frac{1}{n+1}\binom{2n}{n}$$

is the n-th Catalan number.

Proof. We first notice that $\#\mathcal{Q}_n^\bullet = (n+2)\#\mathcal{Q}_n$, since every quadrangulation $\mathbf{q} \in \mathcal{Q}_n$ has $n+2$ vertices, each of which induces a distinct element of $\mathcal{Q}_n^\bullet$. On the other hand, it is obvious that

$$\#\mathbb{T}_n \times \{-1, 1\} = 2 \cdot 3^n \#\mathbf{T}_n = 2 \cdot 3^n \mathrm{Cat}_n \,.$$

The result follows from Theorem 2.3.4. □

The probabilistic counterpart of this can be stated as follows.

Corollary 2.3.6. *Let Q_n be a uniform random element in $\mathcal{Q}_n$, and conditionally on Q_n, let v_* be chosen uniformly at random in $V(Q_n)$. On the other hand, let (T_n, ℓ_n) be chosen uniformly at random in $\mathbb{T}_n$, and let ϵ be independent of (T_n, ℓ_n) and uniformly distributed in $\{-1, 1\}$. Then $\Phi(T_n, \ell_n, \epsilon)$ has the same distribution as (Q_n, v_*).*

The proof is obvious, since the probability that (Q_n, v_*) equals some particular $(\mathbf{q}, v) \in \mathcal{Q}_n^\bullet$ equals $((n+2)\#\mathcal{Q}_n)^{-1} = (\#\mathcal{Q}_n^\bullet)^{-1}$.

Interpretation of the labels

The CVS bijection will be of crucial importance to us when we will deal with metric properties of random elements of $\mathcal{Q}_n$, because the labels on $\mathbf{q}$ that are inherited from a labeled tree through the CVS construction turn out to measure certain distances in $\mathbf{q}$. Recall that the set $V(\mathbf{t})$ is identified with $V(\mathbf{q})\setminus\{v_*\}$ if $(\mathbf{t},\ell)$ and $\mathbf{q}$ are associated through the CVS bijection (the choice of ϵ is irrelevant here). Hence, the function ℓ is also a function on $V(\mathbf{q})\setminus\{v_*\}$, and we extend it by letting, as previously, $\ell(v_*) = \min\{\ell(u) : u \in V(\mathbf{t})\} - 1$. For simplicity, we write

$$\min \ell = \min\{\ell(u) : u \in V(\mathbf{t})\} .$$

Proposition 2.3.7. *For every $v \in V(\mathbf{q})$, we have*

$$d_{\mathbf{q}}(v, v_*) = \ell(v) - \min \ell + 1 , \tag{2.2}$$

where $d_{\mathbf{q}}$ is the graph distance on $\mathbf{q}$.

Proof. Let $v \in V(\mathbf{q}) \setminus \{v_*\} = V(\mathbf{t})$, and let e be a corner (in $\mathbf{t}$) incident to v. Then the chain of arcs

$$e \to s(e) \to s^2(e) \to \ldots \to e_\infty$$

is a chain of length $\ell(e) - \ell(e_\infty) = \ell(v) - \ell(v_*)$ between v and v_*. Therefore, $d_{\mathbf{q}}(v, v_*) \leq \ell(v) - \ell(v_*)$. On the other hand, if $v = v_0, v_1, \ldots, v_d = v_*$ are the consecutive vertices of any chain linking v to v_*, then since $|\ell(e)-\ell(s(e))| = 1$ by definition for any corner e and since the edges of $\mathbf{q}$ all connect a corner to its successor, we get

$$d = \sum_{i=1}^{d} |\ell(v_i) - \ell(v_{i-1})| \geq |\ell(v_0) - \ell(v_d)| = \ell(v) - \ell(v_*) ,$$

as desired. $\square$

Remark. The preceding proof also shows that the chain of arcs $e \to s(e) \to s^2(e) \to \ldots \to e_\infty$ is a geodesic chain linking e^- to v_*. Such a geodesic chain, or more generally a chain of the form $e \to s(e) \to s^2(e) \to \ldots \to s^k(e)$, will be called a successor geodesic chain.

The triangle inequality for $d_{\mathbf{q}}$ (or the second part of the proof) gives the useful bound

$$d_{\mathbf{q}}(u, v) \geq |\ell(u) - \ell(v)| , \tag{2.3}$$

This bound will be improved in the next section.

As a consequence of the proposition, we obtain for instance that the "volume of spheres" around v_* can be interpreted in terms of ℓ: for every $k \geq 0$,

$$|\{v \in V(\mathbf{q}) : d_{\mathbf{q}}(v, v_*) = k\}| = |\{u \in V(\mathbf{t}) : \ell(u) - \min \ell + 1 = k\}| .$$

Two useful bounds

The general philosophy in the forthcoming study of random planar maps is then the following: information about labels in a random labeled tree, if this tree is uniformly distributed over $\mathbb{T}_n$, allows one to obtain information about distances in the associated quadrangulation. One major problem with this approach is that exact information will only be available for distances to a distinguished vertex v_*. There is no simple expression for the distances between two vertices distinct from v_* in terms of the labels in the tree. However, more advanced properties of the CVS bijection allow to get useful bounds on these distances. Recall that $e_0, e_1, e_2, \ldots$ is the contour sequence of corners (or oriented edges) around a tree $\mathbf{t} \in \mathbf{T}_n$, starting from the root (see the beginning of Section 2.3). We view $(e_i, i \geq 0)$ as cyclically ordered, and for any two corners e, e' of $\mathbf{t}$, we let $[e, e']$ be the set of all corners encountered when starting from e, following the cyclic contour order, and stopping when visiting e'.

Proposition 2.3.8. *Let $((\mathbf{t}, \ell), \epsilon)$ be an element in $\mathbb{T}_n \times \{-1, 1\}$, and $(\mathbf{q}, v_*) = \Phi(((\mathbf{t}, \ell), \epsilon))$. Let u, v be two vertices in $V(\mathbf{q}) \setminus \{v_*\}$, and let e, e' be two corners of $\mathbf{t}$ such that $e^- = u, (e')^- = v$.*
(i) *There holds that*

$$d_{\mathbf{q}}(u, v) \leq \ell(u) + \ell(v) - 2 \min_{e'' \in [e, e']} \ell(e'') + 2 \, ,$$

(ii) *There holds that*

$$d_{\mathbf{q}}(u, v) \geq \ell(u) + \ell(v) - 2 \min_{w \in [\![u, v]\!]} \ell(w) \, ,$$

where $[\![u, v]\!]$ is the set of all vertices lying on the geodesic path from u to v in the tree $\mathbf{t}$.

PROOF. For simplicity, let $m = \min_{e'' \in [e,e']} \ell(e'')$. Let e'' be the first corner in $[e,e']$ such that $\ell(e'') = m$. The corner $s^k(e)$, whenever it is well defined (i.e. whenever $d_{\mathbf{q}}(e^-, v_*) \geq k$), is called the k-th successor of e. Then e'' is the $(\ell(e) - m)$-th successor of e. Moreover, by definition, $s(e'')$ does not belong to $[e,e']$ since it has lesser label than e'', and necessarily, $s(e'')$ is also the $(\ell(e') - m + 1)$-st successor of e'. Hence, the successor geodesic chain $e \to s(e) \to s^2(e) \to \cdots \to s(e'')$ from $u = e^-$ to $s(e'')^-$, concatenated with the similar geodesic chain from v to $s(e'')^-$ is a path of length

$$\ell(u) + \ell(v) - 2m + 2\,,$$

and the distance $d_{\mathbf{q}}(u,v)$ is less than or equal to this quantity. This proves (i).

Let us prove (ii). Let $w \in [\![u,v]\!]$ be such that $\ell(w) = \min\{\ell(w') : w' \in [\![u,v]\!]\}$. If $w = u$ or $w = v$ then the statement follows trivially from (2.3). So we exclude this case. We can then write $\mathbf{t}$ as the union $\mathbf{t} = \mathbf{t}_1 \cup \mathbf{t}_2$ of two connected subgraphs of $\mathbf{t}$ such that $\mathbf{t}_1 \cap \mathbf{t}_2 = \{w\}$, $\mathbf{t}_1$ contains u but not v and $\mathbf{t}_2$ contains v but not u. There may be several such decompositions, so we just choose one. We consider a geodesic path γ from u to v in $\mathbf{q}$. If v_* belongs to this path, then this means that $d_{\mathbf{q}}(u,v) = d_{\mathbf{q}}(v_*,u) + d_{\mathbf{q}}(v_*,v)$ and the desired lower bound immediately follows from (2.2). So we may assume that v_* does not belong to γ. From our choice of $\mathbf{t}_1$ and $\mathbf{t}_2$, we can then find two corners $e_{(1)}$ and $e_{(2)}$ of $\mathbf{t}$ such that $e_{(1)}^-$ belongs to $\mathbf{t}_1$ and $e_{(2)}^-$ belongs to $\mathbf{t}_2$, $e_{(1)}^-$ and $e_{(2)}^-$ are consecutive points on γ, and the corners $e_{(1)}$ and $e_{(2)}$ are connected by an edge of $\mathbf{q}$. From the latter property, we must have $e_{(2)} = s(e_{(1)})$ or $e_{(1)} = s(e_{(2)})$. Consider only the first case for definiteness (the other one is treated in a similar fashion). Since the contour exploration of vertices of $\mathbf{t}$ must visit w between any visit of $u = e_{(1)}^-$ and any visit of $v = e_{(2)}^-$, the definition of the successor ensures that $\ell(w) \geq \ell(e_{(2)})$ (with equality only possible if $w = e_{(2)}^-$). Then, using (2.3) once again, we have

$$\begin{aligned} d_{\mathbf{q}}(u,v) &= d_{\mathbf{q}}(u, e_{(2)}^-) + d_{\mathbf{q}}(e_{(2)}^-, v) \\ &\geq \ell(u) - \ell(e_{(2)}^-) + \ell(v) - \ell(e_{(2)}^-) \\ &\geq \ell(u) + \ell(v) - 2\ell(w), \end{aligned}$$

giving the desired result. □

Notes for Chapter 2

The enumerative theory of maps is a very fundamental topic that has many ramifications in various branches of mathematics. We have not mentioned the connections that exist between maps and representation theory. See e.g. the Appendix in [50] for an introduction to this topic. There are also deep connections with integrable hierarchies and algebraic geometry, see e.g. Chapter 3.6 in [50] or [43].

The CVS bijection was discovered by [33] in 1981, with the notable difference that they considered maps that are not pointed. This has the effect of forcing the labels in the tree to be admissible, but also non-negative, and the enumeration of the set $\mathbb{T}_n^+$ of such labeled trees turns out to be quite complicated. The bijective approach for maps was then essentially left aside, but for some notable exceptions like [10], until it was revived in Schaeffer's PhD thesis [76], in which were found new families of "blossoming" trees for which the analogs of these positivity constraints found natural explanations in terms of classical cycle lemmas. However, these encodings do not allow in general to keep track of graph distances in the map. The idea to introduce a pointing in the map, which allows to lift the positivity assumption, appears in Theorem 4 of [31], and is of a great help both combinatorially and in the probabilistic study to come.

Chapter 3

Scaling limits of random trees

In this chapter, we let T_n be a uniform random element in the set $\mathbf{T_n}$ of rooted plane trees with n edges. We also let ℓ_n be a uniform admissible labeling of T_n, so that (T_n, ℓ_n) is a uniform random element of $\mathbb{T}_n$.

In order to emphasize the dependency on n, we let $e_0^n, e_1^n, \ldots$ be the contour exploration of the corners incident to T_n, starting with the root corner, and let $u_i^n = (e_i^n)^-$, so that u_0^n is the root vertex. This notation will be in force in subsequent chapters as well.

3.1 Convergence of the contour process

Let C_n be the contour process associated with the random tree T_n. Recall that C_n is a random variable in the set $\mathcal{E}_n$ of discrete excursions of length $2n$ defined in Section 1.3. Since the mapping that associates with every tree its contour process is a bijection, we see that C_n is a uniform random variable in $\mathcal{E}_n$.

Clearly, we have the following alternative description for C_n. Let $(S_k, k \geq 0)$ be a simple random walk in $\mathbb{Z}$, extended by linear interpolation to a random function $(S_t, t \geq 0)$. Then C_n has same distribution as $(S_t, 0 \leq t \leq 2n)$ conditioned on $A_n = \{S_k \geq 0\,, 0 \leq k \leq 2n\} \cap \{S_{2n} = 0\}$. To see this, simply note that A_n is the event that $(S_t, 0 \leq t \leq 2n)$ belongs to $\mathcal{E}_n$, and that the law of $(S_t, 0 \leq t \leq 2n)$ is uniform among possible random walk trajectories on $[0, 2n]$.

Recall that the celebrated Donsker invariance principle entails the following convergence in distribution in the space $\mathcal{C}([0,1], \mathbb{R})$ with the topology inherited from the uniform norm $\|\cdot\|_\infty$:

$$\left(\frac{S_{nt}}{\sqrt{n}}\right)_{0\le t\le 1} \xrightarrow[n\to\infty]{(d)} (B_t)_{0\le t\le 1}\,,$$

where the limiting process is standard Brownian motion. It is then natural, due to the discussion above, that a similar limiting result should hold for C_n, where the limiting process is "Brownian motion conditioned to remain non-negative on $[0,1]$ and to take the value 0 at time 1". The latter event has probability 0, so this does not make sense *stricto sensu*. However, it is quite natural that the process corresponding to this intuitive description is the *normalized Brownian excursion*. There are several possible explicit descriptions of this process. For instance, let B be a standard Brownian motion, and let

$$g = \sup\{t \le 1 : B_t = 0\}\,, \qquad d = \inf\{t \ge 1 : B_t = 0\}\,.$$

Since $B_1 \neq 0$ a.s., we have that $g < 1 < d$ with probability 1, and the portion of the path B on the interval $[g,d]$ is the excursion of B away from 0 that straddles 1. We renormalize this excursion by Brownian scaling by setting

$$\mathbb{e}_t = \frac{|B_{g+t(d-g)}|}{\sqrt{d-g}}\,, \qquad 0 \le t \le 1\,. \tag{3.1}$$

The process $(\mathbb{e}_t, 0 \le t \le 1)$ is the normalized Brownian excursion.

Theorem 3.1.1. *Let $C_{(n)}$ denote the renormalized contour process of T_n, defined by*

$$C_{(n)}(t) = \frac{C_n(2nt)}{\sqrt{2n}}\,, \qquad 0 \le t \le 1\,.$$

Then the following convergence in distribution holds in $\mathcal{C}([0,1],\mathbb{R})$:

$$C_{(n)} \xrightarrow[n\to\infty]{(d)} \mathbb{e}\,.$$

We will not provide a proof of this result, and rather refer the reader to [58] for instance. However, in order to familiarize ourselves with the limiting process, let us give two other characterizations.

More on the Brownian excursion

We need some notation. For $t > 0$ and $x, y \in \mathbb{R}$, let

$$p_t(x, y) = \frac{1}{\sqrt{2\pi t}} \exp\left(-\frac{|x-y|^2}{2t}\right) ,$$

so that $p_t(x, y)$ is the density of the transition kernel of Brownian motion. For $t > 0$ and $x > 0$, let also

$$q_t(x) = \partial_x p_t = \frac{x}{\sqrt{2\pi t^3}} \exp\left(-\frac{|x|^2}{2t}\right) ,$$

so that $t \mapsto q_t(x)$ is the density of the law of the first hitting time of x by standard Brownian motion (this is a well-known consequence of the reflection principle). Finally, for $t > 0$ and $x, y > 0$, let

$$p_t^+(x, y) = p_t(x, y) - p_t(x, -y) ,$$

so that $p_t^+(x, y)$ is the density of the law at time t of standard Brownian motion started from x and killed at the first hitting time of 0, in restriction to the event that this hitting time is bigger than t.

Proposition 3.1.2. *The marginal distributions of the normalized Brownian excursion are given by the following formula: for every $t_1 < t_2 < \ldots < t_k < 1$, and positive $x_1, x_2, \ldots, x_k$,*

$$\frac{P(\mathbb{e}_{t_1} \in \mathrm{d}x_1, \ldots, \mathbb{e}_{t_k} \in \mathrm{d}x_k)}{\mathrm{d}x_1 \ldots \mathrm{d}x_k} = 2\sqrt{2\pi} q_{t_1}(x_1) \left(\prod_{i=1}^{k-1} p^+_{t_{i+1}-t_i}(x_i, x_{i+1})\right) q_{1-t_k}(x_k) .$$

See for instance [58] for a proof. The reader can check, as a good exercise, that this formula actually defines probability densities! The intuitive picture is that the term $q_{1-t_k}(x_k)$ corresponds to the density of Brownian paths started from x_k at time t_k, and hitting level 0 for the first time at time 1, while dually, the quantity $q_{t_1}(x_1)$ is the density of "Brownian paths" starting from 0, remaining non-negative, and hitting x_1 at time t_1. The intermediate terms correspond to the probability density that a Brownian motion starting from x_1 at time t_1 visits x_i at time t_i for $1 \leq i \leq k$, while never hitting 0 in this time interval.

One final description of $\mathbb{e}$ is given by the so-called *Vervaat transform.* Namely, let $f : [0, 1] \to \mathbb{R}$ be a continuous function with $f(0) = f(1) = 0$.

For every $s \in [0,1]$, let $V_s(f)$ be the path f shifted cyclically at time s, defined by $V_s f(t) = f(s+t \mod 1) - f(s)$, where it should be understood that the representative of $t+s \mod 1$ that is chosen is the unique one in $[0,1)$. Finally, let $s_*(f) = \inf\{t \in [0,1] : f(t) = \inf f\}$, and set $Vf = V_{s_*(f)} f$.

Recall that the normalized Brownian bridge $\mathbb{b}$ is informally the Brownian motion conditioned to hit 0 at time 1. It can be defined by taking a standard Brownian motion B and letting $\mathbb{b}_t = B_t - tB_1, 0 \le t \le 1$. It is a continuous process given by the following marginal distributions: for every $0 < t_1 < t_2 < \ldots < t_k < 1$ and $x_1, \ldots, x_k \in \mathbb{R}$,

$$P(\mathbb{b}_{t_1} \in \mathrm{d}x_1, \ldots, \mathbb{b}_{t_k} \in \mathrm{d}x_k) = \prod_{i=0}^{k} p_{t_{i+1}-t_i}(x_i, x_{i+1}),$$

where by convention we let $t_0 = 0$, $t_{k+1} = 1$ and $x_0 = x_{k+1} = 0$.

Theorem 3.1.3 (Vervaat's theorem)**.** *The two processes* $\mathbb{e}$ *and* $V\mathbb{b}$ *have the same distribution.*

The original proof by Vervaat used an approximation result of both processes by their random walk analogues, and showed that $V\mathbb{b}$ has the same marginal distributions as $\mathbb{e}$. Conceptually simpler proofs can be obtained, still using approximation, with the help of the so-called cyclic lemma, see [74]. See also Biane [20].

Consequences on the geometry of random trees

Theorem 3.1.1 easily entails some results on the geometry of the random tree T_n as n gets large.

Proposition 3.1.4. *The following convergences in law hold:*

1. *Let* $R(T_n) = \max\{d_{T_n}(v, u_0^n) : v \in V(T_n)\}$ *be the maximal graph distance from the root vertex to a vertex in* T_n*.*
$$\frac{R(T_n)}{\sqrt{2n}} \xrightarrow[n\to\infty]{(d)} \sup \mathbb{e}\,.$$
2. *Let* u_* *be a uniformly distributed random vertex in* T_n*, then*
$$\frac{d_{T_n}(u_0^n, u_*)}{\sqrt{2n}} \xrightarrow[n\to\infty]{(d)} \mathbb{e}_U\,,$$
where U *is a uniform random variable in* $[0,1]$*, independent of* $\mathbb{e}$*.*

Proof. The first point is clear from the fact that $R(T_n) = \max C_n$ and Theorem 3.1.1. As for the second point, we first observe that we may slightly change the hypothesis on the distribution of v_*: it clearly suffices to prove the desired convergence when v_* is replaced by a vertex that is uniformly chosen among the n vertices of T_n that are distinct from the root vertex u_0^n of T_n.

Now, for $s \in [0, 2n)$, we let $\langle s \rangle = \lceil s \rceil$ if C_n has slope $+1$ immediately after s, and $\langle s \rangle = \lfloor s \rfloor$ otherwise. Then, if $u \in V(T_n)$, we have $u^n_{\langle s \rangle} = u$ if and only if $u \neq u_0^n$ and s is a time when the contour exploration around T_n explores one of the two oriented edges between u and its parent $\neg u$. Therefore, for every $u \in V(T_n) \setminus \{u_0^n\}$, the Lebesgue measure of $\{s \in [0, 2n) : u^n_{\langle s \rangle} = u\}$ equals 2. Consequently, if U is a uniform random variable in $[0, 1)$, independent of T_n, then $u^n_{\langle 2nU \rangle}$ is uniform in $V(T_n) \setminus \{u_0^n\}$. Hence, it suffices to prove the desired result with $u^n_{\langle 2nU \rangle}$ instead of u_*.

Since $|s - \langle s \rangle| \leq 1$, Theorem 3.1.1 entails that

$$\frac{d_{T_n}(u_0^n, u^n_{\langle 2nU \rangle})}{\sqrt{2n}} = C_{(n)}(\langle 2nU \rangle / 2n)$$

converges in distribution to $\mathbb{e}_U$, as wanted. □

The distributions of the random variables appearing in this theorem can be made explicit. The law of $\mathbb{e}_U$ is the so-called *Rayleigh distribution*, given by the formula

$$P(\mathbb{e}_U > x) = \exp(-x^2/2)\,, \qquad x \geq 0\,.$$

It is a good exercise to derive this from the marginal distributions of $\mathbb{e}$ given in Proposition 3.1.2. The law of $\sup \mathbb{e}$ is the more complicated Theta distribution

$$P(\sup \mathbb{e} > x) = \sum_{j \geq 1} e^{-2j^2x^2}(8j^2x^2 - 2)\,.$$

See [39, Chapter V.4.3] for a proof based on a discrete approximation by random trees.

3.2 The Brownian continuum random tree

In the same way that the contour process C_n encodes the random tree T_n, we can view the normalized Brownian excursion $\mathbb{e}$ as the contour process of a "continuum random tree". Roughly speaking, every time $t \in [0,1]$ will code for a point of this tree at height $\mathbb{e}_t$, where the height should be understood again as "distance to the root", the latter being the point visited at time 0, which naturally has height $\mathbb{e}_0 = 0$.

The latter property is not sufficient to get a tree structure, so let us go back to the discrete picture for a while. Recall that $e^n_0, e^n_1, \ldots, e^n_{2n-1}, e^n_{2n} = e^n_0$ denote the oriented edges in the contour exploration around T_n, and that $u^n_i = (e^n_i)^-$ is the corresponding i-th visited vertex. Let $i, j \in \{0, 1, \ldots, 2n\}$ with say $i \leq j$. The vertices u^n_i, u^n_j have a highest common ancestor in T_n, denoted by $u^n_i \wedge u^n_j$. It is not difficult to convince oneself that $u^n_i \wedge u^n_j = u^n_k$, where k is any integer between i and j such that $C_n(k) = \min\{C_n(r) : i \leq r \leq j\}$. Denoting the last quantity by $\check{C}_n(i,j)$, we see that this quantity is the height in T_n of $u^n_i \wedge u^n_j$. Therefore, the unique simple path from u^n_i to u^n_j in T_n, which goes from u^n_i down to $u^n_i \wedge u^n_j$, and then from $u^n_i \wedge u^n_j$ up to u^n_j, has length

$$d_{T_n}(u^n_i, u^n_j) = C_n(i) + C_n(j) - 2\check{C}_n(i,j)\,. \tag{3.2}$$

Note that the previous formula does not depend on the actual choices of i and j corresponding to the vertices u^n_i, u^n_j: if we choose i', j' such that $u^n_i = u^n_{i'}$ and $u^n_j = u^n_{j'}$ then the right-hand side will give the same result, as it should. The reason is that, between two visits of the same vertex v in contour order, the contour process only explores vertices that are descendants of v in the tree, and thus have larger heights than that of v.

We can use (3.2) to get better intuition of what it means for a general function to be the "contour process of a tree". Let $f : [0,1] \to \mathbb{R}_+$ be a non-negative, continuous function with $f(0) = f(1) = 0$. We call such a function an *excursion function*, and denote by $\mathcal{E}$ the set of excursion functions, that we endow with the uniform norm. For every $s, t \in [0,1]$, let

$$\check{f}(s,t) = \inf\{f(u) : s \wedge t \leq u \leq s \vee t\}\,,$$

and set

$$d_f(s,t) = f(s) + f(t) - 2\check{f}(s,t)\,.$$

Proposition 3.2.1. *The function d_f on $[0,1]^2$ is a pseudo-metric: it is non-negative, symmetric, and satisfies the triangle inequality.*

Proof. The only non-trivial aspect is the triangle inequality. In fact, the following stronger 4-*point condition* is true: for every $s, t, u, v \in [0,1]$,

$$d_f(s,t) + d_f(u,v) \leq \max(d_f(s,u) + d_f(t,v), d_f(s,v) + d_f(t,u))\,. \tag{3.3}$$

One recovers the triangle inequality by taking $u = v$ above. Let us prove this inequality, which after simplification amounts to

$$\check{f}(s,t) + \check{f}(u,v) \geq \min(\check{f}(s,u) + \check{f}(t,v), \check{f}(s,v) + \check{f}(t,u)) .$$

We prove this by case study depending on the relative positions of s, t, u, v.
Case 1: $s \leq t \leq u \leq v$**.** Clearly we have $\check{f}(s,u) \leq \check{f}(s,t)$ and $\check{f}(t,v) \leq \check{f}(u,v)$, giving the result.
Case 2: $s \leq u \leq t \leq v$**.** We discuss further subcases.

If $\check{f}(s,v) = \check{f}(u,t)$, then we see that this latter quantity is also equal to $\check{f}(s,t) = \check{f}(u,v)$, and the result follows.

If $\check{f}(s,v) = \check{f}(s,u)$ and $\check{f}(u,t) \leq \check{f}(t,v)$ then $\check{f}(s,t) = \check{f}(s,v)$ and $\check{f}(u,v) = \check{f}(u,t)$, giving the result. In the case where $\check{f}(s,v) = \check{f}(s,u)$ and $\check{f}(t,v) \leq \check{f}(u,t)$, we rather use $\check{f}(s,t) = \check{f}(s,u)$ and $\check{f}(u,v) = \check{f}(t,v)$.

The remaining configurations within **Case 2** are symmetric to those discussed above.
Case 3: $s \leq u \leq v \leq t$**.** Again there are subcases.

If $\check{f}(s,t) = \check{f}(u,v)$ then these quantities are also equal to $\check{f}(s,u) = \check{f}(t,v)$.

If $\check{f}(s,t) = \check{f}(s,u) \leq \check{f}(u,v) \leq \check{f}(v,t)$ then $\check{f}(s,t) = \check{f}(s,v)$ and $\check{f}(u,v) = \check{f}(u,t)$. Finally, if $\check{f}(s,t) = \check{f}(s,u) \leq \check{f}(v,t) \leq \check{f}(u,v)$ then the result is immediate.

Finally, all remaining configuration are symmetric to those discussed above. $\square$

Note that in (3.3), the following even stronger result is true: among the three quantities $d_f(s,t) + d_f(u,v)$, $d_f(s,u) + d_f(t,v)$ and $d_f(s,v) + d_f(t,u)$, two are equal and larger than or equal to the third. This is a general property of distances satisfying the 4-point condition, that we leave as an exercise to the reader.

Definition 3.2.1. *Let (X, d) be a metric space. We say that X is a geodesic metric space if for every $x, y \in X$, there exists an isometric embedding $\phi : [0, d(x,y)] \to X$ such that $\phi(0) = x$ and $\phi(d(x,y)) = y$. This isometric embedding is called a geodesic path, and its image a geodesic segment, between x and y.*

We say that (X, d) is an $\mathbb{R}$-tree if it is a geodesic metric space, and if there is no embedding (continuous injective mapping) of $\mathbb{S}^1$ into X. Otherwise said, the geodesic segments are the unique injective continuous paths between their endpoints.

Proposition 3.2.2. *Let (X,d) be a connected metric space. Then (X,d) is an $\mathbb{R}$-tree if and only if it satisfies the 4-point condition.*

See for instance [32, 28] for a proof and the relation to the broader concept of Gromov hyperbolic spaces.

Now let f be as above. Since d_f is a pseudo-metric on $[0,1]$, the set $\{d_f = 0\} = \{s,t \in [0,1] : d_f(s,t) = 0\}$ is an equivalence relation on $[0,1]$. We let $\mathcal{T}_f = [0,1]/\{d_f = 0\}$ be the quotient set and $p_f : [0,1] \to \mathcal{T}_f$ the canonical projection. Being a class function for the relation $\{d_f = 0\}$, the function d_f naturally induces a (true) distance function on the set $\mathcal{T}_f$, and we still denote this distance by d_f.

Proposition 3.2.3. *The space $(\mathcal{T}_f, d_f)$ is a compact $\mathbb{R}$-tree.*

Proof. Note that $d_f(s,t) \le 2\omega(f, |s-t|)$, where $\omega(f,\delta) = \sup\{|f(s)-f(t) : |t-s| \le \delta\}$ is the modulus of continuity of f. Since f is continuous, this converges to 0 as $\delta \to 0$, and this clearly implies that p_f is a continuous mapping from $[0,1]$ to $(\mathcal{T}_f, d_f)$. Therefore, $\mathcal{T}_f$ is compact and connected as a continuous image of a compact, connected set. The result follows from Proposition 3.2.2. $\square$

Before going any further, let us give another proof of connectedness, which is interesting in its own right since it will also allow to give explicitly the geodesics paths in the tree.

Proposition 3.2.4. *For every $t \in [0,1]$ and $0 \le r \le f(t)$, let*

$$\gamma_t^+(r) = \inf\{s \ge t : f(s) < f(t)-r\}\,, \quad \gamma_t^-(r) = \sup\{s \le t : f(s) < f(t)-r\}\,,$$

with the convention that $\inf\varnothing = 1$ and $\sup\varnothing = 0$. Then $p_f(\gamma_t^-(r)) = p_f(\gamma_t^+(r))$ for every $r \in [0, f(t)]$, and if we denote by $\Gamma_t(r)$ this common value, then Γ_t is the (unique) geodesic path from $p_f(t)$ to $p_f(0)$.

Proof. It is immediate to check that $d_f(\gamma_t^-(r), \gamma_t^+(r)) = 0$ for every r, and that $d_f(\gamma_t^+(r), \gamma_t^+(r')) = r' - r$ for every $0 \le r \le r' \le f(t)$, since $f(\gamma_t^+(r)) = f(t) - r$ and $f(\gamma_t^+(r')) = \check{f}(\gamma_t^+(r), \gamma_t^+(r')) = f(t) - r'$. $\square$

We let $[\![a,b]\!]$ be the unique geodesic segment from a to b in the tree $\mathcal{T}_f$. Note that $\mathcal{T}_f$ naturally comes with a root, which is the point $\rho_f = p_f(0) = p_f(1)$. In particular, it also carries a genealogical structure, namely, for every $a,b \in \mathcal{T}_f$, there is a unique point $a \wedge b$ (the highest common ancestor) such that $[\![\rho_f, a]\!] \cap [\![\rho_f, b]\!] = [\![\rho_f, a \wedge b]\!]$. The geodesic segment $[\![a,b]\!]$ is then the concatenation of $[\![a, a\wedge b]\!]$ with $[\![a\wedge b, b]\!]$. It is quite easy to find $a \wedge b$ in terms of the contour function f.

Proposition 3.2.5. *If $a = p_f(s)$ and $b = p_f(t)$, then $a \wedge b = p_f(u)$ for any $u \in [s \wedge t, s \vee t]$ such that $f(u) = \check{f}(s,t)$.*

Proof. It suffices to check that $\gamma_s^+(f(s) - r) = \gamma_t^+(f(t) - r)$ for every $r \leq \check{f}(s,t)$ and that $d_f(\gamma_s^+(f(s)-r), \gamma_t^+(f(t)-r)) > 0$ for every $r \in (\check{f}(s,t), f(s) \wedge f(t))$. We leave the details to the reader. □

We can finally define the Brownian continuum random tree by breathing some probabilistic life into the above abstract construction.

Definition 3.2.2. *The Brownian continuum random tree is the random $\mathbb{R}$-tree $(\mathcal{T}_{\mathbb{e}}, d_{\mathbb{e}})$ encoded by a normalized Brownian excursion as above. We see it as a "rooted" tree by distinguishing the "root" $\rho_{\mathbb{e}}$.*

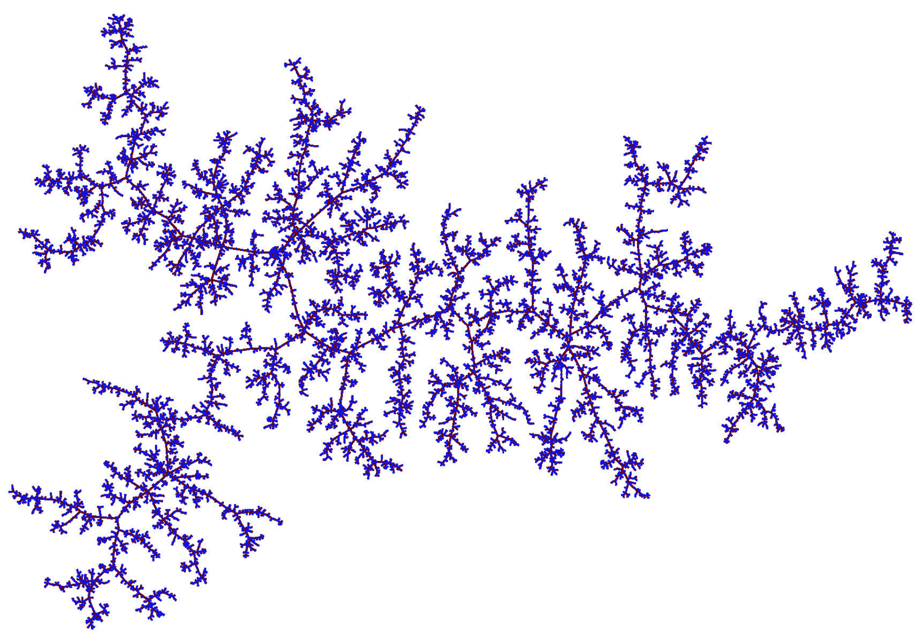

Figure 3.1: A simulation of the Brownian CRT, simulation by Igor Kortchemski

In order to view $\mathcal{T}_{\mathbb{e}}$ as a *bona fide* random variable, we first have to introduce a σ-field on the set of metric spaces in which it takes its values. As was pointed to us by Martin Hairer, a canonical way to achieve this is to endow the image of the mapping Tree : $f \mapsto (\mathcal{T}_f, d_f)$ with the final topology obtained by pushing forward the uniform topology on $\mathcal{E}$, and consider the associated Borel σ-algebra. More precisely, let us consider that Tree takes its values in the set of compact $\mathbb{R}$-trees, seen up to isometries (such identifications will become systematic afterwards). One can remark that Tree is surjective with this interpretation.

Exercise: Let (X, d) be a compact $\mathbb{R}$-tree. Show that there exists an excursion function $f \in \mathcal{E}$ such that (X, d) is isometric to $(\mathcal{T}_f, d_f)$.

Therefore, the final topology described above is a topology on the set of isometry classes of compact $\mathbb{R}$-tree. However, it is not quite clear what this topology looks like. The purpose of the next section is to introduce an important, more explicit topology (in our particular context, both topologies will turn out to be the same in the end).

3.3 The Gromov-Hausdorff topology

The Gromov-Hausdorff topology was introduced in the context of metric geometry as a way to allow smooth structures to approach metrics that, even though they are not smooth anymore, still present certain of the characters of the approximating metrics (in particular, present curvature bounds). See [45, 28] for an introduction to this topic.

Recall that if (Z, δ) is a metric space and $A, B \subset Z$, the Hausdorff distance between A and B is given by

$$\delta_H(A, B) = \max\{\delta(x, B) : x \in A\} \vee \max\{\delta(y, A) : y \in B\},$$

where by definition $\delta(x, C) = \inf\{\delta(x, y) : y \in C\}$ for $C \subset Z$. The function δ_Z defines a distance function on the set of non-empty, closed subsets of Z.

Let (X, d) and (X', d') be two compact metric spaces. The Gromov-Hausdorff distance between these spaces is defined by

$$\mathrm{d_{GH}}((X, d), (X', d')) = \inf \delta_H(\phi(X), \phi'(X')),$$

where the infimum is taken over all metric spaces (Z, δ) and all isometric embeddings ϕ, ϕ' from X, X' respectively into Z.

Clearly, if (X, d) and (X', d') are isometric metric spaces, then their Gromov-Hausdorff distance is 0, so $\mathrm{d_{GH}}$ defines at best a "pseudo distance" between metric spaces. This is actually a good thing, because the class of all compact metric spaces is too big to be a set in the set-theoretic sense, however, the family of compact metric spaces seen up to isometries is indeed a set, in the sense that there exists a set $\mathbb{M}$ such that any compact metric space is isometric to exactly one element of $\mathbb{M}$. Quite surprisingly perhaps, such a set does not require the axiom of choice to be constructed, as it can be realized as a set of doubly infinite non-negative matrices whose entries satisfy the triangle inequality.

Theorem 3.3.1. *The function* $\mathrm{d_{GH}}$ *induces a distance function on the set* $\mathbb{M}$ *of isometry classes of compact metric spaces. Furthermore, the space* $(\mathbb{M}, \mathrm{d_{GH}})$ *is separable and complete.*

The definition of $\mathrm{d_{GH}}$ through a huge infimum makes it quite daunting. However, there is a very useful alternative description via "couplings", which makes this distance quite close in spirit to other distances arising in metric geometry and analysis, such as the Wasserstein distances. If X and X' are two sets, a *correspondence* between X and X' is a subset $R \subset X \times X'$ such that for every $x \in X$, there exists $x' \in X'$ such that $(x, x') \in R$, and for every $x' \in X'$, there exists $x \in X$ such that $(x, x') \in R$. Otherwise said, the restrictions to R of the two projections from $X \times X'$ to X, X' are both surjective. We let $\mathrm{Cor}(X, X')$ be the set of all correspondences between X and X'.

If now (X, d) and (X', d') are metric spaces, and $R \in \mathrm{Cor}(X, X')$, the *distortion* of R with respect to d, d' is defined by

$$\mathrm{dis}\,(R) = \sup\{|d(x, y) - d'(x', y')| : (x, x'), (y, y') \in R\} .$$

Note that we do not mention d, d' in the notation for simplicity, although $\mathrm{dis}\,(R)$ clearly does not only depend on R.

Proposition 3.3.2. *Let* (X, d) *and* (X', d') *be compact metric spaces. Then*

$$\mathrm{d_{GH}}((X, d), (X', d')) = \frac{1}{2} \inf_{R \in \mathrm{Cor}(X, X')} \mathrm{dis}\,(R) .$$

This proposition is left as a good exercise to the reader, who is referred to the above references for help. Before moving on, let us introduce quickly the so-called pointed Gromov-Hausdorff distances. Fix an integer $k \geq 1$. A k-pointed metric space is a triple $(X, d, \mathbf{x})$ where $\mathbf{x} = (x_1, \ldots, x_k) \in$

X^k. We define the Gromov-Hausdorff distance between the k-pointed spaces $(X, d, \mathbf{x}), (X', d', \mathbf{x}')$ by the formula

$$\mathrm{d_{GH}}((X, d, \mathbf{x}), (X', d', \mathbf{x}')) = \frac{1}{2} \inf \left\{ \mathrm{dis}\,(R) : \begin{array}{c} R \in \mathrm{Cor}(X, X') \\ (x_i, x_i') \in R\,, 1 \leq i \leq k \end{array} \right\}.$$

Similarly to Theorem 3.3.1, for every fixed $k \geq 1$, $\mathrm{d_{GH}}$ induces a complete and separable distance on the set $\mathbb{M}^{(k)}$ of k-pointed compact metric spaces considered up to isometries preserving the distinguished points, so that $(X, d, \mathbf{x})$ and $(X', d', \mathbf{x}')$ are considered equivalent if there exists a global isometry $\phi : X \to X'$ such that $\phi(x_i) = x_i'$ for $1 \leq i \leq k$. A 1-pointed $\mathbb{R}$-tree will also be called *rooted*.

From this, we obtain an important consequence for the encoding of $\mathbb{R}$-trees by functions.

Proposition 3.3.3. *The mapping* $f \mapsto (\mathcal{T}_f, d_f, \rho_f)$ *from* $\{f \in \mathcal{C}([0,1], \mathbb{R}_+) : f(0) = f(1) = 0\}$ *to* $(\mathbb{M}^{(1)}, \mathrm{d_{GH}})$ *is* 2*-Lipschitz.*

Proof. Let f, g be continuous non-negative functions on $[0,1]$ with value 0 at times 0 and 1. Recall that p_f, p_g denote the canonical projections from $[0,1]$ to $\mathcal{T}_f, \mathcal{T}_g$. Let R be the image of $[0,1]$ by the mapping $(p_f, p_g) : [0,1] \to \mathcal{T}_f \times \mathcal{T}_g$. Clearly, R is a correspondence between $\mathcal{T}_f$ and $\mathcal{T}_g$ by the surjectivity of p_f and p_g, and the roots $\rho_f = p_f(0)$ and $\rho_g = p_g(0)$ are in correspondence via R. Moreover, its distortion is equal to

$$\mathrm{dis}\,(R) = \sup_{s,t \in [0,1]} \left| (f(s) + f(t) - 2\check{f}(s,t)) - (g(s) + g(t) - 2\check{g}(s,t)) \right|,$$

and it is easy to see that this is bounded from above by $4\|f - g\|_\infty$. This concludes the proof by Proposition 3.3.2. $\square$

Due to this result, we can indeed view $(\mathcal{T}_{\mathbb{e}}, d_{\mathbb{e}}, \rho_{\mathbb{e}})$ as a random variable, being the image of $\mathbb{e}$ by the continuous mapping $f \mapsto (\mathcal{T}_f, d_f, \rho_f)$. At this point, one might wonder whether the Borel σ-field associated with the Gromov-Hausdorff topology is the same as that obtained by pushing forward the uniform topology on $\mathcal{E}$ by Tree, as we discussed at the end of the previous section. This is indeed the case, as a consequence of the next exercise.

Exercise: Let $(\mathcal{T}_n, d_n), n \geq 1$ be a sequence of compact $\mathbb{R}$-trees converging in the Gromov-Hausdorff sense to a limiting $\mathbb{R}$-tree $(\mathcal{T}, d)$. Show that there exist excursion functions $f_n, n \geq 1$ converging uniformly to a limit f, such that $(\mathcal{T}_{f_n}, d_{f_n})$ is isometric to $(\mathcal{T}_n, d_n)$ for every $n \geq 1$, and $(\mathcal{T}_f, d_f)$ is isometric to $(\mathcal{T}, d)$.

We finish this section with a statement for convergence of random trees to the Brownian tree that does not refer anymore to encoding by contour functions. As usual, T_n is a uniform random rooted plane tree with n edges.

Theorem 3.3.4. *We have the following convergence in distribution in the space* $(\mathbb{M}^{(1)}, \mathrm{d}_{\mathrm{GH}})$*:*

$$\left(V(T_n), \frac{d_{T_n}}{\sqrt{2n}}, u_0^n\right) \xrightarrow[n\to\infty]{(d)} (\mathcal{T}_{\mathbb{e}}, d_{\mathbb{e}}, \rho_{\mathbb{e}}) \,.$$

PROOF. By using the Skorokhod representation theorem, let us assume that we are working on a probability space under which the convergence of Theorem 3.1.1 holds in the almost sure sense rather than in distribution. As in many cases where this theorem is utilized, this is not absolutely necessary, but this will ease considerably the arguments. We build a correspondence between $V(T_n)$ and $\mathcal{T}_{\mathbb{e}}$ by letting R_n be the image of the set $\{(i, t : 0 \leq i \leq 2n, 0 \leq t \leq 1, i = \lfloor 2nt \rfloor)\}$ by the mapping $(i,t) \mapsto (u_i^n, p_{\mathbb{e}}(t))$ from $\{0, 1, \ldots, 2n\} \times [0,1]$ to $V(T_n) \times \mathcal{T}_{\mathbb{e}}$. This correspondence associates the roots u_0^n and $\rho_{\mathbb{e}}$ of T_n and $\mathcal{T}_{\mathbb{e}}$. With the help of (3.2), the distortion of R_n with respect to the metrics $d_{T_n}/\sqrt{2n}$ and $d_{\mathbb{e}}$ is then seen to be equal to

$$\sup_{s,t\in[0,1]} \left| d_{C_{(n)}}(\lfloor 2ns \rfloor, \lfloor 2nt \rfloor) - d_{\mathbb{e}}(s,t) \right| \,,$$

which converges to 0 as $n \to \infty$ by the uniform convergence of $C_{(n)}$ to $\mathbb{e}$. $\square$

3.4 Labeled trees and the Brownian snake

We now extend the preceding convergence results for labeled trees. Let (T_n, ℓ_n) be a uniform random element in $\mathbb{T}_n$. Clearly, the tree T_n is then uniform in $\mathbf{T}_n$, and conditionally on $T_n = \mathbf{t}$, the label function ℓ_n is uniform among the 3^n admissible labelings. Yet another way to view ℓ_n is the following. For every $u \in V(\mathbf{t})$ distinct from the root vertex u_0^n, set $Y_u = \ell_n(u) - \ell_n(\neg u)$. Then it is simple to see that the random variables $(Y_u, u \in V(\mathbf{t}) \setminus \{u_0^n\})$ are uniform in $\{-1, 0, 1\}$, and independent. Thus, if we denote by $u \prec v$ the fact that u is an ancestor of v in the tree T_n, then since $\ell_n(e_0^n) = 0$ by convention, we obtain

$$\ell_n(u) = \sum_{v \prec u, v \neq u_0^n} Y_u \,.$$

In other words, the label function ℓ_n can be seen as a random walk along the ancestral lines of T_n, with uniform step distribution in $\{-1, 0, 1\}$, the latter being centered with variance $2/3$. Since a typical branch of T_n, say the one going from u_0^n to $u_{\lfloor 2nt \rfloor}^n$, has a length asymptotic to $\sqrt{2n\mathbb{e}_t}$, according to Theorem 3.1.1, we understand by the central limit theorem that conditionally on $\mathbb{e}$,

$$\frac{\ell_n(u_{\lfloor 2nt \rfloor}^n)}{\sqrt{2/3} \times \sqrt{\sqrt{2n}}} = \left(\frac{9}{8n}\right)^{1/4} \ell_n(u_{\lfloor 2nt \rfloor}^n)$$

should converge to a centered Gaussian random variable with variance $\mathbb{e}_t$.

For this reason, we see a Gaussian field appearing in the labels ℓ_n. To understand its covariance structure, note that given T_n, for $u, v \in T_n$, if $u \wedge v$ denotes their highest common ancestor, then $\ell_n(u) - \ell_n(u \wedge v)$ and $\ell_n(v) - \ell_n(u \wedge v)$ are independent and independent of $\ell_n(u \wedge v)$. Indeed, the latter is the sum of the Y_w for w along the branch from u_0^n to $u \wedge v$, which is the common part of the paths from u_0^n to u and v, while the former involves sums of the variables Y_w with w in two disjoint sets of vertices (the paths from $u \wedge v$ to u and v). For this reason, we see that the covariance of $\ell_n(u)$ and $\ell_n(v)$ is $2/3$ times the height of $u \wedge v$. If $u = u_{\lfloor 2ns \rfloor}^n$ and $v = u_{\lfloor 2nt \rfloor}^n$, then this height is none other than $\check{C}_n(\lfloor 2ns \rfloor, \lfloor 2nt \rfloor)$. Again, this indicates that (conditionally on $\mathbb{e}$), for every $s, t \in [0, 1]$

$$\left(\frac{9}{8n}\right)^{1/4} (\ell_n(u_{\lfloor 2ns \rfloor}^n), \ell_n(u_{\lfloor 2nt \rfloor}^n)) \xrightarrow[n\to\infty]{(d)} (N, N') \,,$$

where (N, N') is a Gaussian vector with variance-covariance matrix

$$\begin{pmatrix} \mathbb{e}_s & \check{\mathbb{e}}(s,t) \\ \check{\mathbb{e}}(s,t) & \mathbb{e}_t \end{pmatrix} .$$

This discussion is the motivation for the following theorem. For $i \in \{0, 1, \ldots, 2n\}$, let $L_n(i) = \ell_n(u_i^n)$ be the label of the i-th visited vertex in contour order exploration around T_n, and let $L_n(t), 0 \le t \le 2n$ be the linear interpolation of

L_n between integer times. Let also

$$L_{(n)}(t) = \left(\frac{9}{8n}\right)^{1/4} L_n(2nt)\,, \qquad 0 \le t \le 1\,.$$

Theorem 3.4.1. *It holds that*

$$(C_{(n)}, L_{(n)}) \xrightarrow[n\to\infty]{(d)} (\mathbb{e}, Z),$$

in distribution in $\mathcal{C}([0,1],\mathbb{R})^2$*, where the pair* $(\mathbb{e}, Z)$ *is described as follows. The process* $\mathbb{e}$ *is a standard Brownian excursion, and conditionally on* $\mathbb{e}$*,* $Z = (Z_t, 0 \le t \le 1)$ *is a continuous, centered Gaussian process with covariance*

$$\mathrm{Cov}\,(Z_s, Z_t) = \check{\mathbb{e}}_{s,t}\,, \qquad s,t \in [0,1]\,.$$

We will not give the rather technical proof of this result here. The argument sketched above can easily be made rigorous to entail convergence of finite marginal distributions, while tightness requires a good control on the modulus of continuity of $L_{(n)}$, e.g. using Kolmogorov's criterion.

The process Z is sometimes called the *head of the Brownian snake*, due to the fact that it can be described alternatively in terms of Le Gall's Brownian snake (really a conditioned version thereof), which is a more elaborate path-valued process. The fact that the second displayed formula of Theorem 3.4.1 does define a covariance function on $[0,1]$, or that a Gaussian process Z with this covariance admits a continuous version, are not obvious, and part of the statement. See [58] for a proof of the first statement, and for the second, note that for every $s,t \in [0,1]$, $Z_t - Z_s$ is, given $\mathbb{e}$, a centered Gaussian random variable with variance $d_{\mathbb{e}}(s,t)$. Therefore, for every $p > 0$, $\alpha > 0$ and $s,t \in [0,1]$,

$$E[|Z_s - Z_t|^p \,|\, \mathbb{e}] = C_p d_{\mathbb{e}}(s,t)^{p/2} \le 2^{p/2} C_p \|\mathbb{e}\|_\alpha^{p/2} |s-t|^{p\alpha/2}\,,$$

where we have denoted by

$$\|f\|_\alpha = \sup_{\substack{s,t\in[0,1]\\ s\neq t}} \frac{|f(t)-f(s)|}{|t-s|^\alpha}$$

the α-Hölder norm of f. It is well-known that Brownian motion is α-Hölder continuous, and hence has a finite α-Hölder norm, for every $\alpha \in (0,1/2)$ a.s., and the same is true of $\mathbb{e}$ from its definition via (3.1). We conclude from Kolmogorov's continuity theorem the following property.

Proposition 3.4.2. *Almost surely, the process Z has a version that is Hölder continuous with any exponent $\beta \in (0, 1/4)$.*

It is useful to understand what Z means in terms of the Brownian tree $\mathcal{T}_{\mathbb{e}}$. By analogy with ℓ_n, which describes a family of random walks indexed by the branches of T_n, it is natural that Z should describe "Brownian motion along the branches of the Brownian tree". To understand this point, let us show that Z is a.s. a class function for $\{d_{\mathbb{e}} = 0\}$.

Proposition 3.4.3. *Almost surely, for every $s, t \in [0, 1]$, $d_{\mathbb{e}}(s, t) = 0$ implies $Z_s = Z_t$. Consequently, the function Z passes to the quotient and defines a function on $\mathcal{T}_{\mathbb{e}}$.*

Proof. We work conditionally on $\mathbb{e}$. For every rational $q \geq 0$, consider the excursion intervals of $\mathbb{e}$ above level q, i.e. maximal intervals of the open set $\{\mathbb{e} > q\}$, and write them as a (possibly empty) countable union $\bigcup_{i \in I_q} (a_i^q, b_i^q)$. Clearly, one has $d_{\mathbb{e}}(a_i^q, b_i^q) = 0$ and therefore $Z_{a_i^q} = Z_{b_i^q}$ almost surely for every $q \in \mathbb{Q}_+, i \in I_q$. The result is obtained by an approximation argument: one should check that for any excursion interval (a, b) of $\mathbb{e}$ above a real number r, the numbers a, b can be approximated arbitrarily closely by numbers of the form a_i^q, b_i^q. This is clearly the case, using the fact that $\mathbb{e}$ is not constant on any non-trivial interval. From Proposition 3.4.2, we obtain that a.s., $Z_a = Z_b$ whenever (a, b) is an excursion interval of $\mathbb{e}$ above any level r. It is not difficult to conclude from there. Indeed, recall the well-known fact that Brownian motion has a.s. pairwise distinct local minima, a fact that is transfered to $\mathbb{e}$ by (3.1). From this, it is not difficult to see that the equivalence classes of $\{d_{\mathbb{e}} = 0\}$ that are not singletons are either of the form $\{a, b\}$, where (a, b) is an excursion interval of $\mathbb{e}$ above a certain level, or $\{a, b, c\}$, where both (a, b) and (b, c) are excursion intervals of $\mathbb{e}$ above a certain level (in the latter case, b is a time at which $\mathbb{e}$ achieves a local minimum). □

3.5 More properties

3.5.1 The tree $\mathcal{T}_Z$

We should note that any function $f : [0, 1] \to \mathbb{R}$ that is continuous and satisfies $f(0) = f(1) = 0$ encodes an $\mathbb{R}$-tree in a similar way to that discussed in this chapter. Indeed, simply perform the Vervaat transform on f to get a function Vf that is non-negative and continuous, and set $(\mathcal{T}_f, d_f, \rho_f) := (\mathcal{T}_{Vf}, d_{Vf}, \rho_{Vf})$. This can be directly defined from a pseudo-metric d_f on $[0,1]$ that is defined by

$$d_f(s,t) = f(s) + f(t) - 2\max\left(\inf_{u\in[s\wedge t, s\vee t]} f(u), \inf_{u\in[s\vee t,1]\cup[0,s\wedge t]} f(u)\right) .$$

Note that the sets $[s\wedge t, s\vee t]$ and $[s\vee t,1]\cup[0,s\wedge t]$ on which the infima are taken can be seen as the two arcs from s to t, if we view f as defined on the circle $\mathbb{S}^1$ rather than the interval $[0,1]$, which is licit since $f(0) = f(1)$. The tree $\mathcal{T}_f$ is naturally rooted at the point $\rho_f = p_f(s_*)$, where s_* is any point in $[0,1]$ at which f achieves its overall minimum: $f(s_*) = \inf f$.

In particular, we can also define a tree $(\mathcal{T}_Z, d_Z, \rho_Z)$ from the random function Z. As we will see later, this is by no means artificial, and has an important role to play in the scaling limit of random maps.

3.5.2 Mass measure and re-rooting invariance

We start by noticing that the root $\rho_{\mathbb{e}}$ of the Brownian tree plays no particular role: had we chosen to distinguish any other “fixed” point, the resulting object would have had the same law.

Proposition 3.5.1. *For every $t \in [0,1]$, the $\mathbb{M}^{(1)}$-valued random variables $(\mathcal{T}_{\mathbb{e}}, d_{\mathbb{e}}, \rho_{\mathbb{e}})$ and $(\mathcal{T}_{\mathbb{e}}, d_{\mathbb{e}}, p_{\mathbb{e}}(t))$ have the same distribution.*

Proof. The mapping ϕ_k which, with every plane tree $\mathbf{t}$ with contour exploration $e_0, e_1, \ldots, e_{2n-1}$ associates the same tree $\mathbf{t}$ re-rooted at e_k, is clearly a bijection from $\mathbf{T}_n$ to itself. Therefore, $(V(T_n), d_{T_n}, u_0^n)$ has same distribution as $(V(T_n), d_{T_n}, u_k^n)$ for every k. Also, the very same argument as for Theorem 3.3.4, using the same correspondence R_n, entails that $(V(T_n), d_{T_n}/\sqrt{2n}, u^n_{\lfloor 2nt\rfloor})$ converges in distribution to $(\mathcal{T}_{\mathbb{e}}, d_{\mathbb{e}}, p_{\mathbb{e}}(t))$. This allows to conclude. $\square$

We now reinterpret this statement, together with the exchangeability properties of the Brownian tree, to argue that the root is “chosen uniformly at random” in the tree. More precisely, in the setting of Section 3.2, the $\mathbb{R}$-trees encoded by “contour functions” $f : [0,1] \to \mathbb{R}$ come with an extra natural object, which is the probability measure given by the image of Lebesgue’s measure on $[0,1]$ by the canonical projection $p_f : [0,1] \to \mathcal{T}_f$. Let us denote by λ_f this measure.

One of the uses of λ_f is that it allows to generate random variables in $\mathcal{T}_f$. If U is a uniform random point in $[0,1]$, then we can indeed see $p_f(U)$

as a λ_f-distributed random variable in $\mathcal{T}_f$. We obtain the following result by randomizing t in Proposition 3.5.1.

Corollary 3.5.2. *Let X be a $\lambda_{\mathbb{e}}$-distributed random variable in the Brownian CRT $\mathcal{T}_{\mathbb{e}}$. Then the two pointed spaces $(\mathcal{T}_{\mathbb{e}}, d_{\mathbb{e}}, \rho_{\mathbb{e}})$ and $(\mathcal{T}_{\mathbb{e}}, d_{\mathbb{e}}, X)$ have the same distribution (as random variables in $\mathbb{M}^{(1)}$).*

3.5.3 More on the topology of the Brownian CRT

Let (X, d) be an $\mathbb{R}$-tree. The degree $\deg(x)$ of $x \in X$ is the (possibly infinite) number of connected components of $X \setminus \{x\}$. We let $\mathrm{Lf}(X)$ be the set of points $x \in X$ such that $X \setminus \{x\}$ is connected, i.e. such that $\deg(x) = 1$. Such a point is called a *leaf* of X. The complement of leaves, denoted by $\mathrm{Sk}(X)$, is the *skeleton* of X.

The Brownian CRT has the somewhat surprising property that "most its points are leaves" in the following sense.

Proposition 3.5.3. *For every $t \in [0, 1]$, the point $p_{\mathbb{e}}(t) \in \mathcal{T}_{\mathbb{e}}$ is almost surely a leaf. In particular, the set of leaves of $\mathcal{T}_{\mathbb{e}}$ is uncountable, and $\lambda_{\mathbb{e}}(\mathrm{Lf}(\mathcal{T}_{\mathbb{e}})) = 1$.*

Moreover, a.s. the skeleton points of $\mathcal{T}_{\mathbb{e}}$ have degree in $\{2, 3\}$. These are points of the form $p_{\mathbb{e}}(s)$ where s is a time of right local minimum of $\mathbb{e}$, i.e. such that there exists $\varepsilon > 0$ with $\check{\mathbb{e}}(s, s + \varepsilon) = \mathbb{e}_s$. Equivalently, these are also points $p_{\mathbb{e}}(s)$ where s is a time of left local minimum, with the obvious definition.

Finally, the set of points with degree exactly 3 is the set of points $p_{\mathbb{e}}(u)$, where u ranges over the countable set of times where $\mathbb{e}$ attains a local minimum.

We leave to the reader the proof of this proposition, which is a good exercise and a good way to get acquainted with random real trees. In particular, note that leaves of $\mathcal{T}_{\mathbb{e}}$ are not all of the form $p_{\mathbb{e}}(u)$ where u is a time at which u attains a local maximum (this confusion is a common mistake). Such leaves are indeed very special: there are only countably many of them.

Chapter 4

First scaling limit results for random quadrangulations

In this chapter, we combine the CVS bijection of Section 2.3 with the scaling limit results for labeled trees obtained in the preceding chapter, to derive our first limiting results for distances in random quadrangulations.

4.1 Radius and profile

The combination of the CVS bijection, together with our scaling limit results for random trees discussed in the previous chapter, gives some interesting results for random quadrangulations without much extra effort.

For $\mathbf{m}$ a map and $u \in V(\mathbf{m})$, let $R(\mathbf{m}, u)$, the radius of $\mathbf{m}$ centered at u, be the quantity

$$R(\mathbf{m}, u) = \max_{v \in V(\mathbf{m})} d_{\mathbf{m}}(u, v).$$

Theorem 4.1.1. *Let Q_n be uniformly distributed in $\mathcal{Q}_n$, and conditionally on Q_n, let v_* be uniform in $V(Q_n)$. Then it holds that*

$$\left(\frac{9}{8n}\right)^{1/4} R(Q_n, v_*) \longrightarrow \sup Z - \inf Z,$$

where Z is the head of the Brownian snake.

Let us make a preliminary remark, since a similar trick will be constantly used later on. Let (T_n, ℓ_n) be uniformly distributed in $\mathbb{T}_n$ as in the previous chapter. Then we can assume without loss of generality that (Q_n, v_*) is the image of (T_n, ℓ_n) by the CVS bijection (here we do not mention explicitly the rooting convention for Q_n, say that the root orientation is chosen uniformly at random among the two possibilities compatible with the CVS bijection). Indeed, the fact that Q_n has $n+2$ vertices a.s. implies that the marginal law of Q_n is unbiased, i.e. uniform over $\mathcal{Q}_n$, and that v_* is uniform in $V(Q_n)$ given Q_n, as in the statement of the theorem.

PROOF. This is an immediate consequence of Theorem 3.4.1, once one notices that

$$R(Q_n, v_*) = \max_{v\in V(Q_n)} \ell(v) - \min_{v\in V(Q_n)} \ell(v) + 1 = \sup L_n - \inf L_n + 1,$$

due to (2.3.7). $\square$

Exercise. Prove that Theorem 4.1.1 remains true if one replaces v_* by the root vertex e_*^-.

The next result deals with the so-called *profile of distances* from the distinguished point. For $(\mathbf{m}, u)$ a pointed map and $r \geq 0$, let

$$I_{\mathbf{m},u}(r) = \#\{v \in V(\mathbf{m}) : d_{\mathbf{m}}(u, v) = r\}.$$

The sequence $(I_{\mathbf{m},u}(r), r \geq 0)$ records the sizes of the "spheres" centered at u in the map $\mathbf{m}$. The profile can be seen as a measure on $\mathbb{Z}_+$ with total volume $n+2$. Our first limit theorem is the following.

Theorem 4.1.2. *Let Q_n be uniformly distributed over $\mathbf{Q}_n$, and conditionally on Q_n, let v_*, v_{**} be chosen uniformly among the $n+2$ vertices of Q_n, and independently of each other. (i) It holds that*

$$\left(\frac{9}{8n}\right)^{1/4} d_{Q_n}(v_*, v_{**}) \xrightarrow[n\to\infty]{(d)} \sup Z\,.$$

(ii) The following convergence in distribution holds for the weak topology on probability measures on $\mathbb{R}_+$:

$$\frac{I_{Q_n,v_*}((8n/9)^{1/4}\cdot)}{n+2} \xrightarrow[n\to\infty]{(d)} \mathcal{I}\,,$$

where $\mathcal{I}$ is the occupation measure of Z above its infimum, defined as follows: for every non-negative, measurable $g : \mathbb{R}_+ \to \mathbb{R}_+$,

$$\langle \mathcal{I}, g\rangle = \int_0^1 \mathrm{d}t\, g(Z_t - \inf Z)\,.$$

The point (ii) is due to Chassaing and Schaeffer [31], and (i) is due to Le Gall [52], although these references state these properties in a slightly different context, namely, in the case where the role of v_* is played by the root vertex e_*^-. This indicates that as $n \to \infty$, the root vertex plays no particular role. We leave to the reader the proof of this alternative version of the above result as an interesting exercise. Some information about the limiting distributions arising in Theorem 4.1.1 and in (i) of Theorem 4.1.2 can be found in Delmas [37].

Property (i) identifies the so-called 2-point function of the Brownian map. An important generalization of this result has been obtained by Bouttier and Guitter [26], who were able to compute the 3-point function, namely the joint asymptotic distribution of the mutual distances between three vertices chosen uniformly at random in $V(Q_n)$. We will discuss this further in the notes at the end of Chapter 6.

Proof. For (i), we argue similarly as in the proof of Proposition 3.1.4. Rather than choosing v_{**} uniformly in $V(Q_n)$, we choose it uniformly in the set of vertices of T_n that are distinct from the root vertex of T_n (recall that $V(T_n) = V(Q_n) \setminus \{v_*\}$). This will not change the result since $n \to \infty$. Now recall the notation $\langle s \rangle$ from the proof of Proposition 3.1.4, and that if U is a random variable in $[0,1]$ independent of (T_n, ℓ_n), then the vertex $u^n_{\langle 2nU \rangle}$ of T_n visited at time $\langle 2nU \rangle$ in contour order is uniform in the set of vertices of T_n distinct from the root. Since $|s - \langle s \rangle| \leq 1$, Theorem 3.4.1 entails that

$$\begin{aligned}\left(\frac{8n}{9}\right)^{-1/4} d_{Q_n}(v_*, u^n_{\langle 2nU \rangle}) &= \left(\frac{8n}{9}\right)^{-1/4} (\ell_n(u^n_{\langle 2nU \rangle}) - \min \ell_n + 1) \\ &= \left(\frac{8n}{9}\right)^{-1/4} (L_n(\langle 2nU \rangle) - \min L_n + 1)\,,\end{aligned}$$

converges in distribution to $Z_U - \inf Z$ (here U is also assumed to be independent of $(\mathbb{e}, Z)$). The fact that $Z_U - \inf Z$ has the same distribution as $\sup Z$, or equivalently as $-\inf Z$, can be derived from the invariance of the CRT under uniform re-rooting, see e.g. [60]. This completes the proof of (i).

Finally, for (ii) we just note that, for every bounded continuous $g : \mathbb{R}_+ \to$

$\mathbb{R}$,

$$\begin{aligned}
&\frac{1}{n+2}\sum_{k\in\mathbb{Z}_+} I_{Q_n,v_*}(k)\, g((8n/9)^{-1/4}k) \\
&= \frac{1}{n+2}\sum_{v\in Q_n} g((8n/9)^{-1/4}d_{Q_n}(v_*,v)) \\
&= E_{**}[g((8n/9)^{-1/4}d_{Q_n}(v_*,v_{**}))] \\
&\xrightarrow[n\to\infty]{(d)} E_U[g(Z_U-\inf Z)] \\
&= \int_0^1 \mathrm{d}t\, g(Z_t-\inf Z)\,,
\end{aligned}$$

where E_{**} and E_U means that we take the expectation only with respect to v_{**} and U in the corresponding expressions (these are conditional expectations given (Q_n,v_*) and $(\mathbb{e},Z)$ respectively). In the penultimate step, we used the convergence established in the course of the proof of (i). □

4.2 Convergence as metric spaces

We would like to be able to understand the full scaling limit picture for random maps, in a similar way as what was done for trees, where we showed, using Theorem 3.3.4, that the distances in discrete trees, once rescaled by $\sqrt{2n}$, converge to the distances in the CRT $(\mathcal{T}_{\mathbb{e}},d_{\mathbb{e}})$. We thus ask if there is an analog of the CRT that arises as the limit of the properly rescaled metric spaces (Q_n,d_{Q_n}). In view of Theorem 4.1.2, the correct normalization for the distance should be $n^{1/4}$.

Assume that (T_n,ℓ_n) is uniformly distributed over $\mathbb{T}_n$, let ϵ be uniform in $\{-1,1\}$ and independent of (T_n,ℓ_n), and let Q_n be the random uniform quadrangulation with n faces and with a uniformly chosen vertex v_*, which is obtained from $((T_n,\ell_n),\epsilon)$ via the CVS bijection. We now follow Le Gall [54][1]. Recall our notation $u_0^n,u_1^n,\ldots,u_{2n}^n$ for the contour exploration of the vertices of T_n, and recall that in the CVS bijection these vertices are also

[1]At this point, it should be noted that [54, 59, 55] consider another version of Schaeffer's bijection, where no distinguished vertex v_* has to be considered. This results in considering pairs (T_n,ℓ_n) in which ℓ_n is conditioned to be positive. The scaling limits of such pairs are still tractable, and in fact, are simple functionals of $(\mathbb{e},Z)$, as shown in [60, 52]. So there will be some differences with our exposition, but these turn out to be unimportant.

viewed as elements of $V(Q_n) \setminus \{v_*\}$. Define a pseudo-metric on $\{0,\ldots,2n\}$ by letting $d_n(i,j) = d_{Q_n}(u_i^n, u_j^n)$. A major problem comes from the fact that $d_n(i,j)$ cannot be expressed as a simple functional of (C_n, L_n). The only distances that we are able to handle in an easy way are distances to v_*, through the following rewriting of (2.2):

$$d_{Q_n}(v_*, u_i^n) = L_n(i) - \min L_n + 1\,. \tag{4.1}$$

We also define, for $i,j \in \{0,1,\ldots,2n\}$,

$$d_n^0(i,j) = L_n(i) + L_n(j) - 2\max\Big(\min_{i\le k\le j} L_n(k), \min_{j\le k\le i} L_n(k)\Big) + 2\,.$$

Here, if $j < i$, the condition $i \le k \le j$ means that $k \in \{i, i+1, \ldots, 2n\} \cup \{0,1,\ldots,j\}$ and similarly for the condition $j \le k \le i$ if $i < j$.

As a consequence of Proposition 2.3.8(i), we have the bound $d_n \le d_n^0$.

We now extend the function d_n to $[0,2n]^2$ by letting

$$\begin{aligned} d_n(s,t) \;\; = \;\; & (\lceil s\rceil - s)(\lceil t\rceil - t) d_n(\lfloor s\rfloor, \lfloor t\rfloor) \\ & + (\lceil s\rceil - s)(t - \lfloor t\rfloor) d_n(\lfloor s\rfloor, \lceil t\rceil) \\ & + (s - \lfloor s\rfloor)(\lceil t\rceil - t) d_n(\lceil s\rceil, \lfloor t\rfloor) \\ & + (s - \lfloor s\rfloor)(t - \lfloor t\rfloor) d_n(\lceil s\rceil, \lceil t\rceil)\,, \end{aligned} \tag{4.2}$$

recalling that $\lfloor s\rfloor = \sup\{k \in \mathbb{Z}_+ : k \le s\}$ and $\lceil s\rceil = \lfloor s\rfloor + 1$. The function d_n^0 is extended to $[0,2n]^2$ by the obvious similar formula.

It is easy to check that d_n thus extended is continuous on $[0,2n]^2$ and satisfies the triangle inequality (although this is not the case for d_n^0), and that the bound $d_n \le d_n^0$ still holds. We define a rescaled version of these functions by letting

$$D_n(s,t) = \left(\frac{9}{8n}\right)^{1/4} d_n(2ns, 2nt)\,, \qquad 0 \le s,t \le 1\,.$$

We define similarly the functions D_n^0 on $[0,1]^2$. Then, as a consequence of Theorem 3.4.1, we have

$$(D_n^0(s,t), 0 \le s,t \le 1) \xrightarrow[n\to\infty]{(d)} (d_Z(s,t), 0 \le s,t \le 1)\,, \tag{4.3}$$

for the uniform topology on $C([0,1]^2, \mathbb{R})$, where by definition

$$d_Z(s,t) = Z_s + Z_t - 2\max\Big(\min_{s\le r\le t} Z_r, \min_{t\le r\le s} Z_r\Big)\,, \tag{4.4}$$

where if $t < s$ the condition $s \leq r \leq t$ means that $r \in [s,1] \cup [0,t]$.

We can now state

Proposition 4.2.1. *The family of laws of* $(D_n(s,t), 0 \leq s,t \leq 1)$*, as* n *varies, is relatively compact for the weak topology on probability measures on* $C([0,1]^2, \mathbb{R})$.

Proof. Let $s,t,s',t' \in [0,1]$. Then by a simple use of the triangle inequality, and the fact that $D_n \leq D_n^0$,

$$|D_n(s,t) - D_n(s',t')| \leq D_n(s,s') + D_n(t,t') \leq D_n^0(s,s') + D_n^0(t,t')\,,$$

which allows one to estimate the modulus of continuity at a fixed $\delta > 0$:

$$\sup_{\substack{|s-s'|\leq\delta \\ |t-t'|\leq\delta}} |D_n(s,t) - D_n(s',t')| \leq 2 \sup_{|s-s'|\leq\delta} D_n^0(s,s')\,. \tag{4.5}$$

However, the convergence in distribution (4.3) entails that for every $\varepsilon > 0$,

$$\limsup_{n\to\infty} P\left(\sup_{|s-s'|\leq\delta} D_n^0(s,s') \geq \varepsilon\right) \leq P\left(\sup_{|s-s'|\leq\delta} d_Z(s,s') \geq \varepsilon\right)\,,$$

and the latter quantity goes to 0 when $\delta \to 0$ (for any fixed value of $\epsilon > 0$) by the continuity of d_Z and the fact that $d_Z(s,s) = 0$. Hence, taking $\eta > 0$ and letting $\varepsilon = \varepsilon_k = 2^{-k}$, we can choose $\delta = \delta_k$ (tacitly depending also on η) such that

$$\sup_{n\geq 1} P\left(\sup_{|s-s'|\leq\delta_k} D_n^0(s,s') \geq 2^{-k}\right) \leq \eta 2^{-k}\,, \qquad k \geq 1,$$

entailing

$$P\left(\bigcap_{k\geq 1}\left\{\sup_{|s-s'|\leq\delta_k} D_n^0(s,s') \leq 2^{-k}\right\}\right) \geq 1 - \eta\,,$$

for all $n \geq 1$. Together with (4.5), this shows that with probability at least $1-\eta$, the function D_n belongs to the set of all functions f from $[0,1]^2$ into $\mathbb{R}$ such that $f(0,0) = 0$ and, for every $k \geq 1$,

$$\sup_{\substack{|s-s'|\leq\delta_k \\ |t-t'|\leq\delta_k}} |f(s,t) - f(s',t')| \leq 2^{-k}\,.$$

The latter set is compact by the Arzelà-Ascoli theorem. The conclusion then follows from Prokhorov's theorem. □

At this point, we are allowed to say that the random distance functions D_n admit a limit in distribution, up to taking $n \to \infty$ along a subsequence:

$$(D_n(s,t), 0 \le s,t \le 1) \xrightarrow{(d)} (D(s,t), 0 \le s,t \le 1) \tag{4.6}$$

for the uniform topology on $C([0,1]^2, \mathbb{R})$. In fact, we are going to need a little more than the convergence of D_n. From the relative compactness of the components, we see that the closure of the collection of laws of the triplets

$$((2n)^{-1}C_n(2n\cdot), (9/8n)^{1/4}L_n(2n\cdot), D_n), \quad n \ge 1$$

is compact in the space of all probability measures on $C([0,1],\mathbb{R})^2 \times C([0,1]^2,\mathbb{R})$. Therefore, it is possible to choose a subsequence $(n_k, k \ge 1)$ so that this triplet converges in distribution to a limit, which is denoted by $(\mathbb{e}, Z, D)$ (from Theorem 3.4.1, this is of course consistent with the preceding notation). The joint convergence to the triplet $(\mathbb{e}, Z, D)$ gives a coupling of D and d_Z such that $D \le d_Z$, since $D_n \le D_n^0$ for every n.

Define a random equivalence relation on $[0,1]$ by letting $s \approx t$ if $D(s,t) = 0$. We let $S = [0,1]/\approx$ be the associated quotient space, endowed with the quotient distance, which we still denote by D. The canonical projection $[0,1] \to S$ is denoted by $\mathbf{p}$.

Finally, let $s_* \in [0,1]$ be such that $Z_{s_*} = \inf Z$. It can be proved that s_* is unique a.s., see [63] or [60], and we will admit this fact (although it is not really needed for the next statement). We set $x_* = \mathbf{p}(s_*)$. We can now state the main result of this section.

Theorem 4.2.2. *The random pointed metric space (S, D, x_*) is the limit in distribution of the spaces $(V(Q_n), (9/8n)^{1/4}d_{Q_n}, v_*)$, for the Gromov-Hausdorff topology, along the subsequence $(n_k, k \ge 1)$. Moreover, we have a.s. for every $x \in S$ and $s \in [0,1]$ such that $\mathbf{p}(s) = x$,*

$$D(x_*, x) = D(s_*, s) = Z_s - \inf Z\,.$$

Note that, in the discrete model, a point at which the minimal label in T_n is attained lies at distance 1 from v_*. Therefore, the point x_* should be seen as the continuous analog of the distinguished vertex v_*. The last identity in the statement of the theorem is then of course the continuous analog of (2.2) and (4.1).

Proof. For the purposes of this proof, it is useful to assume, using the Skorokhod representation theorem, that the convergence

$$((2n)^{-1/2}C_n(2n\cdot), (9/8n)^{1/4}L_n(2n\cdot), D_n) \longrightmap (\mathbb{e}, Z, D)$$

holds a.s. along the subsequence (n_k). In what follows we restrict our attention to values of n in this sequence.

For every n, let $i_*^{(n)}$ be any index in $\{0, 1, \dots, 2n\}$ such that $L_n(i_*^{(n)}) = \min L_n$. Then for every $v \in V(Q_n)$, it holds that

$$|d_{Q_n}(v_*, v) - d_{Q_n}(u^n_{i_*^{(n)}}, v)| \leq 1$$

because $d_{Q_n}(v_*, u^n_{i_*^{(n)}}) = 1$ (v_* and $u^n_{i_*^{(n)}}$ are linked by an arc in the CVS bijection). Moreover, since $(8n/9)^{-1/4}L_n(2n\cdot)$ converges to Z uniformly on $[0, 1]$, and since we know[2] that Z attains its overall infimum at a unique point s_*, it is easy to obtain that $i_*^{(n)}/2n$ converges as $n \to \infty$ towards s_*.

For every integer n, we construct a correspondence R_n between $V(Q_n)$ and S, by putting:

- $(v_*, x_*) \in R_n$;
- $(u^n_{\lfloor 2ns \rfloor}, \mathbf{p}(s)) \in R_n$, for every $s \in [0, 1]$.

We then verify that the distortion of R_n (with respect to the metrics $(9/8n)^{1/4}d_{Q_n}$ on $V(Q_n)$ and D on S) converges to 0 a.s. as $n \to \infty$. We first observe that

$$\begin{aligned}
&\sup_{s\in[0,1]} |(9/8n)^{1/4}d_{Q_n}(v_*, u^n_{\lfloor 2ns \rfloor}) - D(x_*, \mathbf{p}(s))| \\
&\quad \leq (9/8n)^{1/4} + \sup_{s\in[0,1]} |(9/8n)^{1/4}d_{Q_n}(u^n_{i_*^{(n)}}, u^n_{\lfloor 2ns \rfloor}) - D(x_*, \mathbf{p}(s))| \\
&\quad = (9/8n)^{1/4} + \sup_{s\in[0,1]} |D_n(i_*^{(n)}/2n, \lfloor 2ns \rfloor/2n) - D(s_*, s)|,
\end{aligned}$$

[2]We could also perform the proof without using this fact, but it makes things a little easier.

which tends to 0 as $n \to \infty$, by the a.s. uniform convergence of D_n to D, and the fact that $i_*^{(n)}/2n$ converges to s_*. Similarly, we have

$$\begin{aligned}
&\sup_{s,t\in[0,1]} |(9/8n)^{1/4} d_{Q_n}(u^n_{\lfloor 2ns \rfloor}, u^n_{\lfloor 2nt \rfloor}) - D(\mathbf{p}(s), \mathbf{p}(t))| \\
&\quad = \sup_{s,t\in[0,1]} |D_n(\lfloor 2ns \rfloor/2n, \lfloor 2nt \rfloor/2n) - D(s,t)|
\end{aligned}$$

which tends to 0 as $n \to \infty$. We conclude that the distortion of R_n converges to 0 a.s. and that the pointed metric spaces $(V(Q_n), (9/8n)^{-1/4} d_{Q_n}, v_*)$ also converge a.s. to (S, D, x_*) in the Gromov-Hausdorff topology.

Let us prove the last statement of the theorem. Using once again the uniform convergence of D_n to D, we obtain that for every $s \in [0,1]$,

$$\begin{aligned}
D(s_*, s) &= \lim_{n\to\infty} D_n(i_*^{(n)}/2n, \lfloor 2ns \rfloor/2n) \\
&= \lim_{n\to\infty} \left(\frac{8n}{9}\right)^{-1/4} d_{Q_n}(v_*, u^n_{\lfloor 2ns \rfloor}) \\
&= \lim_{n\to\infty} \left(\frac{8n}{9}\right)^{-1/4} (L_n(\lfloor 2ns \rfloor) - \min L_n + 1) \\
&= Z_s - \inf Z\,,
\end{aligned}$$

as desired. □

4.3 The Brownian map

It is tempting to call (S, D) the "Brownian map", or the "Brownian continuum random map", by analogy with the fact that the "Brownian continuum random tree" is the scaling limit of uniformly distributed plane trees with n edges. However, the choice of the subsequence in Theorem 4.2.2 poses a problem of uniqueness of the limit. As we saw in the previous statement, only the distances to x_* are *a priori* defined as simple functionals of the process Z. Distances between other points in S are much harder to handle. Let us discuss these issues a little more.

Proposition 4.3.1. *Almost surely, the random function D is a pseudo-metric on $[0,1]$ that satisfies the following two properties*

1. *for every $s, t \in [0,1]$, if $d_{\mathbb{e}}(s,t) = 0$ then $D(s,t) = 0$*

2. for every $s,t \in [0,1]$, $D(s,t) \le d_Z(s,t)$.

Proof. This is obtained by a simple limiting argument. Again, let us assume, using Skorokhod's representation theorem, that $(\mathbb{e}, Z, D)$ is the almost sure limit of $(C_{(n)}, L_{(n)}, D_n)$. Let us take $s < t$ such that $d_{\mathbb{e}}(s,t) = 0$. Then we can find $i_n, j_n \in \{0, 1, \dots, 2n\}$ such that $i_n/2n \to s$, $j_n/2n \to t$, and $u^n_{i_n} = u^n_{j_n}$: this is a simple consequence of the almost sure convergence of $C_{(n)}$ to $\mathbb{e}$ and of the fact that C_n is the contour process of the tree T_n. Clearly, this implies that $D_n(i_n/2ns, i_n/2nt) = 0$, and we conclude since $D(s,t)$ is the limit of the latter quantity as $n \to \infty$.

The second bound is obtained by a similar but simpler limiting argument, using the fact that $D_n(s,t) \le D^0_n(s,t)$ and the convergence of $D^0_n(s,t)$ to $d_Z(s,t)$. □

From these two properties, one can obtain a refined upper bound for D. Let $s,t \in [0,1]$, and let $s_1, t_1, s_2, t_2, \dots, s_k, t_k$ be points in $[0,1]$ such that $s_1 = s, t_k = t$ and $d_{\mathbb{e}}(t_i, s_{i+1}) = 0$ for every $i \in \{1, \dots, k-1\}$. Then by the triangle inequality, and by Proposition 4.3.1,

$$D(s,t) \le \sum_{i=1}^{k} D(s_i, t_i) + \sum_{i=1}^{k-1} D(t_i, s_{i+1}) \le \sum_{i=1}^{k} d_Z(s_i, t_i)\,.$$

Note that if $k = 1$, we just recover the bound $D \le d_Z$. Therefore, we obtain the upper bound

$$D(s,t) \le D^*(s,t) = \inf\left\{\sum_{i=1}^{k} d_Z(s_i, t_i)\right\}\,,$$

where the infimum is taken over all $k \ge 1$, and $s_1, t_1, \dots, s_k, t_k$ as above. It is now elementary to see that the function D^* thus defined is a pseudo-metric on $[0,1]$. Clearly, it satisfies properties 1. and 2. in Proposition 4.3.1. Moreover, the same argument as previously shows that $d \le D^*$ for any pseudo-metric d on $[0,1]$ that satisfies properties 1. and 2., so that D^* is the maximal pseudo-metric satisfying these two properties. The function D^* is usually called the *quotient pseudo-metric* of d_Z with respect to the equivalence relation $\{d_{\mathbb{e}} = 0\}$: starting from d_Z, it is the pseudo-metric that shrinks the distance as little as it can, while complying to the identifications given by $\{d_{\mathbb{e}} = 0\}$.

The true metric space obtained by taking the quotient $S^* = [0,1]/\{D^* = 0\}$ and equipping it with the class function D^* was called *the Brownian* *map* in Marckert and Mokkadem [63], see also Le Gall [54], from which the above description of D^* is taken. Marckert and Mokkadem conjectured that (S^*, D^*) is the unique scaling limit of rescaled random quadrangulations. This is indeed the case, implying the following result in conjunction with Theorem 4.2.2.

Theorem 4.3.2. *Almost surely, it holds that the functions D and D^* on $[0,1]^2$ are equal. Consequently, the convergence of $(V(Q_n), (9/8n)^{1/4} d_{Q_n})$ to $(S, D) = (S^*, D^*)$ holds in the Gromov-Hausdorff topology, without having to take an appropriate subsequence.*

This result was proved in [56] and [68]. We will give some ideas of the proof in Chapter 7. Before that, and as an intermediate step towards the proof of this result, we will derive some properties that any subsequential limit of the form (S, D) must satisfy. We thus fix a subsequence as appearing in Theorem 4.2.2 once and for all, and call the space (S, D) "the Brownian map", despite the ambiguity that it might represent. This ambiguity will be finally lifted in Chapter 7.

4.4 Basic properties of the Brownian map

4.4.1 Simple geodesics to x_*

The properties 1. and 2. depicted in Proposition 4.3.1 imply that the projection $\mathbf{p} : [0,1] \to S$ factorizes through $p_{\mathbb{e}} : [0,1] \to \mathcal{T}_{\mathbb{e}}$ and $p_Z : [0,1] \to \mathcal{T}_Z$. Specifically, there exist surjective maps $\mathbf{p}_{\mathbb{e}} : \mathcal{T}_{\mathbb{e}} \to S$ and $\mathbf{p}_Z : \mathcal{T}_Z \to S$ such that $\mathbf{p} = \mathbf{p}_{\mathbb{e}} \circ p_{\mathbb{e}}$ and $\mathbf{p} = \mathbf{p}_Z \circ p_Z$. Yet otherwise said, the trees $\mathcal{T}_{\mathbb{e}}$ and $\mathcal{T}_Z$ are naturally "immersed" in S (see below for why we chose this term).

It is interesting to understand what these immersed trees represent in the space (S, D). The tree $\mathcal{T}_{\mathbb{e}}$ is of course the natural analog of the tree T_n, we will see later that it has a natural interpretation as a geometric locus in S. On the other hand, it is natural to figure out what the tree $\mathcal{T}_Z$ is. Indeed, in the discrete picture, the geodesics in $\mathcal{T}_Z$ from $p_Z(t)$ to ρ_Z correspond exactly to the geodesic chains in Q_n obtained by taking the consecutive successors of a corner of T_n in the CVS bijection. It is thus natural that the "branches of $\mathcal{T}_Z$" should encode geodesic paths in (S, D).

For every $t \in [0,1]$, let $\Gamma_t^Z : [0, Z_t - \inf Z] \to \mathcal{T}_Z$ be the geodesic path from $p_Z(t)$ to ρ_Z in $\mathcal{T}_Z$. We let

$$\Gamma_t(r) = \mathbf{p}_Z(\Gamma_t^Z(r)), \qquad 0 \le r \le Z_t - \inf Z,$$

which is a continuous path in (S, D).

Proposition 4.4.1. *For every $t \in [0,1]$, the path Γ_t is a geodesic path in (S, D) from $\mathbf{p}(t)$ to x_*. The elements of the set $\{\Gamma_t, t \in [0,1]\}$ are called the* simple geodesics *to x_* in (S, D).*

PROOF. Recall that $\Gamma_t^Z(r) = p_Z(\gamma_t(r))$, where $\gamma_t(r)$ is the infimum over all s coming in cyclic order after t (where $[0,1]$ is identified with the circle $\mathbb{S}^1$) such that $Z_s < Z_t - r$. In particular, $Z_{\gamma_t(r)} = Z_t - r$ and $\Gamma_t(r) = \mathbf{p}(\gamma_t(r))$.

By definition, for every $r, r' \in [0, Z_t - \inf Z]$, we have $d_Z(\gamma_t(r), \gamma_t(r')) = |r - r'|$. Therefore, $D(\gamma_t(r), \gamma_t(r')) \le d_Z(\gamma_t(r), \gamma_t(r')) = |r - r'|$. Since we know that $D(s, s_*) = Z_s - \inf Z = d_Z(s, s_*)$, this implies, assuming for instance that $r \le r'$,

$$r' - r \ge D(\gamma_t(r), \gamma_t(r')) \ge D(\gamma_t(r), x_*) - D(\gamma_t(r'), x_*) = r' - r,$$

by the triangle inequality. Hence the result. □

4.4.2 Mass measure and re-rooting invariance

The random metric space (S, D) comes with a natural measure μ, the image of the Lebesgue measure on $[0,1]$ by the projection $\mathbf{p} : [0,1] \to S$. By the preceding considerations, this is also the image of the mass measures on $\mathcal{T}_{\mathbb{e}}$ and $\mathcal{T}_Z$ by $\mathbf{p}_{\mathbb{e}}$ and $\mathbf{p}_Z$ respectively, with the notation discussed in the preceding paragraph.

The pointed metric space (S, D, x_*) satisfies an important re-rooting invariance property, similar to that of the Brownian CRT, and that we now explain.

Proposition 4.4.2. *Let $k \ge 1$ and $U_1, U_2, \ldots, U_{k+1}$ be independent uniform random variables in $[0,1]$, independent of other random variables considered so far. Set $X_i = \mathbf{p}(U_i), 1 \le i \le k+1$. Then $(S, D, x_*, X_1, \ldots, X_k)$ and $(S, D, X_1, X_2, \ldots, X_{k+1})$ have the same distribution, seen as random variables taking values in the set of $k+1$-pointed metric spaces.*

PROOF. By using the same method of proof as for Theorem 4.2.2, it is easy to obtain that $(S, D, x_*, X_1, \ldots, X_k)$ is the limit in distribution in $\mathbb{M}^{(k+1)}$ of $(V(Q_n), (9/8n)^{1/4} d_{Q_n}, v_*, u^n_{\langle 2nU_1 \rangle}, \ldots, u^n_{\langle 2nU_k \rangle})$, where the notation $\langle s \rangle$ comes from the proof of Proposition 3.1.4. Clearly, this has a distribution close to that of $(V(Q_n), (9/8n)^{1/4} d_{Q_n}, u^n_{\langle 2nU_1 \rangle}, \ldots, u^n_{\langle 2nU_{k+1} \rangle})$, in the sense that v_* is uniform on $V(Q_n)$ while $u^n_{\langle 2nU \rangle}$ is uniform on $V(Q_n) \setminus \{v_*, u^n_0\}$ conditionally on v_*. We leave the details to the reader. $\square$

4.4.3 Hausdorff dimension

One can show that the Brownian map (S, D) has Hausdorff dimension 4. This was proved in [54].

Theorem 4.4.3. *Almost surely, the space (S, D) has Hausdorff dimension 4.*

PROOF. We start with a preliminary lemma. Recall that $\lambda_{\mathbb{e}}(\cdot)$ denotes the mass measure on $\mathcal{T}_{\mathbb{e}}$, which simply is the image of the Lebesgue measure on $[0, 1]$ under the projection $p_{\mathbb{e}} : [0, 1] \longrightarrow \mathcal{T}_{\mathbb{e}}$.

Lemma 4.4.4. *Almost surely, for every $\delta \in (0, 1)$, there exists a (random) constant $C_\delta(\omega)$ such that, for every $r > 0$ and every $a \in \mathcal{T}_{\mathbb{e}}$,*

$$\lambda_{\mathbb{e}}(\{b \in \mathcal{T}_{\mathbb{e}} : D(a, b) \leq r\}) \leq C_\delta \, r^{4-\delta}.$$

We omit the proof of this lemma. The lower bound is an easy consequence of Lemma 4.4.4. Indeed, Lemma 4.4.4 implies that a.s., for every $\delta \in (0, 1)$, and every $x \in S$, it holds that

$$\limsup_{r \downarrow 0} \frac{\mu(B_D(x, r))}{r^{4-\delta}} = 0\,,$$

where $B_D(x, r) = \{y \in S : D(x, y) < r\}$ is the open ball centered at x with radius r. This last fact, combined with standard density theorems for Hausdorff measures, implies that a.s. the Hausdorff dimension of (S, D) is greater than or equal to $4 - \delta$, for every $\delta \in (0, 1)$.

For the upper bound, we recall from Proposition 3.4.2 that Z is a.s. Hölder continuous with exponent α for every $\alpha \in (0, 1/4)$. From this, we deduce that the projection $\mathbf{p} : [0, 1] \to S$ is a.s. Hölder continuous with index $\alpha \in (0, 1/4)$

as well. Indeed, using the fact that $D \leq d_Z$, where d_Z is defined in (4.4), we get

$$\begin{aligned} D(\mathbf{p}(s), \mathbf{p}(t)) &= D(s,t) \\ &\leq Z_s + Z_t - 2 \inf_{s \wedge t \leq u \leq s \vee t} Z_u \\ &\leq 2 \sup_{s \wedge t \leq u, v \leq s \vee t} |Z_u - Z_v| \\ &\leq C_p'' |s-t|^\alpha , \end{aligned}$$

for some $C_p'' \in (0, \infty)$. The fact that the Hausdorff dimension of (S, D) is bounded above by $1/\alpha$ is then a classical consequence of this last property. This completes the proof of Theorem 4.4.3.

Chapter 5

The topology of the Brownian map

The main goal of this chapter is to prove that the Brownian map is almost surely homeomorphic to the 2-sphere, a theorem by Le Gall and Paulin [59] that justifies calling the Brownian map a “random surface”. To this end, we will have to identify the equivalence relation $\{D = 0\}$ on $[0, 1]$ in terms of the pair $(\mathbb{e}, Z)$.

5.1 Identifying the Brownian map

5.1.1 The Brownian map as a quotient of the CRT

In the previous section, we wrote the scaling limit of rescaled random quadrangulations (along a suitable subsequence) as a quotient space $S = [0, 1]/\approx$ where the equivalence relation $\approx$ is defined by $s \approx t$ iff $D(s, t) = 0$. Here, we provide a more explicit description of this quotient.

Recall the notation of the previous section. In particular, $((T_n, \ell_n), \epsilon)$ is uniformly distributed over $\mathbf{T}_n \times \{-1, 1\}$, and (Q_n, v_*) is the pointed quadrangulation that is the image of $((T_n, \ell_n), \epsilon)$ under the CVS bijection. For every $n \geq 1$, $u_0^n, u_1^n, \ldots, u_{2n}^n$ is the contour exploration of the vertices of T_n. Thus, $C_n(i) = d(u_i^n, u_0^n)$ and $L_n(i) = \ell_n(u_i^n)$ for $0 \leq i \leq 2n$.

As in the proof of Theorem 4.2.2, we may assume that, along the sequence

(n_k) we have the almost sure convergence

$$\begin{aligned}((2n)^{-1/2}C_n(2ns), (9/8n)^{1/4}L_n(2ns), D_n(s,t))_{s,t\in[0,1]} \\ \underset{n\to\infty}{\longrightarrow} (\mathbb{e}_s, Z_s, D(s,t))_{s,t\in[0,1]}\end{aligned} \tag{5.1}$$

uniformly over $[0,1]^2$. Recall from the proof of Theorem 4.2.2 that this implies the almost sure convergence

$$\left(V(Q_n), \left(\frac{9}{8n}\right)^{1/4} d_{Q_n}\right) \underset{n\to\infty}{\longrightarrow} (S, D)$$

in the Gromov-Hausdorff sense, along the sequence (n_k).

As in Section 3.2, introduce the random pseudo-metric $d_{\mathbb{e}}$ and the equivalence relation $\sim_{\mathbb{e}} = \{d_{\mathbb{e}} = 0\}$ on $[0,1]$, so that

$$s \sim_{\mathbb{e}} t \text{ iff } \mathbb{e}_s = \mathbb{e}_t = \min_{s\wedge t\leq r\leq s\vee t} \mathbb{e}_r$$

and recall that the CRT $\mathcal{T}_{\mathbb{e}}$ is defined as the quotient space $[0,1]/\sim_{\mathbb{e}}$ equipped with the distance $d_{\mathbb{e}}$. Also recall that $p_{\mathbb{e}} : [0,1] \longrightarrow \mathcal{T}_{\mathbb{e}}$ denotes the canonical projection. By Proposition 4.3.1, $D(s,t)$ only depends on $p_{\mathbb{e}}(s)$ and $p_{\mathbb{e}}(t)$. We can therefore put for every $a, b \in \mathcal{T}_{\mathbb{e}}$,

$$D(a,b) = D(s,t)$$

where s, resp. t, is an arbitrary representative of a, resp. of b, in $[0,1]$. Then D is (again) a pseudo-metric on $\mathcal{T}_{\mathbb{e}}$. With a slight abuse of notation we keep writing $a \approx b$ iff $D(a,b) = 0$, for $a, b \in \mathcal{T}_{\mathbb{e}}$. Then the Brownian map S can be written as

$$S = [0,1]/\approx\, = \mathcal{T}_{\mathbb{e}}/\approx$$

where the first equality was the definition of S and the second one corresponds to the fact that there is an obvious canonical isometry between the two quotient spaces.

One may wonder why it is more interesting to write the Brownian map S as a quotient space of the CRT $\mathcal{T}_{\mathbb{e}}$ rather than as a quotient space of $[0,1]$. The point is that it will be possible to give a simple intuitive description of $\approx$ viewed as an equivalence relation on $\mathcal{T}_{\mathbb{e}}$. This is indeed the main goal of the next subsection.

5.1.2 Identifying the equivalence relation $\approx$

We noticed in Proposition 3.4.3 that the process Z (the head of the Brownian snake driven by $\mathbb{e}$) can be viewed as indexed by $\mathcal{T}_{\mathbb{e}}$. This will be important in what follows: for $a \in \mathcal{T}_{\mathbb{e}}$, we will write $Z_a = Z_t$ for any choice of t such that $a = p_{\mathbb{e}}(t)$. We also set $a_* = p_{\mathbb{e}}(s_*)$: a_* is thus the unique vertex of $\mathcal{T}_{\mathbb{e}}$ such that

$$Z_{a_*} = \min_{a \in \mathcal{T}_{\mathbb{e}}} Z_a.$$

We first need to define intervals on the tree $\mathcal{T}_{\mathbb{e}}$. For simplicity we consider only leaves of $\mathcal{T}_{\mathbb{e}}$. Recall that a point a of $\mathcal{T}_{\mathbb{e}}$ is a leaf if $\mathcal{T}_{\mathbb{e}} \backslash \{a\}$ is connected. Equivalently a vertex a distinct from the root $\rho_{\mathbb{e}}$ is a leaf if and only if $p_{\mathbb{e}}^{-1}(a)$ is a singleton. Note in particular that a_* is a leaf of $\mathcal{T}_{\mathbb{e}}$.

Let a and b be two (distinct) leaves of $\mathcal{T}_{\mathbb{e}}$, and let s and t be the unique elements of $[0, 1)$ such that $p_{\mathbb{e}}(s) = a$ and $p_{\mathbb{e}}(t) = b$. Assume that $s < t$ for definiteness. We then set

$$\begin{aligned} [a, b] &= p_{\mathbb{e}}([s, t]) \\ [b, a] &= p_{\mathbb{e}}([t, 1] \cup [0, s]). \end{aligned}$$

It is easy to verify that $[a, b] \cap [b, a] = [\![a, b]\!]$ is the line segment between a and b in $\mathcal{T}_{\mathbb{e}}$.

Theorem 5.1.1. *Almost surely, for every distinct $a, b \in \mathcal{T}_{\mathbb{e}}$,*

$$a \approx b \quad \Leftrightarrow \begin{cases} a, b \text{ are leaves of } \mathcal{T}_{\mathbb{e}} \text{ and} \\ Z_a = Z_b = \max\Big(\min_{c \in [a,b]} Z_c, \min_{c \in [b,a]} Z_c\Big) \end{cases}$$

Remark. We know that the minimum of Z over $\mathcal{T}_{\mathbb{e}}$ is reached at the unique vertex a_*. If a and b are (distinct) leaves of $\mathcal{T}_{\mathbb{e}} \backslash \{a_*\}$, exactly one of the two intervals $[a, b]$ and $[b, a]$ contains the vertex a_*. Obviously the minimum of Z over this interval is equal to Z_{a_*} and thus cannot be equal to Z_a or Z_b.

The proof of the implication $\Leftarrow$ in the theorem is easy. Suppose that $a = p_{\mathbb{e}}(s)$ and $b = p_{\mathbb{e}}(t)$ with $s < t$ (for definiteness). If

$$Z_a = Z_b = \max\Big(\min_{c \in [a,b]} Z_c, \min_{c \in [b,a]} Z_c\Big)$$

this means that

$$Z_s = Z_t = \max\Big(\min_{r \in [s,t]} Z_r, \min_{r \in [t,1] \cup [0,s]} Z_r\Big).$$

The last identity is equivalent to saying that $d_Z(s,t) = 0$, and since $D \le d_Z$ we have also $D(s,t) = 0$, or equivalently $a \approx b$.

Unfortunately, the proof of the converse implication is much harder, and we will only give some key ideas of the proof, referring to [54] for additional details.

The first ingredient of the proof is the re-rooting invariance property given in Proposition 4.4.2, which makes it possible to reduce the proof to the case $a = a_*$. In that case we can use the formula $D(a_*, b) = Z_b - \min Z$ and explicit moment calculations for the Brownian snake (see Corollary 6.2 in [55] for a detailed proof).

Let us come to the proof of the implication $\Rightarrow$ in Theorem 5.1.1. For simplicity we consider only the case when a and b are leaves of $\mathcal{T}_{\mathbb{e}}$ (it would be necessary to also show that the equivalence class of any vertex of $\mathcal{T}_{\mathbb{e}}$ that is not a leaf is a singleton – this essentially follows from Lemma 2.2 in [54]). We let $s, t \in [0,1]$ be such that $a = p_{\mathbb{e}}(s)$ and $b = p_{\mathbb{e}}(t)$, and assume for definiteness that $0 \le s_* < s < t \le 1$.

We assume that $a \approx b$, and our goal is to prove that

$$Z_a = Z_b = \min_{c \in [a,b]} Z_c.$$

We already know that $Z_a = Z_b$, because

$$Z_a - \min Z = D(a_*, a) = D(a_*, b) = Z_b - \min Z.$$

First step. We first establish that

$$Z_a = Z_b = \min_{c \in [\![a,b]\!]} Z_c. \tag{5.2}$$

To see this, we go back to the discrete picture. We can find $a_n, b_n \in T_n$ such that $a_n \longrightarrow a$ and $b_n \longrightarrow b$ as $n \to \infty$ (strictly speaking these convergences make no sense: what we mean is that $a_n = u^n_{i_n}$, $b_n = u^n_{j_n}$ with $i_n/2n \longrightarrow s$ and $j_n/2n \longrightarrow t$). Then the condition $D(a,b) = 0$ implies that

$$n^{-1/4}\, d_{Q_n}(a_n, b_n) \longrightarrow 0. \tag{5.3}$$

Recall, from Proposition 2.3.8, the notation $[\![a_n, b_n]\!]$ for the set of vertices lying on the geodesic path from a_n to b_n in the tree T_n. By Proposition 2.3.8(ii), we have

$$d_{Q_n}(a_n, b_n) \ge \ell_n(a_n) + \ell_n(b_n) - 2 \min_{c \in [\![a_n, b_n]\!]} \ell_n(c).$$

We multiply both sides of this inequality by $n^{-1/4}$ and let n tend to ∞, using (5.3). Modulo some technical details that we omit (essentially one needs to check that any vertex of $\mathcal{T}_{\mathbb{e}}$ belonging to $[\![a,b]\!]$ is of the form $p_{\mathbb{e}}(r)$, where $r = \lim k_n/2n$ and the integers k_n are such that $u^n_{k_n}$ belongs to $[\![a_n,b_n]\!]$), we get that

$$Z_a + Z_b - 2 \min_{c\in[\![a,b]\!]} Z_c \leq 0$$

from which (5.2) immediately follows.

Second step. We argue by contradiction, assuming that

$$\min_{c\in[a,b]} Z_c < Z_a = Z_b.$$

Let γ_n be a discrete geodesic from a_n to b_n in the quadrangulation Q_n (here we view a_n and b_n as vertices of the quadrangulation Q_n, and this geodesic is of course different from the geodesic from a_n to b_n in the tree T_n). From (5.3) the maximal distance between a_n (or b_n) and a vertex visited by γ_n is $o(n^{1/4})$ as $n \to \infty$. As a consequence, using the triangle inequality and (2.2), we have

$$\sup_{u\in\gamma_n} |\ell_n(u) - \ell_n(a_n)| = o(n^{1/4})$$

as $n \to \infty$.

To simplify the presentation of the argument, we assume that, for infinitely many values of n, the geodesic path γ_n from a_n to b_n stays in the lexicographical interval $[a_n, b_n]$. This lexicographical interval is defined, analogously to the continuous setting, as the set of all vertices visited by the contour exploration sequence $(u^n_i)_{0\leq i\leq 2n}$ between its last visit of a_n and its first visit of b_n. Note that the preceding assumption may not hold, and so the real argument is slightly more complicated than what follows.

We use the previous assumption to prove the following claim. If $x \in [a,b]$, we denote by $\phi_{a,b}(x)$ the last ancestor of x that belongs to $[\![a,b]\!]$ (the condition $x \in [a,b]$ ensures that the ancestral line $[\![\rho_{\mathbb{e}}, x]\!]$ intersects $[\![a,b]\!]$). Alternatively, $\phi_{a,b}(x)$ is the point of $[\![a,b]\!]$ at minimal $d_{\mathbb{e}}$-distance of x in the tree $\mathcal{T}_{\mathbb{e}}$.

Claim. Let $\varepsilon > 0$. For every $c \in [a,b]$ such that

$$\begin{cases} Z_c < Z_a + \varepsilon & \\ Z_x > Z_a + \varepsilon/2 & \forall x \in [\![\phi_{a,b}(c), c]\!] \end{cases}$$

we have $D(a,c) \leq \varepsilon$.

The claim eventually leads to the desired contradiction: using the first step of the proof (which ensures that $Z_c \geq Z_a$ for $c \in [\![a,b]\!]$) and the properties of the Brownian snake, one can check that, under the condition

$$\min_{c\in[a,b]} Z_c < Z_a = Z_b,$$

the volume of the set of all vertices c that satisfy the assumptions of the claim is bounded from below by a (random) positive constant times ε^2, at least for sufficiently small $\varepsilon > 0$ (see Lemma 2.4 in [54] for a closely related statement). The desired contradiction follows since Lemma 4.4.4 implies that, for every $\delta \in (0,1)$,

$$\lambda_{\mathbb{e}}(\{c : D(a,c) \leq \varepsilon\}) \leq C_\delta\, \varepsilon^{4-\delta}.$$

To complete this sketch, we explain why the claim holds. Again, we need to go back to the discrete setting. We consider a vertex $u \in [a_n, b_n]$ such that

(i) $\ell_n(u) < \ell_n(a_n) + \varepsilon n^{1/4}$;

(ii) $\ell_n(v) > \ell_n(a_n) + \frac{\varepsilon}{2}\, n^{1/4}\ , \qquad \forall v \in [\![\phi^n_{a_n,b_n}(u), u]\!]$

where $\phi^n_{a_n,b_n}(u)$ is the last ancestor of u in the tree T_n that belongs to $[\![a_n, b_n]\!]$.

Condition (ii) guarantees that the vertex u lies "between" $[\![a_n, b_n]\!]$ and the geodesic γ_n: if this were not the case, the geodesic γ_n would contain a point in $[\![\phi^n_{a_n,b_n}(u), u]\!]$, which is impossible by (ii) (we already noticed that the label of a vertex of the geodesic γ_n must be $\ell_n(a_n) + o(n^{1/4})$.

Consider the geodesic path from u to v_* in Q_n that is obtained from the successor geodesic chain $e \to s(e) \to s^2(e) \to \cdots$ starting from any corner e of u in T_n. Since arcs in the CVS bijection do not cross edges of the tree and since we know that the vertex u lies in the area between $[\![a_n, b_n]\!]$ and the geodesic γ_n, the geodesic we have just constructed cannot "cross" $[\![a_n, b_n]\!]$ and so it must intersect γ_n at a vertex w. This vertex w is such that

$$\ell_n(u) - \ell_n(w) = d_{Q_n}(u, w).$$

Since w belongs to γ_n, we have $d_{Q_n}(w, a_n) = o(n^{1/4})$, and therefore

$$\ell_n(u) - \ell_n(a_n) = d_{Q_n}(u, a_n) + o(n^{1/4}).$$

By (i), we now get

$$d_{Q_n}(u, a_n) \leq \varepsilon n^{1/4} + o(n^{1/4}).$$

We have thus obtained a discrete analog of the claim. To get the continuous version as stated above, we just need to do a careful passage to the limit $n \to \infty$. This finishes the sketch of the proof of Theorem 5.1.1.

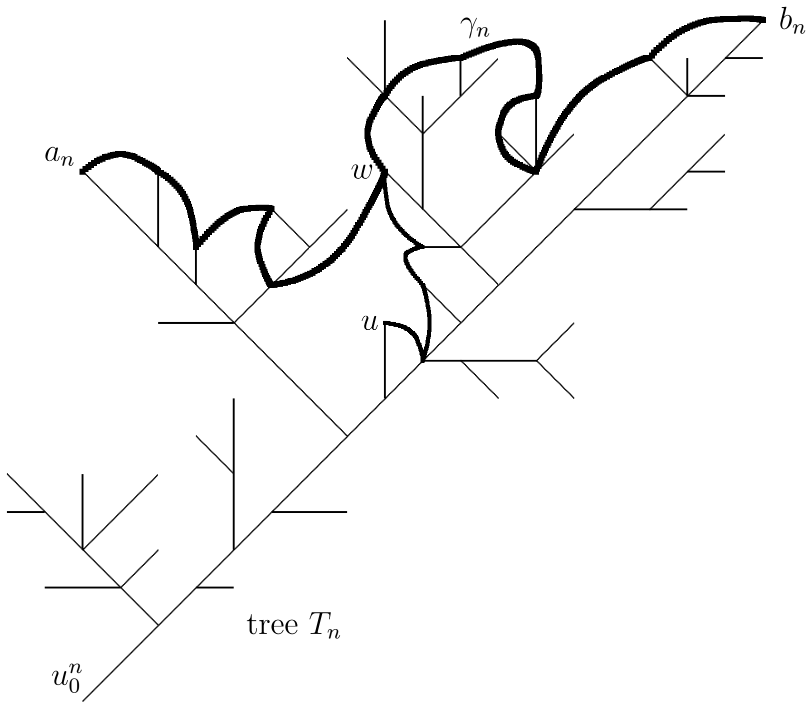

Figure 5.1: Illustration of the proof: The geodesic path γ_n from a_n to b_n is represented by the thick curves. The thin curves correspond to the beginning of the successor geodesic chain starting from u. This chain does not cross the line segment $[[a_n, b_n]]$ and thus has to meet the path γ_n at some point w.

5.2 The sphericity theorem

Theorem 5.2.1. *Almost surely, the Brownian map (S, D) is homeomorphic to the 2-sphere $\mathbb{S}^2$.*

This result was first obtained by Le Gall and Paulin [59], by arguing directly on the quotient space $S = \mathcal{T}_{\mathbb{e}}/\approx$. More precisely, Le Gall and Paulin observed that the equivalence relations $\sim_{\mathbb{e}}$ and $\approx$ may be viewed as equivalence relations on the sphere $\mathbb{S}^2$. Upon showing that the associated classes are closed, arcwise connected, and have connected complements, one can then apply a theorem due to Moore [72], showing that under these hypotheses, the quotient $\mathbb{S}^2/\approx$ is itself homeomorphic to $\mathbb{S}^2$. Here, we will adopt a different approach, introduced in Miermont [66], which relies heavily on the discrete approximations described in these notes. The idea is roughly as follows: even though the property of being homeomorphic to $\mathbb{S}^2$ is not preserved under Gromov-Hausdorff convergence, this preservation can be deduced under an additional property, called *regular convergence*, introduced by Whyburn. This property heuristically says that the spaces under consideration do not have small bottlenecks, i.e. cycles of vanishing diameters that separate the spaces into two macroscopic components.

In this section, when dealing with elements of the space $\mathbb{M}^{(1)}$ of isometry classes of pointed compact metric spaces, we will often omit to mention the distinguished point, as its role is less crucial than it was in Chapter 4 and in Section 5.1.

5.2.1 Geodesic spaces and regular convergence

The set $\mathbb{M}_{\mathrm{geo}}$ of isometry classes of (rooted) compact geodesic metric spaces is closed in $(\mathbb{M}, \mathrm{d}_{\mathrm{GH}})$, as shown in [28].

Definition 5.2.1. *Let $((X_n, d_n), n \geq 1)$ be a sequence of compact geodesic metric spaces, converging to (X, d) in $(\mathbb{M}, \mathrm{d}_{\mathrm{GH}})$. We say that the convergence is* regular *if for every $\varepsilon > 0$, one can find $\delta > 0$ and $N \in \mathbb{N}$ such that, for every $n > N$, every closed path γ in X_n with diameter at most δ is homotopic to 0 in its ε-neighborhood.*

For instance, let Y_n be the complement in the unit sphere $\mathbb{S}^2 \subset \mathbb{R}^3$ of the open $1/n$-neighborhood of the North pole, and endow Y_n with the intrinsic distance induced from the usual Euclidean metric on $\mathbb{R}^3$ (so that the distance between $x, y \in Y_n$ is the minimal length of a path from x to y in Y_n). Let X_n be obtained by gluing two (disjoint) copies of Y_n along their boundaries, and endow it with the natural intrinsic distance. Then X_n converges in the Gromov-Hausdorff sense to a bouquet of two spheres, i.e. two (disjoint) copies of $\mathbb{S}^2$ whose North poles have been identified. However, the convergence is not regular, because the path γ that consists in the boundary of (either copy of) Y_n viewed as a subset of X_n has vanishing diameter as $n \to \infty$, but is not homotopic to 0 in its ε-neighborhood for any $\varepsilon \in (0, 1)$ and for any n. Indeed, such an ε-neighborhood is a cylinder, around which γ makes one turn.

Theorem 5.2.2. *Let $((X_n, d_n), n \geq 1)$ be a sequence of $\mathbb{M}_{\mathrm{geo}}$ that converges regularly to a limit (X, d) that is not reduced to a point. If (X_n, d_n) is homeomorphic to $\mathbb{S}^2$ for every $n \geq 1$, then so is (X, d).*

This theorem is an easy reformulation of a result of Whyburn in the context of Gromov-Hausdorff convergence; see the paper by Begle [11]. In the latter, it is assumed that every X_n should be a compact subset of a compact metric space (Z, δ), independent of n, and that X_n converges in the Hausdorff sense to X. This transfers to our setting, because, if (X_n, d_n) converges to (X, d) in the Gromov-Hausdorff sense, then one can find a compact metric space (Z, δ) containing isometric copies $X'_n, n \geq 1$ and X' of $X_n, n \geq 1$ and X, such that X'_n converges in the Hausdorff sense to X', see for instance [44, Lemma A.1]. In [11], it is also assumed in the definition of regular convergence that for every $\varepsilon > 0$, there exist $\delta > 0$ and $N \in \mathbb{N}$ such that, for every $n \geq N$, any two points of X_n that lie at distance $\leq \delta$ are in a connected subset of X_n of diameter $\leq \varepsilon$. This condition is tautologically satisfied for geodesic metric spaces, which is the reason why we work in this context.

5.2.2 Quadrangulations seen as geodesic spaces

Theorem 5.2.2 gives a natural method to prove Theorem 5.2.1, using the convergence of quadrangulations to the Brownian map, as stated in Theorem 4.2.2. However, the finite space $(V(Q_n), d_{Q_n})$ is certainly not a geodesic space, nor homeomorphic to the 2-sphere. Hence, we have to modify a little these spaces so that they satisfy the hypotheses of Theorem 5.2.2. We will achieve this by constructing a particular[1] graphical representation of any fixed plane quadrangulation $\mathbf{q}$.

[1] The way we do this is by no means canonical. For instance, the emptied cubes X_f used to fill the faces of $\mathbf{q}$ below could be replaced by unit squares for the l^1 metric. However,

Let $(X_f, d_f), f \in F(\mathbf{q})$ be disjoint copies of the emptied unit cube "with bottom removed"

$$\mathcal{C} = [0,1]^3 \setminus \big((0,1)^2 \times [0,1)\big) \ ,$$

endowed with the intrinsic metric d_f inherited from the Euclidean metric (the distance between two points of X_f is the minimal Euclidean length **of a path in** X_f). Obviously each (X_f, d_f) is a geodesic metric space homeomorphic to a closed disk of $\mathbb{R}^2$. We will write elements of X_f in the form $(s,t,r)_f$, where $(s,t,r) \in \mathcal{C}$ and the subscript f is used to differentiate points of the different spaces X_f. The boundary ∂X_f is then the collection of all points $(s,t,r)_f$ for $(s,t,r) \in ([0,1]^2 \setminus (0,1)^2) \times \{0\}$.

Let $f \in F(\mathbf{q})$ and let e_1, e_2, e_3, e_4 be the four oriented edges incident to f enumerated in a way consistent with the counterclockwise order on the boundary (here the labeling of these edges is chosen arbitrarily among the 4 possible labelings preserving the cyclic order). We then define

$$\begin{array}{lcl} c_{e_1}(t) = (t,0,0)_f & , & 0 \le t \le 1 \\ c_{e_2}(t) = (1,t,0)_f & , & 0 \le t \le 1 \\ c_{e_3}(t) = (1-t,1,0)_f & , & 0 \le t \le 1 \\ c_{e_4}(t) = (0,1-t,0)_f & , & 0 \le t \le 1\,. \end{array}$$

In this way, for every oriented edge e of the map $\mathbf{q}$, we have defined a path c_e which goes along one of the four edges of the square ∂X_f, where f is the face located to the left of e.

We define an equivalence relation $\equiv$ on the disjoint union $\amalg_{f \in F(\mathbf{q})} X_f$, as the coarsest equivalence relation such that, for every oriented edge e of $\mathbf{q}$, and every $t \in [0,1]$, we have $c_e(t) \equiv c_{\overline{e}}(1-t)$. By identifying points of the same equivalence class, we glue the oriented sides of the squares ∂X_f pairwise, in a way that is consistent with the map structure. More precisely, the topological quotient $\mathcal{S}_{\mathbf{q}} := \amalg_{f \in F(\mathbf{q})} X_f / \equiv$ is a surface which has a 2-dimensional cell complex structure, whose 1-skeleton $\mathcal{E}_{\mathbf{q}} := \amalg_{f \in F(\mathbf{q})} \partial X_f / \equiv$ is a representative of the map $\mathbf{q}$, with faces (2-cells) $X_f \setminus \partial X_f$. In particular, $\mathcal{S}_{\mathbf{q}}$ is homeomorphic to $\mathbb{S}^2$ by [71, Lemma 3.1.4]. With an oriented edge e of $\mathbf{q}$ one associates an edge of the graph drawing $\mathcal{E}_{\mathbf{q}}$ in $\mathcal{S}_{\mathbf{q}}$, more simply called an edge of $\mathcal{S}_{\mathbf{q}}$, made of the equivalence classes of points in $c_e([0,1])$ (or $c_{\overline{e}}([0,1])$). We also let $\mathcal{V}_{\mathbf{q}}$ be the 0-skeleton of this complex, i.e. the vertices

our choice avoids the existence of too many geodesic paths between vertices of the map in the surface in which it is embedded.

of the graph — these are the equivalence classes of the corners of the squares ∂X_f. We call them the vertices of $\mathcal{S}_{\mathbf{q}}$ for simplicity.

We then endow the disjoint union $\amalg_{f\in F(\mathbf{q})} X_f$ with the largest pseudo-metric $D_{\mathbf{q}}$ that is compatible with $d_f, f \in F(\mathbf{q})$ and with $\equiv$, in the sense that $D_{\mathbf{q}}(x,y) \leq d_f(x,y)$ for $x, y \in X_f$, and $D_{\mathbf{q}}(x,y) = 0$ for $x \equiv y$. Therefore, the function $D_{\mathbf{q}} : \amalg_{f\in F(\mathbf{q})} X_f \times \amalg_{f\in F(\mathbf{q})} X_f \to \mathbb{R}_+$ is compatible with the equivalence relation $\equiv$, and its quotient mapping defines a pseudo-metric on the quotient space $\mathcal{S}_{\mathbf{q}}$, which is still denoted by $D_{\mathbf{q}}$.

Proposition 5.2.3. *The space $(\mathcal{S}_{\mathbf{q}}, D_{\mathbf{q}})$ is a geodesic metric space homeomorphic to $\mathbb{S}^2$. Moreover, the space $(\mathcal{V}_{\mathbf{q}}, D_{\mathbf{q}})$ is isometric to $(V(\mathbf{q}), d_{\mathbf{q}})$, and any geodesic path in $\mathcal{S}_{\mathbf{q}}$ between two elements of $\mathcal{V}_{\mathbf{q}}$ is a concatenation of edges of $\mathcal{S}_{\mathbf{q}}$. Last,*

$$\mathrm{d_{GH}}((V(\mathbf{q}), d_{\mathbf{q}}), (\mathcal{S}_{\mathbf{q}}, D_{\mathbf{q}})) \leq 3\,.$$

Proof. We first check that $D_{\mathbf{q}}$ is a true metric on $\mathcal{S}_{\mathbf{q}}$, i.e. that it separates points. To see this, we use the fact [28, Theorem 3.1.27] that $D_{\mathbf{q}}$ admits the constructive expression:

$$\begin{aligned} &D_{\mathbf{q}}(a,b) \\ &\quad = \inf\left\{\sum_{i=0}^{n} d(x_i,y_i) : n \geq 0, x_0 = a, y_n = b, y_i \equiv x_{i+1} \text{ for } 0 \leq i \leq n-1\right\}, \end{aligned}$$

where we have set $d(x,y) = d_f(x,y)$ if $x, y \in X_f$ for some f, and $d(x,y) = \infty$ otherwise. It follows that, for $a \in X_f \setminus \partial X_f$ and $b \neq a$, $D_{\mathbf{q}}(a,b) > \min(d(a,b), d_f(a, \partial X_f)) > 0$, so a and b are separated.

To verify that $D_{\mathbf{q}}$ is a true metric on $\mathcal{S}_{\mathbf{q}}$, it remains to treat the case where $a \in \partial X_f, b \in \partial X_{f'}$ for some $f, f' \in F(\mathbf{q})$. The crucial observation is that a shortest path in X_f between two points of ∂X_f is entirely contained in ∂X_f. It is then a simple exercise to check that if a, b are in distinct equivalence classes, the distance $D_{\mathbf{q}}(a,b)$ will be larger than the length of some fixed non-trivial path with values in $\mathcal{E}_{\mathbf{q}}$. More precisely, if (the equivalence classes of) a, b belong to the same edge of $\mathcal{S}_{\mathbf{q}}$, then we can find representatives a', b' in the same X_f and we will have $D_{\mathbf{q}}(a,b) \geq d_f(a',b')$. If the equivalence class of a is not a vertex of $\mathcal{S}_{\mathbf{q}}$ but that of b is, then $D_{\mathbf{q}}(a,b)$ is at least equal to the distance of $a \in X_f$ to the closest corner of the square ∂X_f. Finally, if the (distinct) equivalence classes of a, b are both vertices, then $D_{\mathbf{q}}(a,b) \geq 1$. One deduces that $D_{\mathbf{q}}$ is a true distance on $\mathcal{S}_{\mathbf{q}}$, which makes it a geodesic metric space by [28, Corollary 3.1.24]. Since $\mathcal{S}_{\mathbf{q}}$ is a compact topological space, the metric $D_{\mathbf{q}}$ induces the quotient topology on $\mathcal{S}_{\mathbf{q}}$ by [28, Exercise 3.1.14], hence $(\mathcal{S}_{\mathbf{q}}, D_{\mathbf{q}})$ is homeomorphic to $\mathbb{S}^2$.

From the observations in the last paragraph, a shortest path between vertices of $\mathcal{S}_{\mathbf{q}}$ takes values in $\mathcal{E}_{\mathbf{q}}$. Since an edge of $\mathcal{S}_{\mathbf{q}}$ is easily checked to have length 1 for the distance $D_{\mathbf{q}}$, such a shortest path will have the same length as a geodesic path for the (combinatorial) graph distance between the two vertices. Hence $(\mathcal{V}_{\mathbf{q}}, D_{\mathbf{q}})$ is indeed isometric to $(V(\mathbf{q}), d_{\mathbf{q}})$. The last statement follows immediately from this and the fact that $\operatorname{diam}(X_f, d_f) \leq 3$, entailing that $\mathcal{V}_{\mathbf{q}}$ is 3-dense in $(\mathcal{S}_{\mathbf{q}}, D_{\mathbf{q}})$, i.e. its 3-neighborhood in $(\mathcal{S}_{\mathbf{q}}, D_{\mathbf{q}})$ equals $\mathcal{S}_{\mathbf{q}}$. $\square$

As a consequence of the proposition, we can view $D_{\mathbf{q}}$ as an extension to $\mathcal{S}_{\mathbf{q}}$ of the graph distance $d_{\mathbf{q}}$ on $V(\mathbf{q})$. For this reason, we will denote $D_{\mathbf{q}}$ by $d_{\mathbf{q}}$ from now on, which should not cause any ambiguity.

5.2.3 Proof of the sphericity theorem

We now work in the setting of the beginning of Section 5.1.1. Recall that the uniform pointed quadrangulation (Q_n, v_*) is encoded by a uniform random element (T_n, ℓ_n) of $\mathbf{T}_n$ via the CVS bijection (the parameter $\epsilon \in \{-1, 1\}$ will play no role here), and that C_n and L_n are the contour and label processes of (T_n, ℓ_n). We assume that the almost sure convergence (5.1) holds uniformly on $[0,1]^2$, along the sequence (n_k), which is fixed. In what follows, all convergences as $n \to \infty$ hold along this sequence, or along some further subsequence.

We can also assume that $(V(Q_n), d_{Q_n})$ is actually the (isometric) space $(\mathcal{V}_{Q_n}, d_{Q_n})$, i.e. the subspace of vertices of the space $(\mathcal{S}_{Q_n}, d_{Q_n})$ constructed in the previous Section. Recall from Section 2.3 that, in the CVS bijection, each edge of the tree T_n lies in exactly one face of Q_n. Therefore, we may and will assume that T_n is also embedded in the surface $\mathcal{S}_{Q_n}$, in such a way that the set of its vertices is $\mathcal{V}_{Q_n} \setminus \{v_*\}$, and that each edge of T_n lies entirely in the corresponding face of $\mathcal{S}_{Q_n}$ via the CVS bijection. Here, we identified $v_* \in V(Q_n)$ with its counterpart in $\mathcal{V}_{Q_n}$.

We will rely on the following lemma. Recall that $\mathrm{Sk}(\mathcal{T}_{\mathbb{e}})$ denotes the skeleton of $\mathcal{T}_{\mathbb{e}}$ (see Section 3.5.3).

Lemma 5.2.4. *The following property is true with probability* 1. *Let* $a \in \mathrm{Sk}(\mathcal{T}_{\mathbb{e}})$*, and let* $s \in (0,1)$ *be such that* $a = p_{\mathbb{e}}(s)$*. Then for every* $\varepsilon > 0$*, there* *exists* $t \in (s, (s+\varepsilon) \wedge 1)$ *such that* $Z_t < Z_s$.

This lemma is a consequence of [59, Lemma 3.2] (see also [54, Lemma 2.2] for a slightly weaker statement). The proof relies on a precise study of the label function Z, and we refer the interested reader to [59]. Note that this result (and the analogous statement derived by time-reversal) implies that a.s., if $a \in \mathrm{Sk}(\mathcal{T}_{\mathbb{e}})$, then in each component of $\mathcal{T}_{\mathbb{e}} \setminus \{a\}$, one can find points b that are arbitrarily close to a and such that $Z_b < Z_a$.

Lemma 5.2.5. *Almost surely, for every* $\varepsilon > 0$, *there exists* $\delta \in (0, \varepsilon)$ *such that, for* n *large enough, any simple loop* γ_n *made of edges of* $\mathcal{S}_{Q_n}$, *with diameter* $\leq n^{1/4}\delta$, *splits* $\mathcal{S}_{Q_n}$ *in two Jordan domains, one of which has diameter* $\leq n^{1/4}\varepsilon$.

Proof. We argue by contradiction. Assume that, with positive probability, along some (random) subsequence of (n_k) there exist simple loops γ_n made of edges of $\mathcal{S}_{Q_n}$, with diameters $o(n^{1/4})$ as $n \to \infty$, such that the two Jordan domains bounded by γ_n are of diameters $\geq n^{1/4}\varepsilon$, where $\varepsilon > 0$ is some fixed constant. From now on we argue on this event. By abuse of notation we will sometimes identify the chain γ_n with the set of vertices it visits, or with the union of its edges, in a way that should be clear from the context.

By the Jordan curve theorem, the path γ_n splits $\mathcal{S}_{Q_n}$ into two Jordan domains, which we denote by $\mathcal{D}_n$ and $\mathcal{D}'_n$. Since the diameters of both these domains are at least $n^{1/4}\varepsilon$, and since every point in $\mathcal{S}_{Q_n}$ is at distance at most 3 from some vertex, we can find vertices y_n and y'_n belonging to $\mathcal{D}_n$ and $\mathcal{D}'_n$ respectively, and which lie at distance at least $n^{1/4}\varepsilon/4$ from γ_n. Since $V(Q_n) = T_n \cup \{v_*\}$, we can always assume that y_n and y'_n are distinct from v_*. Now, consider the geodesic path from y_n to y'_n in T_n, and let x_n be the first vertex of this path that belongs to γ_n.

In the contour exploration around T_n, the vertex x_n is visited at least once in the interval between y_n and y'_n, and another time in the interval between y'_n and y_n. More precisely, let j_n and j'_n be such that $y_n = u^n_{j_n}, y'_n = u^n_{j'_n}$, and assume first that $j_n < j'_n$ for infinitely many n. For such n, we can find integers $i_n \in (j_n, j'_n)$ and $i'_n \in (0, j_n) \cup (j'_n, 2n)$ such that $x_n = u^n_{i_n} = u^n_{i'_n}$. Up to further extraction, we may and will assume that

$$\frac{i_n}{2n} \to s\,, \qquad \frac{i'_n}{2n} \to s'\,, \qquad \frac{j_n}{2n} \to t\,, \qquad \frac{j'_n}{2n} \to t'\,, \tag{5.4}$$

for some $s, s', t, t' \in [0,1]$ such that $t \leq s \leq t'$ and $s' \in [0,t] \cup [t,1]$. Since

$$d_{Q_n}(x_n, y_n) \wedge d_{Q_n}(x_n, y'_n) \geq n^{1/4}\varepsilon/4\,,$$

we deduce from (5.1) that $D(s,t), D(s',t), D(s,t'), D(s',t') > 0$, and in particular, s, s', t, t' are all distinct. Since $u^n_{i_n} = u^n_{i'_n}$, we conclude that $s \sim_{\mathbb{e}} s'$, so that $p_{\mathbb{e}}(s) \in \mathrm{Sk}(\mathcal{T}_{\mathbb{e}})$. One obtains the same conclusion by a similar argument if $j_n > j'_n$ for every n large. We let $x = p_{\mathbb{e}}(s)$ and $y = p_{\mathbb{e}}(t)$. Note that $y \neq x$ because $D(s,t) > 0$ (recall 1. in Proposition 4.3.1).

Since $x \in \mathrm{Sk}(\mathcal{T}_{\mathbb{e}})$, by Theorem 5.1.1 we deduce that $D(a_*, x) = D(s_*, s) > 0$, where $a_* = p_{\mathbb{e}}(s_*)$ is as before the a.s. unique leaf of $\mathcal{T}_{\mathbb{e}}$ where Z attains its minimum. In particular, we obtain by (4.1), (5.1) and the fact that $\mathrm{diam}\,(\gamma_n) = o(n^{1/4})$ that

$$\liminf_{n\to\infty} n^{-1/4} d_{Q_n}(v_*, \gamma_n) = \liminf_{n\to\infty} n^{-1/4} d_{Q_n}(v_*, x_n) > 0\,.$$

Therefore, for n large enough, v_* does not belong to γ_n, and for definiteness, we will assume that for such n, $\mathcal{D}_n$ is the component of $\mathcal{S}_{Q_n} \setminus \gamma_n$ that does not contain v_*.

Now, we let $\ell^+_n = \ell_n - \min \ell_n + 1$, and in the rest of this proof, we call $\ell^+_n(v) = d_{Q_n}(v_*, v)$ the *label of the vertex* v in Q_n. Let $l_n = d_{Q_n}(v_*, \gamma_n) = \min_{v\in\gamma_n} \ell^+_n(v)$ be the minimal distance from v_* to a point visited by γ_n. Note that, for every vertex $v \in \mathcal{D}_n$, the property $\ell^+_n(v) \geq l_n$ holds, since any geodesic chain from v_* to v in Q_n has to cross γ_n.

Recalling that the vertex x_n was chosen so that the simple path in T_n from x_n to y_n lies entirely in $\mathcal{D}_n$, we conclude that the labels of vertices on this path are all greater than or equal to l_n. By passing to the limit, one concludes that for every c in the path $[\![x,y]\!]$ in $\mathcal{T}_{\mathbb{e}}$, there holds that $Z_c \geq Z_x$. Since the process Z evolves like a Brownian motion along line segments of the tree $\mathcal{T}_{\mathbb{e}}$, we deduce that for every $c \in [\![x,y]\!]$ close enough to x, we have in fact $Z_c > Z_x$. From the interpretation of line segments in $\mathcal{T}_{\mathbb{e}}$ in terms of the encoding function $\mathbb{e}$, we can find $\overline{s} \in (0,1)$ such that $p_{\mathbb{e}}(\overline{s}) = x$, and such that, for every $u > \overline{s}$ sufficiently close to $\overline{s}$, the intersection of $[\![x, p_{\mathbb{e}}(u)]\!]$ with $[\![x,y]\!]$ will be of the form $[\![x, p_{\mathbb{e}}(r)]\!]$ for some $r \in (\overline{s}, u]$. By Lemma 5.2.4, and the fact that $Z_c \geq Z_x$ for every $c \in [\![x,y]\!]$ close enough to x, we can find $u > \overline{s}$ encoding a point $a = p_{\mathbb{e}}(u)$ and some $\eta > 0$ such that $Z_a \leq Z_x - (9/8)^{1/4}\eta$, and such that $[\![x,a]\!] \cap [\![x,y]\!] = [\![x,b]\!]$ for some $b \neq x$ such that $Z_b \geq Z_x + (9/8)^{1/4}\eta$.

We then go back once again to the discrete approximations of the Brownian map, by considering k_n such that $k_n/2n$ converges to u. From the fact that $Z_a < Z_x$, we deduce that the vertex $a_n = u^n_{k_n}$ has label $L^+_n(a_n) < l_n$ for every n large enough. Indeed, the convergence (5.1) and the fact that

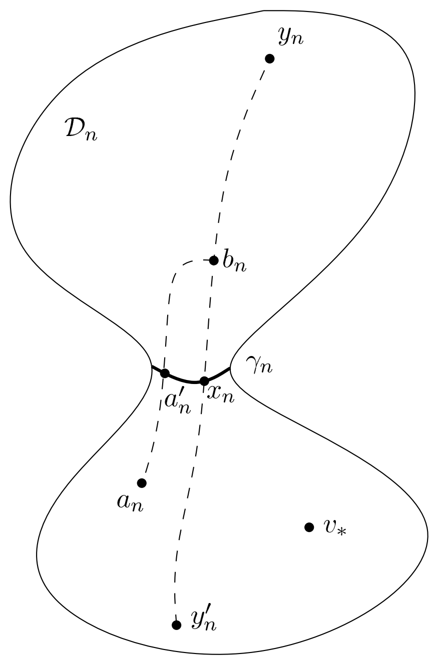

Figure 5.2: Illustration of the proof. The surface $\mathcal{S}_{Q_n}$ is depicted as a sphere with a bottleneck circled by γ_n (thick line). The dashed lines represent paths of T_n that are useful in the proof: one enters the component $\mathcal{D}_n$, and the other goes out after entering, identifying in the limit a point of the skeleton with another.

$\mathrm{diam}\,(\gamma_n) = o(n^{1/4})$ imply that $(9/8n)^{1/4} l_n \to Z_x - \inf Z$. Consequently, the point a_n does not belong to $\mathcal{D}_n$. Moreover, the path in T_n from a_n to x_n meets the path from x_n to y_n at a point b_n such that $\ell_n^+(b_n) \geq l_n + \eta n^{1/4}$. The path from a_n to b_n has to cross the loop γ_n at some vertex, and we let a'_n be the first such vertex. By letting $n \to \infty$ one last time, we find a vertex $a' \in \mathcal{T}_{\mathbb{e}}$, which in the appropriate sense is the limit of a'_n as $n \to \infty$, such that $[\![a', x]\!]$ meets $[\![x, y]\!]$ at b. In particular, $a' \neq x$. But since a'_n and x_n are both on γ_n, we deduce that $D(a', x) = 0$. This contradicts Theorem 5.1.1 because x is not a leaf of $\mathcal{T}_{\mathbb{e}}$. This contradiction completes the proof of the lemma. □

We claim that Lemma 5.2.5 suffices to verify that the convergence of $(V(Q_n), (9/8n)^{1/4} d_{Q_n})$ to (S, D) is regular, and hence to conclude by Theorem 5.2.2 that the limit (S, D) is a topological sphere. To see this, we first choose $\varepsilon < \mathrm{diam}\,(S)/3$ to avoid trivialities. Let γ_n be a loop in $\mathcal{S}_{Q_n}$ with diameter $\leq n^{1/4}\delta$. Consider the union of the closures of faces of $\mathcal{S}_{Q_n}$ that are visited by γ_n. The boundary of this union is a collection $\mathcal{L}$ of pairwise disjoint simple loops made of edges of $\mathcal{S}_{Q_n}$. If x, y belong to the preceding union of faces, the fact that a face of $\mathcal{S}_{Q_n}$ has diameter less than 3 implies that there exist points x' and y' of γ_n at distance at most 3 from x and y respectively. Therefore, the diameters of the loops in $\mathcal{L}$ all are $\leq n^{1/4}\delta + 6$.

By the Jordan Curve Theorem, each of these loops splits $\mathcal{S}_{Q_n}$ into two simply connected components. By definition, one of these two components contains γ_n entirely. By Lemma 5.2.5, one of the two components has diameter $\leq n^{1/4}\varepsilon$. If we show that the last two properties hold simultaneously for one of the two components associated with (at least) one of the loops in $\mathcal{L}$, then obviously γ_n will be homotopic to 0 in its ε-neighborhood in $(\mathcal{S}_{Q_n}, n^{-1/4} d_{Q_n})$. So assume the contrary: the component not containing γ_n associated with every loop of $\mathcal{L}$ is of diameter $\leq n^{1/4}\varepsilon$. If this holds, then any point in $\mathcal{S}_{Q_n}$ must be at distance at most $n^{1/4}\varepsilon + 3$ from some point in γ_n. Take x, y such that $d_{Q_n}(x, y) = \mathrm{diam}\,(\mathcal{S}_{Q_n})$. Then there exist points x' and y' in γ_n at distance at most $n^{1/4}\varepsilon + 3$ respectively from x and y, and we conclude that $d_{Q_n}(x', y') \geq \mathrm{diam}\,(\mathcal{S}_{Q_n}) - 6 - 2n^{1/4}\varepsilon > n^{1/4}\delta \geq \mathrm{diam}\,(\gamma_n)$ for n large enough by our choice of ε. This contradiction completes the proof.

Notes for Chapter 5

As we mentioned above, the original approach to the proof of the sphericity theorem, due to Le Gall and Paulin, does not rely on a particular approximation of the Brownian map, but directly argues in the continuous world. It relies on beautiful ideas coming from the world of complex dynamics, and in particular techniques introduced by Thurston in the context of the so-called "mating of polynomials" of Douady-Hubbard. We sketch it briefly here. First of all, the Brownian map is homeomorphic to the topological quotient $[0,1]/\{D=0\}$, and Theorem 5.1.1 shows that

$$\{D=0\} = \{d_{\mathbb{e}}=0\} \cup \{d_Z=0\}.$$

There is a way to look at the quotient $[0,1]/\{D=0\}$ that starts with augmenting the base space $[0,1]$ in the following way. Since $d_{\mathbb{e}}(0,1)=0$, we can in fact view $\{D=0\}$ as an equivalence relation on the circle $\mathbb{S}^1=[0,1]/0\sim 1$, and then we can in turn see $\mathbb{S}^1=\partial\mathbb{D}$ as the boundary of the open unit circle in $\mathbb{R}^2$. Now whenever $x,y\in\mathbb{S}^1$ are such that $d_{\mathbb{e}}(x,y)=0$, let us draw a chord between x and y in the closure $\overline{\mathbb{D}}$ of $\mathbb{D}$ (for reasons pertaining to hyperbolic geometry, the chord in question is often chosen to be an arc of a circle, that is a geodesic in the Poincaré disc model for hyperbolic space — it also yields nicer pictures). The resulting collection of geodesics is called a *geodesic lamination*, that is a collection of pairwise disjoint geodesics in $\mathbb{D}$. It can be shown that it is also maximal: any geodesic in the Poincaré disc intersects at least one of the geodesics of the lamination. Let $\sim_{\mathbb{e}}$ be the smallest equivalence relation on $\overline{\mathbb{D}}$ such that $x\sim_{\mathbb{e}} y$ if and only if

- either $x=y$,
- or x and y belong to the same geodesic of the lamination (plus the endpoints in $\mathbb{S}^1$),
- or x and y belong to the same (filled-in) ideal geodesic triangle of the lamination (plus the endpoints).

It can be seen that $[0,1]/\{d_{\mathbb{e}}=0\}=\mathcal{T}_{\mathbb{e}}=\overline{\mathbb{D}}/\sim_{\mathbb{e}}$. We define similarly an equivalence relation $\sim_Z$ on $\mathbb{D}$ by drawing chords between points x,y of $\mathbb{S}^1$ such that $d_Z(x,y)=0$, and identifying all the points of these respective chords (together with the incident geodesic triangle if need be). Then $\mathcal{T}_Z=[0,1]/\{d_Z=0\}=\overline{\mathbb{D}}/\sim_Z$.

Now let us distinguish the two constructions above by letting $\overline{\mathbb{D}}$ be identified with the upper hemisphere $\mathbb{H}_+$ of $\mathbb{S}^2$ in the first construction, and with the lower hemisphere in the second. We see that

$$[0,1]/\{D=0\} \simeq [0,1]/(\{d_{\mathbb{e}}=0\} \cup \{d_Z=0\}) \simeq \mathbb{S}^2/(\sim_{\mathbb{e}} \cup \sim_Z),$$

where $X \simeq Y$ means that the topological spaces X and Y are homeomorphic.

Almost surely, the equivalence relation $\sim=\sim_{\mathbb{e}} \cup \sim_Z$ clearly does not identify all points together, and satisfies the property that its equivalence classes are closed connected subsets of $\mathbb{S}^2$, whose complements are connected: this property comes from the fact that points identified in $\sim_{\mathbb{e}}$ are almost surely the unique elements of their equivalence classes for $\sim_Z$, and vice-versa. This is exactly what is needed to apply a celebrated theorem by Moore, entailing that $\mathbb{S}^2/\sim$ is homeomorphic to $\mathbb{S}^2$.

Chapter 6

The multi-pointed bijection and some applications

6.1 The multi-pointed CVS bijection

In order to prove further properties of the Brownian map, we need to introduce another bijection, that gives information on *multi-pointed* maps. Let $\mathbf{q}$ be a rooted quadrangulation and $\mathbf{v} = (v_1, \ldots, v_k) \in V(\mathbf{q})^k$ for some $k \geq 1$. A delay vector is an element $\mathbf{d} \in \mathbb{Z}^k/\mathbb{Z}$, that is, a integer vector $(d_1, \ldots, d_k)$ defined up to a common additive constant. We usually note $[d_1, \ldots, d_k]$ such equivalence classes. We say that the delay vector is compatible with $(\mathbf{q}, \mathbf{v})$ if

1. for every distinct $i, j \in \{1, 2, \ldots, k\}$, we have $|d_i - d_j| < d_{\mathbf{q}}(v_i, v_j)$
2. for every $i, j \in \{1, 2, \ldots, k\}$, we have $d_{\mathbf{q}}(v_i, v_j) + d_i - d_j \in 2\mathbb{N}$.

Note that in these conditions, the quantity $d_i - d_j$ does not depend on the particular choice of a representative of the delay vector. We will soon explain the meaning of these two conditions.

For $k \geq 1$ we let $\mathcal{Q}^{(k)}$ be the set of triples $(\mathbf{q}, \mathbf{v}, \mathbf{d})$, where

- $\mathbf{q}$ is a rooted quadrangulation
- $\mathbf{v} = (v_1, \ldots, v_k)$ are k vertices of $\mathbf{q}$
- $\mathbf{d} = [d_1, \ldots, d_k]$ is a delay vector compatible with $(\mathbf{q}, \mathbf{v})$.

Let also $\mathcal{Q}_n^{(k)}$ be those elements of $\mathcal{Q}^{(k)}$ that have n faces. Note that for instance, $\mathcal{Q}^{(1)}$ is nothing but the set $\mathcal{Q}^\bullet$ of rooted, pointed quadrangulations. Indeed, for $k = 1$, the set of delay vector is a singleton and can be forgotten.

Remark. For a given marked quadrangulation $(\mathbf{q}, \mathbf{v})$, it might be the case that there are no delay vectors compatible with $(\mathbf{q}, \mathbf{v})$. This is the clearly the case if $v_i = v_j$ for some $i \neq j$ (because of the first condition), or if v_i and v_j are neighbors for some $i \neq j$ (the first condition forces $d_i - d_j = 0$, and the second condition cannot be true).

These are in fact the only cases: as soon as

$$\min\{d_{\mathbf{q}}(v_i, v_j) : 1 \leq i < j \leq k\} \geq 2,$$

one can find a compatible delay vector. To see this, simply take d_i to be the representative in $\{0, 1\}$ of $d_{\mathbf{q}}(v_i, v_j)$ modulo 2. Then clearly,

$$|d_i - d_j| \leq 1 < 2 \leq d_{\mathbf{q}}(v_i, v_j)$$

for every $i \neq j$ by assumption, so the first condition is satisfied. Moreover,

$$d_i - d_j \equiv d_{\mathbf{q}}(v_1, v_i) - d_{\mathbf{q}}(v_1, v_j) \equiv d_{\mathbf{q}}(v_i, v_j)$$

modulo 2, where the last congruence uses the fact that $\mathbf{q}$ is a bipartite map.

Given an element of $Q^{(k)}$, we can define a label function ℓ on $V(\mathbf{q})$ in the following way. Choose a representative $(d_1, \ldots, d_k)$ of $\mathbf{d}$ arbitrarily, and let

$$\ell(v) = \min_{1 \leq i \leq k} (d_{\mathbf{q}}(v, v_i) + d_i).$$

We let $\ell_i(v) = d_{\mathbf{q}}(v, v_i) + d_i$ be the quantity that must be minimized.

What does the quantity $\ell(v)$ represent? Imagine that $d_i = 0$ for every $i \in \{1, 2, \ldots, k\}$ (which requires that $d_{\mathbf{q}}(v_i, v_j)$ is even for every i, j). Then $\ell(v)$ is simply the distance of v to the set $\{v_1, \ldots, v_k\}$, that is the distance to the closest point in this set. Yet otherwise said, $\ell(v)$ is the distance of v to the center of its Voronoi cell with respect to the metric space $(V(\mathbf{q}), d_{\mathbf{q}})$ marked by $v_1, \ldots, v_k$. Moreover, $\ell(v) = \ell_i(v)$ iff v is in the Voronoi cell centered at v_i.

The idea to introduce delay vectors is, roughly speaking, to move away from Voronoi cells and to allow more general geometric loci that can be described in terms of competing cells, which start growing at vertex v_i at time d_i and expand at unit speed without being able to overlap. Let us make this idea more precise, starting with a lemma. From here on, we fix an element $(\mathbf{q}, \mathbf{v}, \mathbf{d}) \in \mathcal{Q}^{(k)}$.

Lemma 6.1.1. *For every adjacent $u, v \in V(\mathbf{q})$, it holds that*

$$|\ell(u) - \ell(v)| = 1.$$

PROOF. The same property with ℓ_i instead of ℓ is clearly satisfied by bipartition of $\mathbf{q}$. Since $\ell = \min_i \ell_i$, we immediately deduce that $|\ell(u)-\ell(v)| \leq 1$ for every u, v adjacent. Let us imagine that we can find adjacent u, v such that $\ell(u) = \ell(v)$. Then we can find i, j such that $\ell(u) = \ell_i(u) = \ell(v) = \ell_j(v)$, and necessarily, $i \neq j$ by the remark made at the start of the proof. By definition of ℓ_i, we obtain

$$d_{\mathbf{q}}(u, v_i) - d_{\mathbf{q}}(v, v_j) + d_i - d_j \equiv 0[\text{mod } 2].$$

Since u and v are adjacent, we have $|d_{\mathbf{q}}(u, v_i) - d_{\mathbf{q}}(v, v_i)| = 1$, and again by bipartition, $d_{\mathbf{q}}(v, v_i) - d_{\mathbf{q}}(v, v_j) \equiv d_{\mathbf{q}}(v_i, v_j)$ modulo 2. Putting all together, we obtain

$$d_{\mathbf{q}}(v_i, v_j) + d_i - d_j \equiv 1[\text{mod } 2],$$

contradicting the definition of a compatible delay vector. □

With this lemma at hand, note that every edge in $E(\mathbf{q})$ can be canonically oriented in such a way that it points toward the vertex of lesser label:

$$\ell(e^+) = \ell(e^-) - 1.$$

We adopt this convention here. Let e be an edge thus oriented, and consider the oriented path starting from e, that turns to the left as much as possible (we call this the leftmost path started from e). More precisely, if e' is an oriented edge in this path, then the next one is the first outgoing edge from the target of e' (if any), when going around the latter in clockwise order starting from e'.

In fact, the function ℓ decreases strictly along oriented paths in $\mathbf{q}$, and therefore have to get stuck at vertices to which all incident edges point, that are the local minima of ℓ. It is not difficult to convince oneself that the vertices satisfying this property exactly $v_1, \ldots, v_k$. For instance, v_1 has this property: indeed, for every $i \neq 1$,

$$\ell_1(v_1) = d_1 < d_{\mathbf{q}}(v_1, v_i) + d_i = \ell_i(v_1),$$

and the parity condition implies that $\ell_i(v_1) - \ell_1(v_1) \geq 2$. Since ℓ_i varies by ± 1 when going to an adjacent vertex, this shows that $\ell_1(v) \leq \ell_i(v)$ for all vertices v adjacent to v_1, so $\ell(v) = \ell_1(v) = \ell(v_1) + 1$, and indeed, v_1 is a local minimum of ℓ, the same being of course true of all v_i's. Conversely, suppose that v is a local minimum of ℓ. Let i be an index such that $\ell_i(v) = \ell(v)$. Necessarily, it must hold that $\ell_i(u) = \ell_i(v) + 1$ for every u adjacent to v, otherwise we would have $\ell(u) = \ell(v) - 1$.

As a consequence, any (oriented) edge is the start of a unique, maximal leftmost part that finishes at one element of $\{v_1, \dots, v_k\}$. We let $E_i(\mathbf{q}, \mathbf{v}, \mathbf{d})$ be the set of edges whose associated leftmost path ends at v_i; these sets partition the set of edges of $\mathbf{q}$. We call these sets the $\mathbf{d}$-delayed Voronoi cells of $(\mathbf{q}, \mathbf{v})$. We leave as a simple exercise to check that for every $e \in E_i(\mathbf{q}, \mathbf{v}, \mathbf{d})$, it holds that

$$\ell(e^-) = \ell_i(e^-) = d_{\mathbf{q}}(e^-, v_i) + d_i \,. \tag{6.1}$$

Let us now apply the same construction as in the Cori-Vauquelin-Schaeffer bijection, adding edges in the faces of $\mathbf{q}$ depending on the labels around each face, according to Figure 6.1. To differentiate these extra edges, we call them the "red edges" due to the color used in this figure. We let $\mathbf{q}'$ be the embedded graph obtained by augmenting the quadrangulation $\mathbf{q}$ with the red edges, and it is clearly a map.

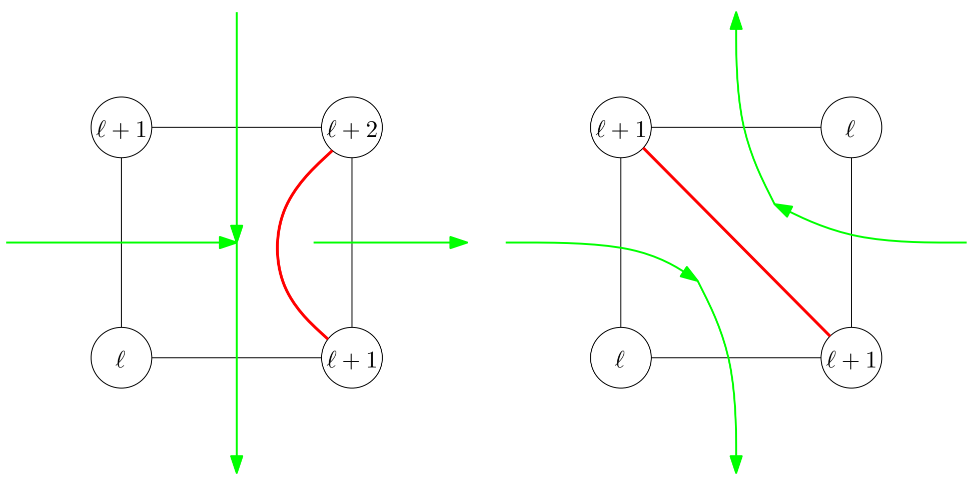

Figure 6.1: The CVS convention, and the orientation convention on the dual of $\mathbf{q}'$.

Claim. The embedded graph formed by the red edges, together with the vertices incident to them, is a plane map $\mathbf{m}$ with k faces, and the label function ℓ induces an admissible labeling of this graph, in the same sense as for trees: for every adjacent vertices $u, v \in V(\mathbf{m})$, it holds that $|\ell(u) - \ell(v)| \leq 1$.

To prove this claim, we introduce a method that can be found in Chapuy-Marcus-Schaeffer [29]. This method is more general than the one we used to study the CVS bijection, and in particular it works in arbitrary genera. The idea consists in orienting the edges of the dual graph of $\mathbf{q}'$ that do not cross the red edges, in such a way that the oriented edges cross the original edges of $\mathbf{q}$ leaving the vertex of lesser label always to the right. These correspond to the green oriented edges of Figure 6.1.

These green oriented edges form oriented paths in the dual graph of $\mathbf{q}'$, and note that from any vertex of this graph, there is only one oriented path starting from this vertex, because there is exactly one outgoing edge from this vertex (i.e. from every face of $\mathbf{q}'$). Ultimately, these oriented paths must end up cycling. Moreover, consider two consecutive green oriented edges, and assume that the vertices of $\mathbf{q}$ located to the right of these edges have the same label (i.e. does not decrease). Then a quick inspection of the cases shows that these vertices must in fact be the same, that is, the two green edges are dual to two consecutive edges of $\mathbf{q}$ in clockwise order around some vertex of $\mathbf{q}$. For this reason, the green oriented cycles must circle around a single vertex of locally minimal label, that is, one of $\{v_1, \ldots, v_k\}$. Conversely, every vertex in the latter set is a center of a green oriented cycle.

Hence, the green oriented graph is formed of k components, and these components, seen as unoriented graphs, are unicycles, i.e. graphs with only one cycle. Now, the graph $\mathbf{m}$ is obtained from $\mathbf{q}'$ by removing the edges of $\mathbf{q}$, and another way to say this is that we glue together the faces of $\mathbf{q}'$ along the green edges. In order to show that $\mathbf{m}$ is a map, we must see that this gluing generates only simply connected domains. Clearly this is so when we glue the faces around the green cycle around the vertex v_i, which amounts to make all the faces incident to v_i into a single face. Next, we can remove edges of $\mathbf{q}$ inductively, e.g. following the green tree components in depth-first order. In doing so, at each step we glue a simply connected domain with a face of $\mathbf{q}$ along an edge of $\mathbf{q}$, and by a theorem by Van Kampen, the result is still simply connected. At the end of this procedure, we have produced k simply connected domains, one for each green component, and only the red edges remain. We deduce that $\mathbf{m}$ is indeed a map with k faces.

We can make $\mathbf{m}$ into a rooted map by adopting a rooting convention similar to that of the CVS bijection, we will not mention it here as it will not be very important to us.

Each face of $\mathbf{m}$ naturally receives the label of the corresponding v_i, so we let $f_1, \ldots, f_k$ be the accordingly labeled faces. The fact that ℓ is an admissible labeling on $\mathbf{m}$ is clear by definition.

Let $\mathbb{LM}^{(k)}$ be the set of rooted maps with k labeled faces $f_1, \ldots, f_k$ endowed with an admissible labeling (defined up to an additive constant). We also let $\mathbb{LM}_n^{(k)}$ be the subset of those maps that have n edges, this subset is not empty as soon as $n \geq k-1$. The analog of the "CVS bijection theorem" is the following.

Theorem 6.1.2. *The mapping $\Phi^{(k)} : \mathcal{Q}_n^{(k)} \to \mathbb{LM}_n^{(k)}$ is two-to-one.*

We will not give the proof here, referring the interested reader to [67], but let us describe the inverse construction. The latter simply consists in applying the CVS construction inside each face of $\mathbf{m}$. More precisely, let $(\mathbf{m}, \ell)$ be an element of $\mathbb{LM}^{(k)}$. For every $i \in \{1, \ldots, k\}$, add an extra vertex v_i inside the face f_i of $\mathbf{m}$. The corners incident to f_i are cyclically ordered in counterclockwise order, and the successor $s(e)$ of corner e is the next corner, in this cyclic order, to have a smaller label, or v_i if no such corner exists (the latter naturally receives label $\ell(v_i) = \min \ell(e) - 1$ where the minimum is taken over the corners incident to f_i).

We let $\mathbf{q}$ be the graph form by the arcs $e \to s(e)$ from corners of $\mathbf{m}$ to their successors. This is a quadrangulation, and it is naturally marked with the vertices $\mathbf{v} = (v_1, \ldots, v_k)$. The delays are simply given by the labels of the v_i's: $d_i = \ell(v_i)$. Of course the vector $\mathbf{d} = (d_1, \ldots, d_k)$ is defined, like ℓ, only up to an additive constant. We let $(\mathbf{q}, \mathbf{v}, \mathbf{d}) = \Psi^{(k)}(\mathbf{m}, \ell)$, and the claim is that $\Phi^{(k)}$ and $\Psi^{(k)}$ are inverse of each other (modulo the rooting convention which, like in the CVS bijection, necessitates an extra sign parameter). Before moving on, we mention the following fact.

Claim. The arcs $e \to s(e)$, for e incident to f_i in $\mathbf{m}$, are exactly the elements of the delayed Voronoi cell $E_i(\mathbf{q}, \mathbf{v}, \mathbf{d})$.

The proof of this claim is easy by inspection: the path of arcs

$$e^- \to s(e)^- \to s(s(e))^- \to \ldots \to v_i$$

ends at v_i, the labels ℓ are decreasing along this path, and it is also the left-most path in $\mathbf{q}$ starting from the oriented arc from e^- to $s(e)^-$.

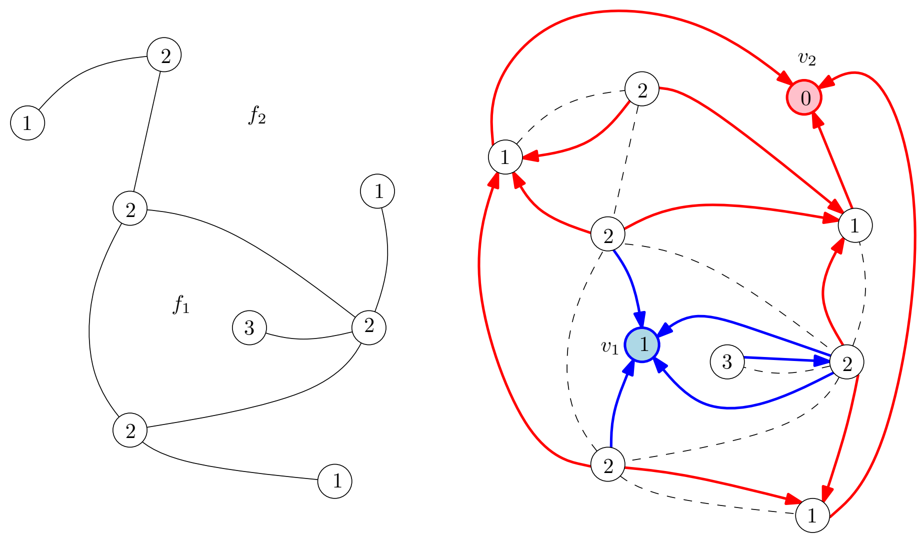

Figure 6.2: Illustration of the two-point bijection

6.2 Uniqueness of typical geodesics

Let us now give some applications of this bijection, starting from a result of uniqueness of geodesics in the Brownian map. Recall that (S, D) denotes the Brownian map, and that μ denotes the uniform measure on S.

Theorem 6.2.1. *Almost surely, for $\mu \otimes \mu$-almost every (x, y) in S, there exists a unique geodesic from x to y.*

We are going to follow an approach developed in [67], and simplified in [18], that uses the 2-pointed bijection described above. See also [55] for a rather different approach.

6.2.1 Discrete considerations

Let us first discuss the discrete setting. Let $(\mathbf{q}, (v_1, v_2), [d_1, d_2])$ be an element of $\mathcal{Q}^{(k)}$. The delay vector is in fact given by the parameter $d_{12} = d_2 - d_1$ which does not depend on the choice of d_1, d_2. This parameter is in $]-d_{\mathbf{q}}(v_1, v_2), d_{\mathbf{q}}(v_1, v_2)[$ and has the same parity as $d_{\mathbf{q}}(v_1, v_2)$: there are $d_{\mathbf{q}}(v_1, v_2) - 1$ possible choices for it. Consider the labeled map $(\mathbf{m}, \ell) = \Phi^{(2)}(\mathbf{q}, \mathbf{v}, \mathbf{d})$. This is a plane map with two faces f_1, f_2, which naturally contain the vertices v_1 and v_2. Such a map contains exactly one cycle, which is formed by the edges and vertices that are incident both to f_1 and f_2. Let $E(f_1, f_2), V(f_1, f_2)$ denote these sets.

Let γ be a geodesic path from v_1 to v_2 in $\mathbf{q}$. This path has to intersect $V(f_1, f_2)$ for topological reasons (it starts in f_1 and ends in f_2, while using only edges of $\mathbf{q}$, that do not cross the edges of $\mathbf{m}$ except at vertices). Let v_0 a vertex that belongs both to γ and $V(f_1, f_2)$. We claim that

$$\ell(v_0) = \min\{\ell(v) : v \in V(f_1, f_2)\} . \tag{6.2}$$

To see this, take any $v \in V(f_1, f_2)$ and let e, e' be two oriented edges incident to v, such that f_1 is incident to e and f_2 is incident to e'. Consider the two chains of consecutive successor arcs $e \to s(e) \to \ldots \to v_1$ and $e' \to s(e') \to \ldots \to v_2$. These have respective lengths $\ell(v) - \ell(v_1)$ and $\ell(v) - \ell(v_2)$. Thus, we can construct a path of length $2\ell(v) - \ell(v_1) - \ell(v_2)$ from v_1 to v_2 passing through v. In particular, by definition of a geodesic path,

$$d_{\mathbf{q}}(v_1, v_2) \leq 2\ell(v) - \ell(v_1) - \ell(v_2) , \qquad \text{for every } v \in V(f_1, f_2) .$$

On the other hand, we clearly have $d_{\mathbf{q}}(u, v) \geq |\ell(u) - \ell(v)|$ for every $u, v \in V(\mathbf{q})$, because the edges of $\mathbf{q}$ link vertices with labels differing by 1 in absolute value. Therefore, since v_0 is on a geodesic path from v_1 to v_2, we have

$$d_{\mathbf{q}}(v_1, v_2) = d_{\mathbf{q}}(v_1, v_0) + d_{\mathbf{q}}(v_0, v_2) \geq 2\ell(v_0) - \ell(v_1) - \ell(v_2) . \tag{6.3}$$

Hence, we conclude that $d_{\mathbf{q}}(v_1, v_2) = 2\ell(v_0) - \ell(v_1) - \ell(v_2)$ and that (6.2) holds. In fact, we have even shown that for every $v \in V(f_1, f_2)$ such that $\ell(v)$ is minimal, there exists a geodesic from v_1 to v_2 that passes through v. Moreover, using (6.1) and the fact vertices of $V(f_1, f_2)$ are by definition incident to both f_1 and f_2, and therefore are incident to edges in $E_1(\mathbf{q}, \mathbf{v}, \mathbf{d})$ and $E_2(\mathbf{q}, \mathbf{v}, \mathbf{d})$ by the claim at the end of Section 6.1, we obtain that for every $v \in V(f_1, f_2)$,

$$d_{\mathbf{q}}(v, v_1) + d_1 = \ell(v) = d_{\mathbf{q}}(v, v_2) + d_2 .$$

If v is also on a geodesic path from v_1 to v_2, we also have $d_{\mathbf{q}}(v_1, v_2) = d_{\mathbf{q}}(v_1, v) + d_{\mathbf{q}}(v, v_2)$, so that

$$d_{\mathbf{q}}(v_1, v) = \frac{d_{\mathbf{q}}(v_1, v_2) + d_{12}}{2} ,$$

and by varying d_{12} the latter quantity can take any integer value in $]0, d_{\mathbf{q}}(v_1, v_2)[$. Putting all this together, we have shown that for every integer $r \in]0, d_{\mathbf{q}}(v_1, v_2)[$

the set of vertices on some geodesic from v_1 to v_2 at distance r from v_1 are exactly those vertices on $V(f_1, f_2)$ with minimal label, for the choice of $d_{12} = 2r - d_{\mathbf{q}}(v_1, v_2)$.

This simple remark gives us a natural approach to the uniqueness of geodesics in random maps: we see that if there is a unique vertex that achieves the minimal label on $V(f_1, f_2)$ for a given choice of d_{12}, then all geodesics from v_1 to v_2 must necessarily pass through this point. Considering the fact that the labels of the vertices on the cycle $V(f_1, f_2)$ form a discrete (periodic) walk, we see that it is plausible that in a uniform random setting, these walks approximate a Brownian bridge, which attains its overall infimum at only one point, indicating that geodesics are asymptotically forced to pass through many imposed points.

6.2.2 The scaling limit of labeled unicycles

The facts discussed above entice to study the asymptotic structure of labels on sets $V(f_1, f_2)$ as discussed above, where the role of $(\mathbf{q}, \mathbf{v}, d_{12})$ is performed by a uniform random quadrangulation Q_n with n faces, v_1, v_2 are independent uniform random vertices in Q_n, and d_{12} is some parameter whose choice will be discussed later. In fact, it will be more convenient for us to use $R = (d_{Q_n}(v_1, v_2) + d_{12})/2$ as the parameter.

To this effect, let us consider first a uniform rooted labeled map (M_n, ℓ_n) in $\mathbb{LM}_n^{(2)}$. We want to consider an appropriate scaling limit for this object. As mentioned above, M_n is a unicycle, that is a map with only one cycle bounding the two faces f_1, f_2. Such a map can be described as follows: start with a cycle of edges with a certain length K_n, which bounds two faces, and to each of the $2K_n$ corners incident to this cycle (K_n for each face), attach a plane tree by its root, in such a way that the total number of edges is n. Distinguish one of the edges as the root. Finally, label this map with an admissible labeling. Comparing with the situation for trees, where typical distances are of order $\sqrt{n}$, it is natural that K_n should be of this order as well. This is what we argue now.

Roughly speaking, the combinatorial information that is contained in the previous description is the data of the length of the cycle, say $K_n = k$, the labels of the vertices along this cycle, and the sequence of $2k$ rooted (labeled) trees $(T_n^{1,1}, \ldots, T_n^{1,k}, T_n^{2,1}, \ldots, T_n^{2,k})$ with a total of $n - k$ edges, the first k of which are grafted on the corners of the cycle incident to f_1 say, the k following ones being grafted in f_2, in this cyclic order. This involves a choice of which corner e_0 of the cycle incident to f_1 we choose to start the grafting, resulting in a symmetry factor $1/k$. The truth is slightly more complicated as there can be further rotational symmetries, for instance if all the trees $T_n^{1,i}, T_n^{2,j}$ turn out to be the same, or more generally if there exists some i such that $(T_n^{1,1+i}, \dots, T_n^{1,k+i}) = (T_n^{1,1}, \dots, T_n^{1,k})$, where the addition $j+i$ should be understood modulo k, and similarly $(T_n^{2,1+i}, \dots, T_n^{2,k+i}) = (T_n^{2,1}, \dots, T_n^{2,k})$. However, we will disregard them, as they become asymptotically unlikely as $n \to \infty$: indeed, with high probability there will be a single tree in the whole family that has maximal height. We will not delve further into the details here.

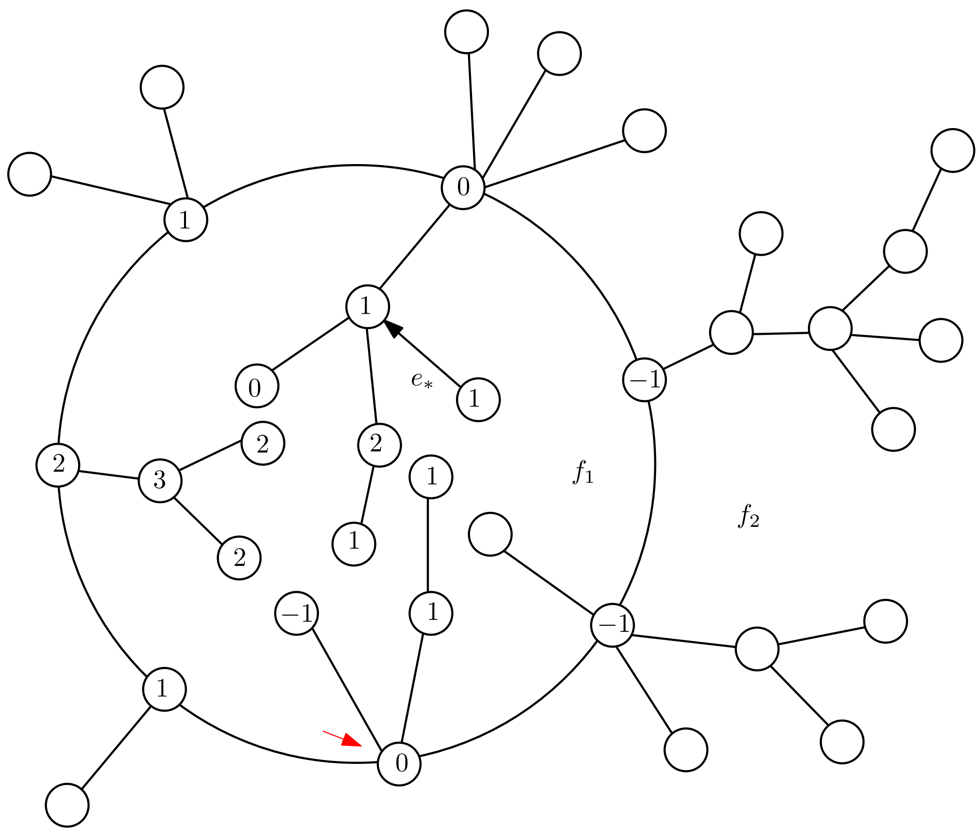

Figure 6.3: A labeled unicycle: labeled trees planted on both sides of a cycle (we only represent some of the labels). The little red arrow represents which tree has been selected as the first to be grafted on the cycle bounding f_1.

Fixing the corner e_0 as above means that we have rooted the cycle in the unicycle. It also allows one to adopt the following convention for ℓ: we take the representative such that $\ell(e_0) = 0$.

Let $\mathrm{M}(k)$ be the number of admissible labelings of the vertices along a rooted cycle of length k: this is the number of sequences $(x_1, \dots, x_k) \in \{-1, 0, 1\}^k$ with sum 0. It can also be written in the form $3^k P(Y_k = 0)$ where Y is a random walk with i.i.d. uniform steps in $\{-1, 0, 1\}$, and from the so-called local limit theorem, we have

$$\sqrt{k} P(Y_k = 0) \underset{n\to\infty}{\longrightarrow} \sqrt{\frac{3}{4\pi}} \,.$$

We also let $\mathrm{F}(m, n)$ be the number of forests with m trees and n edges in total. A well-known combinatorial formula gives

$$\mathrm{F}(m, n) = \frac{m}{2n + m}\binom{2n + m}{n} ,$$

and in particular one can check that $\mathrm{F}(1, n)$ is indeed the n-th Catalan number. The preceding discussion implies that as $n \to \infty$,

$$\begin{aligned} P(K_n = k) &= \frac{2n}{\#\mathbb{LM}_n^{(2)}} \frac{3^{n-k}}{k} \mathrm{M}(k)\mathrm{F}(2k, n - k)(1 + o(1)) \\ &= \frac{2 \cdot 3^n}{\#\mathbb{LM}_n^{(2)}} P(Y_k = 0)\binom{2n}{n - k}(1 + o(1)) \,. \end{aligned} \tag{6.4}$$

where in the first line, the factor $2n$ accounts for the choice of the root edges, and 3^{n-k} accounts for the number of admissible labelings of a forest with $n-k$ edges. Here again, we will be rather sketchy, as the details are a bit tedious. Recall that $\mathbb{LM}_n^{(2)}$ is in 2-to-1 correspondence with $\mathcal{Q}_n^{(2)}$. To understand the asymptotic cardinality of this set, note that an element in this set is given by a bi-pointed quadrangulation, and a choice of a parameter between 0 and the distance between the two distinguished points. As $n \to \infty$, we see that there are typically of order $n^{1/4}$ possible choice for the parameter, and $n^2 \#\mathcal{Q}_n \sim C\, 12^n/n^{1/2}$ choices for the two-pointed quadrangulation, for some $C > 0$, using (2.1) (we will use the same notation C for a positive finite constant, the value of which is allowed to change from line to line). This suggests that

$$\#\mathbb{LM}_n^{(2)} \underset{n\to\infty}{\sim} C \frac{12^n}{n^{1/4}} \,.$$

Recalling that we expect K_n to scale like $\sqrt{n}$, let us take $k = \lfloor x\sqrt{2n} \rfloor$, in (6.4) and perform some routine asymptotic developments, to obtain

$$\sqrt{2n} P(K_n = \lfloor x\sqrt{2n} \rfloor) = C\frac{e^{-x^2/2}}{\sqrt{x}}(1 + o(1)) .$$

Lemma 6.2.2. *Let (M_n, ℓ_n) be a uniform random element of $\mathbb{LM}_n^{(2)}$. Let K_n be the length of the unique cycle of M_n. Then $K_n/\sqrt{2n}$ converges in distribution to a random variable K with density $2^{3/4}e^{-x^2/2}/(\Gamma(1/4)\sqrt{x})$.*

The complete proof of the lemma requires either to identify the constant C as the correct one, or to prove a tightness statement, namely that the probability that $K_n/\sqrt{2n}$ does not belong to $[\varepsilon, 1/\varepsilon]$ can be made arbitrarily small uniformly in n if one chooses ε small enough.

Conditionally on K_n, the labels along the cycle form a random walk $(Y_0, Y_1, \dots, Y_{K_n})$ with uniform steps in $\{-1, 0, 1\}$, and conditioned to come back to 0 at time K_n (the walk is canonically started at the vertex of the cycle of M_n where the first tree was grafted, as discussed above). As is customary, we extend this walk to a function on $[0, K_n]$ by linear interpolation between integer times. Clearly, such a walk will converge after an appropriate rescaling to a Brownian bridge with random duration K, as defined in the preceding Lemma. On top of this, we can also study the scaling limit for labeled trees that are grafted on both sides of the cycle. We will not do this in detail, and content ourselves with stating the following result.

Proposition 6.2.3. *Conditionally on K, on has*

$$\left(\left(\frac{9}{8n}\right)^{1/4} Y_{\sqrt{2n}t}\right)_{0 \le t \le K_n/\sqrt{2n}} \xrightarrow[n\to\infty]{(d)} (\mathcal{Y}_t)_{0 \le t \le K} ,$$

in distribution on the space $\mathcal{W}$ of continuous functions with a finite lifetime[1]*.*

[1] More precisely, let $\mathcal{W} = \bigcup_{\sigma \ge 0} \mathcal{C}([0, \sigma], \mathbb{R})$ be the set of continuous real-valued functions with a compact lifetime interval of the form $[0, \sigma]$. If $w \in \mathcal{W}$, we let $\sigma(w)$ be the length of the lifetime interval. The space $\mathcal{W}$ is separable and complete when endowed with the following distance

$$\mathrm{dist}(w, w') = \sup_{t \ge 0} |w(t \wedge \sigma(w)) - w'(t \wedge \sigma(w'))| + |\sigma(w) - \sigma(w')| .$$

6.2.3 Conclusion of the proof of 6.2.1 (draft)

Using the re-rooting invariance of the Brownian map, Proposition 4.4.2, it suffices to show that if $X = \mathbf{p}(U)$ is a μ-distributed random point in S (here U is a uniform random variable in $[0,1]$), then a.s. there exists a unique geodesic from x_* to X.

Let us assume that this is not the case. Then there exist rational numbers $0 < q < q'$ such that with positive probability, for every $r \in (q, q')$ there exist at least two distinct points $x, x' \in S$, such that $D(x_*, x) = D(x_*, x') = r$ and

$$D(x_*, x) + D(x, X) = D(x_*, x') + D(x', X) = D(x_*, X)\,.$$

We work in restriction to the event of positive probability described above. Let R be an independent random variable, uniform in (q, q'), and let x, x' be as above for $r = R$. Let $t, t' \in [0,1]$ be such that $x = \mathbf{p}(t), x' = \mathbf{p}(t')$.

Going back to the discrete picture, consider bi-pointed approximations of (S, D, x_*, X) by say $(Q_n, (v_*, u^n_{\langle 2nU\rangle}), \mathbf{d})$, where $\mathbf{d}$ is given by

$$d_{12} = 2\left\lfloor \left(\frac{8n}{9}\right)^{1/4} \right\rfloor R - d_{Q_n}(v_*, u^n_{\langle 2nU\rangle})\,.$$

We take also approximations a_n, b_n of the points x, x', say $a_n = \mathbf{p}(\lfloor 2nt \rfloor)$ and similarly for b_n. Finally, we let a'_n, b'_n be the closest points to a_n, b_n on $V(f_1, f_2)$, where the notation is for the map with labeled faces (M_n, ℓ_n) associated with $(Q_n, \mathbf{v}, \mathbf{d})$ as above. Using the considerations above, one sees that a'_n, b'_n must be close to the overall minimum of the label function ℓ_n on $V(f_1, f_2)$, and since the latter converges to a Brownian bridge, there is asymptotically only one such minimum. One concludes that $x = x'$.

6.3 The cut-locus of the Brownian map

A consequence of Theorem 6.2.1, using the re-rooting Proposition 4.4.2, is that almost surely, for μ-almost every $x \in S$, there exists a unique geodesic from x_* to x. Since we already know one geodesic from Section 4.4.1, namely the simple geodesic between x_* and x, we deduce that this is the one and only, for μ-almost every x. In fact, this result can be considerably strengthened, as shown by Le Gall [55].

Theorem 6.3.1. *Almost surely, the only geodesics in (S, D) to x_* are the simple geodesics $\Gamma_t, 0 \leq t \leq 1$.*

This, together with the fact that the simple geodesics to x_* are given by the images by $\mathbf{p}_Z$ of the geodesics to $p_Z(s_*)$ in the tree $\mathcal{T}_Z$, whose root is a leaf, has the following important consequence that the geodesics in the Brownian map tend to coalesce [55].

Proposition 6.3.2. *For every $\varepsilon > 0$, there exists $\eta > 0$ such that for every $x, x' \in S$ such that $D(x, x_*) \wedge D(x', x_*) > \varepsilon$, and every geodesic paths Γ, Γ' from x_* to x, x', it holds that $\gamma(r) = \gamma(r')$ for every $r \in [0, \eta]$.*

This proposition says that there is an "essentially unique" geodesic path leaving the point x_* (more precisely, there is a unique germ of geodesic paths from x_*). In this way, one can argue that (S, D) looks much more like an $\mathbb{R}$-tree than a sphere, but the reason for this phenomenon if of course the very rough structure of the metric, which is also responsible for the (at first surprisingly high) value of the Hausdorff dimension.

We end with a description of the points of (S, D) from where there exist more than one geodesic to x_*. For every $x \in S$, let $\mathrm{Geod}(x \to x_*)$ be the set of geodesic paths from x to x_*.

Theorem 6.3.3. *The following properties hold almost surely.*

(i) *For every $x \in S$ one has $\#\mathrm{Geod}(x \to x_*) \in \{1, 2, 3\}$, and for $k \in \{1, 2, 3\}$, $\#\mathrm{Geod}(x \to x_*) = k$ if and only if $x = \mathbf{p}_{\mathbb{e}}(a)$, where a has degree k in $\mathcal{T}_{\mathbb{e}}$.*

(ii) *The mapping $\mathbf{p}_{\mathbb{e}} : \mathcal{T}_{\mathbb{e}} \to S$ realizes a homeomorphism from the skeleton $\mathrm{Sk}(\mathcal{T}_{\mathbb{e}})$ onto its image, which we denote by $\mathbf{T}_{\mathbb{e}} = \mathbf{p}_{\mathbb{e}}(\mathrm{Sk}(\mathcal{T}_{\mathbb{e}}))$. The latter is the* cut-locus *of (S, D) with respect to the point x_*, i.e. the set of points from which there exist at least two distinct geodesics to x_*.*

(iii) *The mapping $\mathbf{p}_Z : \mathcal{T}_Z \to S$ realizes a homeomorphism from the skeleton $\mathrm{Sk}(\mathcal{T}_Z)$ onto its image, which we denote by $\mathbf{T}_Z$. The latter is the union of the relative interiors of all geodesic segments started from x_*, i.e. the union of all geodesic segments from x_* with their endpoints removed.*

Proof. We first prove the statement about the restrictions of $\mathbf{p}_{\mathbb{e}}, \mathbf{p}_Z$ to $\mathrm{Sk}(\mathcal{T}_{\mathbb{e}}), \mathrm{Sk}(\mathcal{T}_Z)$ realizing homeomorphisms with their respective images. By definition, these two maps are continuous. Clearly, $\mathbf{p}_{\mathbb{e}}$ is injective when restricted to $\mathrm{Sk}(\mathcal{T}_{\mathbb{e}})$ since two points of the skeleton of $\mathcal{T}_{\mathbb{e}}$ are not identified in (S, D). One has to check that if $\mathbf{p}_{\mathbb{e}}(a_n)$ converges to $\mathbf{p}_{\mathbb{e}}(a)$ with $a_n, a \in$

$\mathrm{Sk}(\mathcal{T}_{\mathbb{e}})$, then $a_n \to a$ in $\mathcal{T}_{\mathbb{e}}$. By up to extraction one can assume that $a_n \to a'$ in $\mathcal{T}_{\mathbb{e}}$, so that $\mathbf{p}_{\mathbb{e}}(a) = \mathbf{p}_{\mathbb{e}}(a')$. Since a is a point of the skeleton, it is not identified with any other point in S, so $a = a'$.

The argument for the tree $\mathcal{T}_Z$ is similar, and one deduces (iii) by Theorem 6.3.1.

One obtains (ii) by noticing that we can start as many distinct simple geodesics from a point $x \in \mathbf{T}_{\mathbb{e}}$ as the degree of the corresponding point a in $\mathcal{T}_{\mathbb{e}}$: these correspond to the paths Γ_t where t is such that $p_{\mathbb{e}}(t) = a$. The fact that these geodesics are all distinct amounts to the fact that they correspond to different branches in the tree $\mathcal{T}_Z$. $\square$

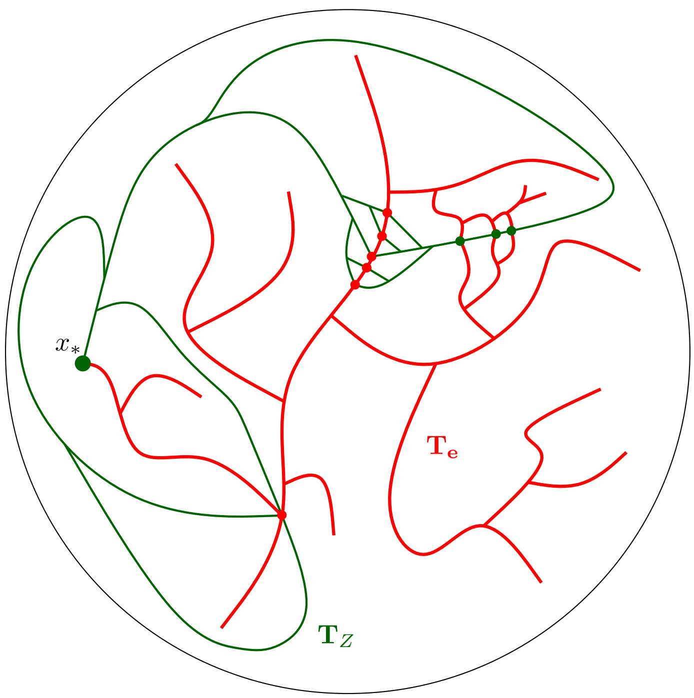

Figure 6.4: The cut-locus of the Brownian map. In green are the geodesics to x_*, in red is the tree $\mathbf{T}_{\mathbb{e}}$. In a small portion at the top right, we emphasize how the trees $\mathbf{T}_{\mathbb{e}}$ and $\mathbf{T}_Z$ are intertwined together, with $\mathbf{T}_{\mathbb{e}} \subset \mathbf{p}_Z(\mathrm{Lf}(\mathcal{T}_Z))$ and $\mathbf{T}_Z \subset \mathbf{p}_{\mathbb{e}}(\mathrm{Lf}(\mathcal{T}_{\mathbb{e}}))$.

It is quite striking that the Brownian tree and its partner tree $\mathcal{T}_Z$, or more precisely the skeletons thereof, can be seen as properly embedded in the Brownian map, with such concrete geometric descriptions (this is one of the reasons why we said earlier that $\mathcal{T}_{\mathbb{e}}$ is naturally immersed in (S, D)). We encourage the reader to go back to the description of the CVS bijection, and to convince her/himself that the role of T_n is in fact similar within the pointed quadrangulation (Q_n, v_*).

Notes for Chapter 6

The multi-pointed bijection for $k = 3$ was notably used in Bouttier and Guitter [26] to derive the three-point function for random quadrangulations, that is, the asymptotics of the distance statistics for the triple

$$n^{-1/4}(d_{Q_n}(v_1, v_2), d_{Q_n}(v_2, v_3), d_{Q_n}(v_1, v_3)),$$

where v_1, v_2, v_3 are three independent uniform vertices in Q_n. The idea is to find a nice bijection for the family of three-pointed quadrangulations $(\mathbf{q}, (v_1, v_2, v_3))$ such that $d_{\mathbf{q}}(v_i, v_j) = a_{ij}$, where $(a_{ij}, 1 \leq i, j \leq 3)$ is a fixed matrix of integers that is null on the diagonal and whose coefficients satisfy the triangle inequality. The idea is that such coefficients can be uniquely written as $a_{ij} = u_i + u_j$ for some $(u_1, u_2, u_3) \in \mathbb{Z}_+^3$. If one chooses the delay vector $[d_1, d_2, d_3]$ so that $d_i = -u_i$, then (excepting asymptotically negligible degenerate cases that we ignore here) the labeled map $(\mathbf{m}, \ell)$ with 3 faces f_1, f_2, f_3 that one gets from the 3-pointed bijection from $(\mathbf{q}, \mathbf{v}, \mathbf{d})$ are those satisfying the following properties: for every $i \neq j \in \{1, 2, 3\}$

- there exists at least one vertex v in the set $V(f_i, f_j)$ of vertices incident to both f_i and f_j is not empty, such that $\ell(v) = 0$,
- the label function ℓ is non-negative on $V(f_i, f_j)$,
- the minimal label on f_i is $-u_i + 1$.

Note that the label function is well-defined here (not up to an additive constant), since we chose a particular value for the components of the delay vector. The enumeration of labeled maps described above can be performed (this requires splitting these maps into simpler pieces) and yields the wanted 3-point function after a careful Laplace/saddle point analysis.

Chapter 7

Uniqueness of the Brownian map

The aim of this chapter is to give some elements of the proof of Theorem 4.3.2, that we recall now: it says that D, seen as a pseudo-metric on $[0,1]$, is identified a.s. with the quotient pseudo-metric D^*, which is the largest pseudo-metric d on $[0,1]$ such that

- $\{d_{\mathbb{e}} = 0\} \subset \{d = 0\}$, and
- $d \leq d_Z$.

Recall also that D^* is given by the explicit formula

$$D^*(s,t) = \inf \left\{ \sum_{i=1}^{k} d_Z(s_i, t_i) \right\} , \tag{7.1}$$

the infimum being taken over all $k \geq 1$ and $s = s_1, t_1, s_2, t_2 \ldots, s_k, t_k = t$ such that $d_{\mathbb{e}}(t_i, s_{i+1}) = 0$ for $1 \leq i \leq k-1$. Recalling that $d_Z(s,t)$ is the combined length of the disjoint parts of the simple geodesics Γ_s, Γ_t to x_*, we arrive at the rather surprising fact that

> *geodesic paths in the Brownian map are well-approximated by paths that are piecewise parts of a geodesic to or away from x_*.*

This can look surprising in many respects. Indeed, by re-rooting invariance, the same is true for a μ-chosen point X instead of x_*. So at this point it seems like any geodesic path in the Brownian map is locally a geodesic to, or away from any typical point in the space! This apparent paradox can be explained by the fact that the Brownian map is like a mountain area, of which typical points (a set of full measure for μ) is made of peaks, and geodesics follow deep valleys: since these valleys are so deep and rare, locally, one has to follow them regardless of the peak one is aiming at.

We have already mentioned the coalescence of geodesics, which is a way to see that there is essentially only one geodesic path that leaves a typical point (e.g. the root x_*). This goes in the right direction, however, we have to show that this property also holds well inside geodesic paths, and the difficulty is that points on geodesics are not "typical".

One of the difficulties in handling the Brownian map is that it is defined in terms of the various functions $d_{\mathbb{e}}, d_Z, D, D^*$ that can be seen as defined either on $[0,1], \mathcal{T}_{\mathbb{e}}, \mathcal{T}_Z, S$. In particular, in the sequel we let d_Z denote either the pseudo-metric on $[0,1]$, the associated distance function on $\mathcal{T}_Z$, or the function $d_Z(x,y) = \inf\{d_Z(s,t) : \mathbf{p}(s) = x, \mathbf{p}(t) = y\}$ defined on S (not a pseudo-metric!). In this way, (7.1) becomes, if we now see D as a distance on S,

$$D^*(x,y) = \inf\left\{\sum_{i=1}^{k} d_Z(x_{i-1}, x_i)\right\}, \qquad x, y \in S,$$

the infimum being over all $x = x_0, x_1, \ldots, x_k = y$ in S.

7.1 Bad points on geodesics

Let x_1, x_2 be two independent μ-distributed points in the Brownian map, say $x_1 = \mathbf{p}(U_1), x_2 = \mathbf{p}(U_2)$ where U_1, U_2 are independent uniform random variables in $[0,1]$. We know from Chapter 6 that there is a.s. a unique geodesic path γ from x_1 to x_2. We identify this path with its image, i.e. the unique geodesic segment between x_1 and x_2.

Since we want to show that most of γ is made of pieces of geodesics going to x_*, we will declare a point $x \in \gamma$ to be *bad* if for any geodesic segment γ' between x and x_*, one has $\gamma \cap \gamma' = \{x\}$. We let B be the set of bad points of γ.

We want to show that bad points are rare in a certain sense. It turns out that the right notion here is that of Minkowski dimension. The key lemma is the following.

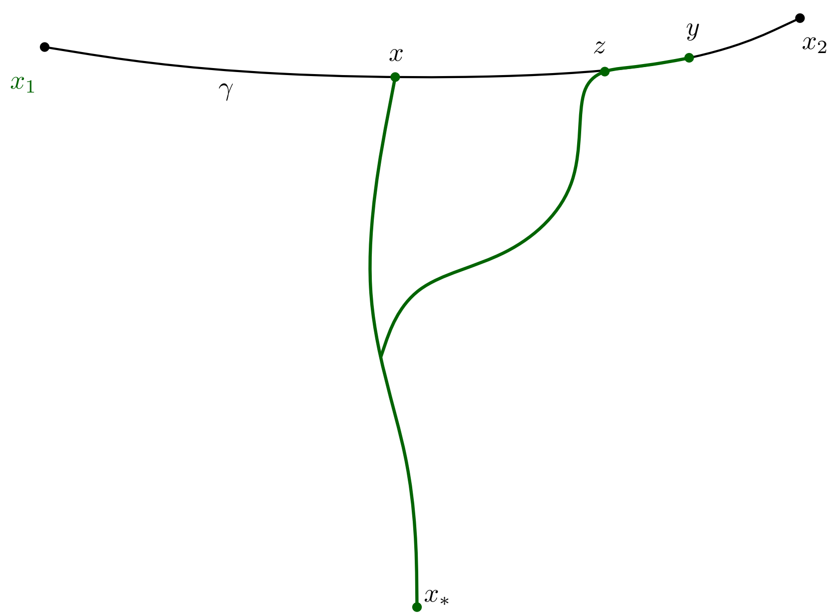

Figure 7.1: A bad point x, and a good point y, with highlighted geodesics from x and y to x_*. Note that the last point of contact z between γ and the geodesic from y to x_* has to be a bad point (why?)

Lemma 7.1.1. *There exists $\alpha \in (0,1)$ for which the following holds almost surely. There exists a (random) $\varepsilon_0 > 0$ such that for every $\varepsilon \in (0, \varepsilon_0)$, the set B can be covered with at most $\varepsilon^{-\alpha}$ balls of radius ε in (S, D).*

In particular, the Hausdorff dimension of B in (S, D) is strictly less than 1 a.s. In this sense B is sufficiently rare, since clearly the Hausdorff or Minkowski dimension of γ is 1 (γ being isometric to a real segment).

This lemma is what takes most of the efforts in the proof of Theorem 4.3.2. Let us now explain why it is essentially sufficient to conclude. The second important remark is to recall that the distances D and D^* satisfy $D \le D^*$, so that it only remains to prove the reverse inequality. Even though this is not achievable directly, it turns out that the distance D^* and the *snowflaked*[1] version of D, that is the distance D^β for some $\beta \in (0,1)$, are equivalent.

Lemma 7.1.2. *For every $\beta \in (0,1)$, there exists a (random) constant $C = C_\beta \in (0, \infty)$ such that for every $s, t \in [0,1]$,*

$$D^*(s,t) \le C\, D(s,t)^\beta \,.$$

We will soon explain the reason for this lemma, but let us first conclude the proof of Theorem 4.3.2 assuming the previous two lemmas. Let $\alpha \in (0,1)$ be as in Lemma 7.1.1. Fix $\beta \in (\alpha, 1)$ and let $C = C_\beta$ be the random constant appearing in Lemma 7.1.2. Fix $\varepsilon \in (0, \varepsilon_0)$, where ε_0 is as in Lemma 7.1.1, and consider a covering $\mathcal{C}$ of B by open balls with centers $b_1, \ldots, b_N$ and radius ε, where $N = N(\varepsilon) \le \varepsilon^{-\alpha}$. The complement $\gamma \setminus \mathcal{C}$ is free of bad points. It is made of a finite collection of geodesic sub-segments of γ, which we write $[\![a_i, b_i]\!]_S, i \in I$. The notation is unambiguous since γ itself was the unique geodesic segment between its extremities.

We claim that $D^*(a_i, b_i) = D(a_i, b_i)$ for every $i \in I$. To see this, let us consider one of these segments, let us write it in the form $[\![a, b]\!]_S$ for simplicity. We claim that there exists a unique $c \in [\![a, b]\!]_S$ such that a is on a geodesic from c to x_* and b is on a geodesic from c to x_*. This is relatively simple to check: for any c' in $[\![a, b]\!]_S$ there exists by definition a geodesic γ' from c' to x_* that intersects γ at a point $c'' \neq c$: let us choose γ' and c'' so that c'' is at maximal distance from c'. Then necessarily, a or b must belong to $[\![c', c'']\!]_S$,

[1]If (X, d) is a metric space, and $\beta \in (0,1)$, then (X, d^β) is still a metric space, called the β-snowflaking of (X, d). Although this might seem harmless, the snowflaking destroys many regularity properties: in particular, snowflaked spaces are never geodesic spaces, excepting the trivial case of a one-point space.

otherwise, c'' would be a bad point in $[\![a,b]\!]_S$. For simple metric reasons, it is easy to see that the set of c' such that $a \in [\![c'',c']\!]_S$ is a sub-segment of $[\![a,b]\!]_S$ containing c', and it suffices to let c be the right-extremity of this segment.

We see that $D(a,b) = D(a,c) + D(c,b)$ since c is on a geodesic from a to b, and also that $D(a,c) = d_Z(a,c)$ and $D(c,b) = d_Z(c,b)$ since a and b are on geodesics from c to x_*, which must be simple geodesics. Hence, we deduce that $D(a,b) = D^*(a,b)$, as wanted.

By the triangle inequality, we have proved

$$\begin{aligned} D^*(x_1,x_2) &\leq \sum_{i\in I} D^*(a_i,b_i) + N \sup_{x,y\in S, D(x,y)\leq 2\varepsilon} D^*(x,y) \\ &= \sum_{i\in I} D(a_i,b_i) + N \sup_{x,y\in S, D(x,y)\leq 2\varepsilon} D^*(x,y)\,, \end{aligned}$$

where the remainder comes from the fact that we have removed N balls of D-radius ε (hence diameter 2ε) from γ. Since $\bigcup_{i\in I}[\![a_i,b_i]\!]_S \subset \gamma$ is a disjoint union and γ is a geodesic, we have that the first sum is less than $D(x_1,x_2)$, and by Lemma 7.1.2 and the fact that $N \leq \varepsilon^{-\alpha}$, we have

$$D^*(x_1,x_2) \leq D(x_1,x_2) + C\varepsilon^{-\alpha}(2\varepsilon)^\beta\,.$$

By our choice of β, letting $\varepsilon \to 0$ gives $D^*(x_1,x_2) \leq D(x_1,x_2)$, and hence these quantities are equal since $D \leq D^*$. Since x_1,x_2 are μ-distributed points and μ clearly has full support, we deduce that $D = D^*$ by a density argument.

7.2 The snowflaking lemma

Let us give some ideas of the proof of Lemma 7.1.2. This relies on volume estimates for balls in the metrics D and D^*. We already saw in Lemma 4.4.4 that for every $\delta \in (0,1)$, it holds that

$$\sup_{x\in S} \mu(B_D(x,r)) \leq C\,r^{4-\delta}\,, \tag{7.2}$$

for some (random) constant $C \in (0,\infty)$ and every $r > 0$.

Somehow, the converse bound holds for the distance D^*, showing that balls for this distance cannot be too small.

Lemma 7.2.1. *Let $\eta \in (0,1)$ be fixed. Then almost surely, there exist random $c \in (0,\infty)$ and $r_0 > 0$ such that for every $r \in [0, r_0]$ and every $x \in S$, one has*

$$\mu(B_{D^*}(x,r)) \geq c\, r^{4+\eta}\,.$$

Proof. Let $x = \mathbf{p}(s) \in S$. Then, if $d_Z(s,t) < r$ we have $D^*(s,t) < r$ as well by definition of D^*, otherwise said

$$\mathbf{p}(\{t \in [0,1] : d_Z(s,t) < r\}) \subseteq B_{D^*}(x,r)\,,$$

and thus

$$\mu(B_{D^*}(x,r)) \geq \mathrm{Leb}(\{t \in [0,1] : d_Z(s,t) < r\})\,.$$

Since Z is a.s. Hölder-continuous with exponent $1/(4+\eta)$, there exists a finite constant C such that $d_Z(s,t) \leq C|t-s|^{1/(4+\eta)}$. Hence,

$$\mu(B_{D^*}(x,r)) \geq (r/C)^{4+\eta} \wedge 1\,,$$

independently on x and s, and this yields the result. □

We can now prove the snowflaking lemma 7.1.2. Fix $\alpha \in (0,1)$, and let $\eta = (1-\alpha)/3$. We argue by contradiction, assuming the existence of sequences $x_n, y_n, n \geq 1$ of points in S such that $D(x_n,y_n) \to 0$ as $n \to \infty$ and $D^*(x_n,y_n) > D(x_n,y_n)^\alpha$. Let γ_n be a geodesic path from x_n to y_n for the metric D.

Since $\{D = 0\} = \{D^* = 0\}$ from Theorem 5.1.1 and $D \leq D^*$ by definition, we have that the identity map $(S,D^*) \to (S,D)$ is continuous, hence a homeomorphism by compactness of (S,D^*) (the latter is itself a simple consequence of the fact that $D^* \leq d_Z$.)

In particular, γ_n is a continuous path in the space (S,D^*). Clearly, it is possible to find at least $D^*(x_n,y_n)/2D(x_n,y_n)$ points $z_1,\ldots,z_k$ in the image of γ_n such that $D^*(z_i,z_j) \geq D(x_n,y_n)$ for every $i \neq j$ with $i,j \in \{1,2,\ldots,k\}$. In particular, the D^*-balls of radii $D(x_n,y_n)/2$ centered at the points $z_i, 1 \leq i \leq k$ are pairwise disjoint. This shows that the $D(x_n,y_n)$-neighborhood of the image of γ_n in (S,D^*) has μ-measure at least

$$ckD(x_n,y_n)^{4+\eta} \geq cD^*(x_n,y_n)D(x_n,y_n)^{3+\eta}/2 > cD(x_n,y_n)^{3+\eta+\alpha}\,,$$

for every n large enough, by Lemma 7.2.1 and the fact that $D(x_n,y_n) \to 0$. On the other hand, the $D(x_n,y_n)$-neighborhood of γ_n in (S,D^*) is included in the $D(x_n, y_n)$-neighborhood of γ_n in (S, D), because $D \le D^*$. The latter is included in $B_D(x_n, 2D(x_n, y_n))$, which by (7.2) has volume at most $C(2D(x_n, y_n))^{4-\eta}$. Finally, we have proved that there exists a random constant C such that for every n large enough,

$$D(x_n, y_n)^{3+\eta+\alpha} \le C\, D(x_n, y_n)^{4-\eta}\,.$$

Since $D(x_n, y_n) \to 0$, we have a contradiction since $\alpha + 2\eta - 1 < 0$.

7.3 The covering lemma (draft)

It remains to explain how to prove Lemma 7.1.1. To this end, the idea is to try understanding the behavior of geodesics near the bad points B of the geodesic segment between x_1 and x_2. We will be very sketchy here, giving only a couple of the important ideas. From now on, all mentions of distances will concern the distance D, and not D^*.

The first idea is to sample yet another μ-chosen point $x_0 = \mathbf{p}(U_0)$ in (S, D), and estimate the probability that it falls within (D-)distance ε of a bad point B. Since we know that balls in (S, D) with radius r have volumes of order $r^{4+o(1)}$, we can hope to get the fractal dimension of B by a codimension estimate: if $P(D(x_0, B) < \varepsilon) \le \varepsilon^{4-\alpha}$ for small $\varepsilon > 0$, then one can expect that the dimension of B is at most α.

To be more precise, the discussion of Lemma 7.2.1 shows that for every point $x \in S$ and any geodesic to Γ from x to x_*, which is thus say the simple geodesic Γ_t for some t such that $\mathbf{p}(t) = x$, setting $F(t, \varepsilon) = \mathbf{p}([(t-h) \vee 0, (t+h) \wedge 1])$ where $h = \varepsilon^{4+\eta}$, it holds that for every ε small enough, for every $y \in F(t, \varepsilon)$, all geodesics from y to x_* coalesce with Γ at a distance at most $\varepsilon/2$ from both x and y. In particular, $F(t, \varepsilon)$ is included in $B_D(x, \varepsilon)$, but $F(t, \varepsilon)$ need not be a neighborhood of x, we generally picture it rather as a "fan" with apex at x.

Let $U_{(1)}, U_{(2)}, \ldots$ be independent uniforms. With overwhelming probability, the intervals of width $\varepsilon^{4+\eta}$ around $U_{(i)}$ for $1 \le i \le \varepsilon^{-4-2\eta} = N$ cover all $[0, 1]$. Therefore, with overwhelming probability we see that the number $\mathcal{N}_\varepsilon$ of ε-balls necessary to cover B is at most $\sum_{i=1}^N \mathbf{1}_{\{\mathbf{p}(U_{(i)}) \in \bigcup_{y \in B} F(y, \varepsilon)\}}$. Indeed, for every $t \in [0, 1]$, the interval of radius $\varepsilon^{4+\delta}$ around t must contain one of the points $U_{(i)}, 1 \le i \le N$, and if $\mathbf{p}(t) = y$ we get that $y \in B(\mathbf{p}(U_{(i)}), \varepsilon)$.

This means that the probability that $\mathcal{N}_\varepsilon > \varepsilon^{-\alpha}$ can be bounded from

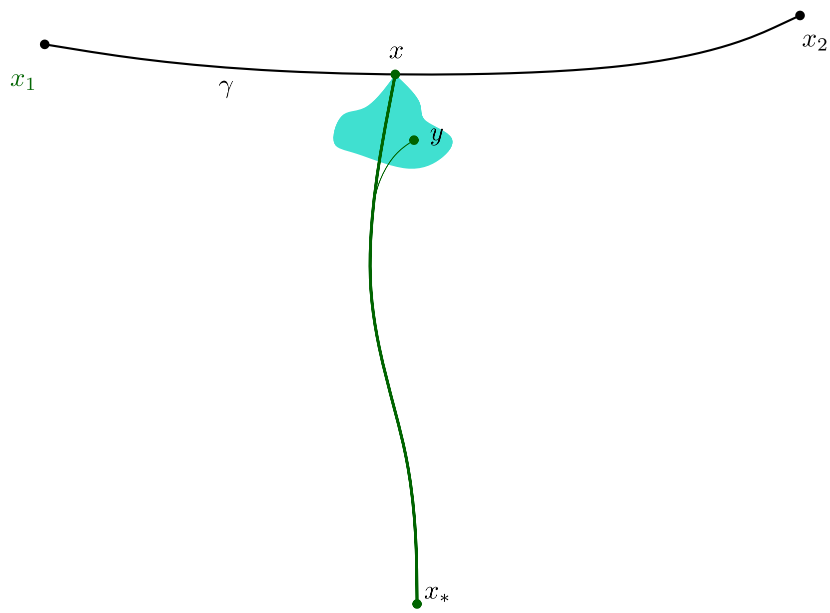

Figure 7.2: Illustration of the set $F(t,\varepsilon)$ from which geodesics to x_* coalesce quickly with Γ_t.

above by

$$\varepsilon^\alpha N\, P\left(x_0 \in \bigcup_{y\in B} F(y,\varepsilon)\right) .$$

Now, if we take the leap of faith that for every $y \in F(t,\varepsilon)$, where $\mathbf{p}(t) \in B$, the geodesics from y to x_1 and x_2 also coalesce with γ within an ε-neighborhood of y, we obtain, using the fact that x is a bad point, that the probability of the last event is less than $\varepsilon^{\alpha-4-2\eta}P(\mathcal{A}(\varepsilon))$, where $\mathcal{A}(\varepsilon)$ is the event that

- γ intersects $B(x_0,\varepsilon)$
- the (unique geodesics) from x_0 to x_1, x_2, x_* do not intersect outside of $B(x_0,\varepsilon)$.

It remains to show that for some $\delta > 0$ and some $C > 0$

$$P(\mathcal{A}(\varepsilon)) < C\varepsilon^{3+\delta} \tag{7.3}$$

for every $\varepsilon > 0$: in this case, choosing $\eta = \delta/4$ gives that $P(\mathcal{N}_\varepsilon > \varepsilon^{-\alpha}) \leq C\varepsilon^{\alpha-1+\eta/2}$, and so if we choose $\alpha = 1-\eta/4 < 1$ we obtain a bound $C\varepsilon^{\eta/4}$.

Taking $\varepsilon = 2^{-k}$ and applying the Borel-Cantelli Lemma entails that $\mathcal{N}_\varepsilon < C'\varepsilon^{-\alpha}$ for every ε small, for some $C' > 0$, and up to taking α slightly larger we can get rid of the factor C' up to taking ε small enough.

From a combinatorial point of view, proving (7.3) really amounts to a counting problem. Namely, going back to the discrete picture, we have to count the number of quadrangulations with n faces and with 4 distinguished vertices v_*, v_1, v_2, v_0 such that there exists a geodesic from v_1 to v_2 that visits the ball of radius $\varepsilon(8n/9)^{1/4}$ around v_0, and such that every triple geodesics from v_0 to v_1, v_2, v_* respectively are disjoint outside this ball. This requires quite a bit of work and can be done using the multi-pointed bijection, setting sources at v_0, v_1, v_2, v_* and delays $d_0 = \lfloor -(2+U)\varepsilon(8n/9)^{1/4} \rfloor$ (where U is uniform on $[0, 1]$),

$$d_1 = -d_{\mathbf{q}}(v_1, v_0), \quad d_2 = -d_{\mathbf{q}}(v_2, v_0), \quad d_* = -d_{\mathbf{q}}(v_*, v_0)\,.$$

The idea is that, with this choice, the “delayed Voronoi cell” of v_0 has the time to swallow the ball of radius $\varepsilon(8n/9)^{1/4}$ around v_0 before it is reached by the cells growing from the three other vertices. Choosing d_0 random allows to keep some flexibility in the counting problems to come.

Notes for Chapter 7

Theorem 4.3.2, which legitimates calling the space (S, D) “the” Brownian map, since it is the well-defined limit of rescaled quadrangulations without having to take a subsequence, was proved at about the same time in the two independent works by Le Gall [56] and the author [68]. The approach by Le Gall does not make use of the multi-pointed bijection, and exploits various symmetries of the Brownian map. An important idea of [56] is also the introduction of a decomposition of the Brownian map into elementary parts that are made of pieces of Brownian map with a geodesic boundary. The study of how the geodesics penetrate in these submaps allows one to conclude that they have the required structure, that is, they tend to go to or away from x_*.

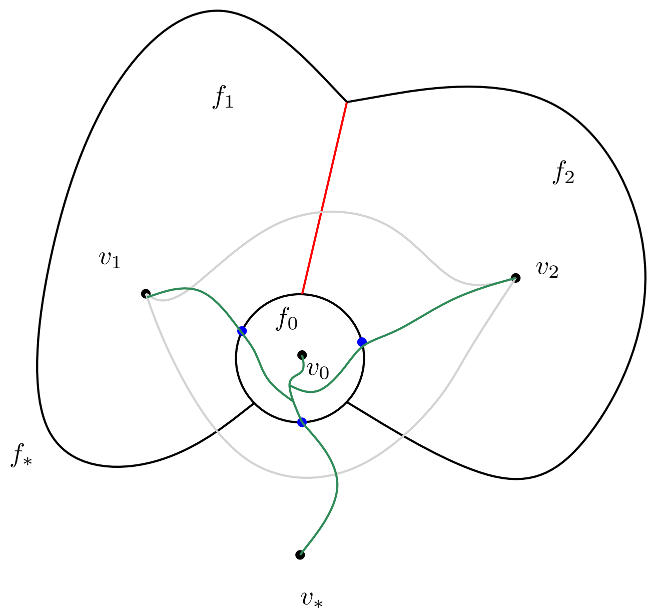

Figure 7.3: A typical configuration of the 4-point bijection associated with a configuration of $\mathcal{A}(\varepsilon)$. In this picture, the three geodesics from v_0 to v_*, v_1, v_2 separate before exiting f_0. Also, a geodesic from v_1 to v_2 should traverse f_0, which is expressed by the condition that grey paths not passing through f_0 cannot be geodesics.

Chapter 8

Universality of the Brownian map

8.1 Introduction

We have focused so far on the sole model of uniform random plane quadrangulations with a given size. The reason is of course that the CVS bijection is specific to this context, but one can legitimately ask whether similar limit results hold for more general models. It turns out that the CVS bijection admits a natural extension to general maps, called the Bouttier-Di Francesco-Guitter (BDG) bijection [25]. This is at the cost of using more elaborate trees to encode maps: only in the quadrangulation case do these trees simplify to re-obtain the CVS bijection. Hence, with this approach, the case of quadrangulations is really the simplest combinatorially.

It was observed in [62], then in [65, 80, 69] that the results of Chassaing-Schaeffer [31] given in particular in Theorem 4.1.2 can be vastly extended to models of random maps with “Boltzmann laws”, using the BDG bijection. The latter models include the case of uniform maps in the set of p-angulations with n faces, where a p-angulation is a map with only faces of degree p, where $p \geq 3$. A large class of these models, including uniform p-angulations with n faces for $p \in \{3\} \cup \{4, 6, 8, 10, \ldots\}$, was then shown by Le Gall [56] to converge after appropriate rescaling to the Brownian map. This is a sign of the so-called *universality* of the Brownian map, namely that the latter is the scaling limit for a large class of natural models of random maps. The argument of Le Gall relies on a beautiful and simple idea using the symmetry of the Brownian map under re-rooting, and that we will describe here.

For technical simplicity, we will only discuss the BDG bijection for *bipartite maps*, see Section 2.3. We let $\mathcal{B}$ be the set of rooted bipartite plane maps. Non-bipartite maps are technically harder, and only partial results are known so far for this class. We will further simplify the exposition by presenting the scaling limit result for $2p$-angulations only. The main goal of this chapter is thus to prove the following theorem.

Theorem 8.1.1. *Fix an integer $p \geq 2$, and for every $n \geq 1$, let M_n be a uniform random rooted $2p$-angulation with n faces. Then*

$$\left(V(M_n), \left(\frac{9}{4p(p-1)n}\right)^{1/4} d_{M_n}\right) \xrightarrow[n\to\infty]{(d)} (S, D)\,,$$

for the Gromov-Hausdorff topology, where (S, D) is the Brownian map.

8.2 The Bouttier-Di Francesco-Guitter bijection

The BDG bijection generalizes the CVS bijection, and in particular, it associates a labeled tree with a rooted planar map that is also pointed. Here, we will describe the family of trees that are involved, and how to get a map starting from such trees.

Let $\mathbf{t}$ be a rooted plane tree, with root e_0 and root vertex u_0. We let $V_\circ(\mathbf{t}), V_\bullet(\mathbf{t})$ be respectively the set of vertices of $\mathbf{t}$ that are at even (resp. odd) distance from the root vertex. In doing so, we can see $\mathbf{t}$ as the genealogical tree of a family in which individuals have two possible types, say "white, $\circ$" or "black, $\bullet$", and where these two types alternate over generations. This point of view will prove useful later on.

Given a tree $\mathbf{t}$, a mobile-admissible labeling is a function $\ell : V_\circ(\mathbf{t}) \to \mathbb{Z}$ such that $\ell(u_0) = 0$, and with the following property. For every $u \in V_\bullet(\mathbf{t})$, if $u_{(0)}, u_{(1)}, \ldots, u_{(p-1)} \in V_\circ(\mathbf{t})$ are the vertices that are adjacent to u, ordered cyclically around u in clockwise order, and with the convention that $u_{(0)} = \neg u$ is the parent of u, then

$$\ell(u_{(i)}) - \ell(u_{(i-1)}) \geq -1\,, \qquad \text{for every } i \in \{1, \ldots, p\}\,, \tag{8.1}$$

with the convention that $u_{(p)} = u_{(0)}$.

Definition 8.2.1. *A labeled mobile is a pair* $(\mathbf{t}, \ell)$ *formed by a rooted tree* $\mathbf{t}$ *and a mobile-admissible labeling* ℓ.

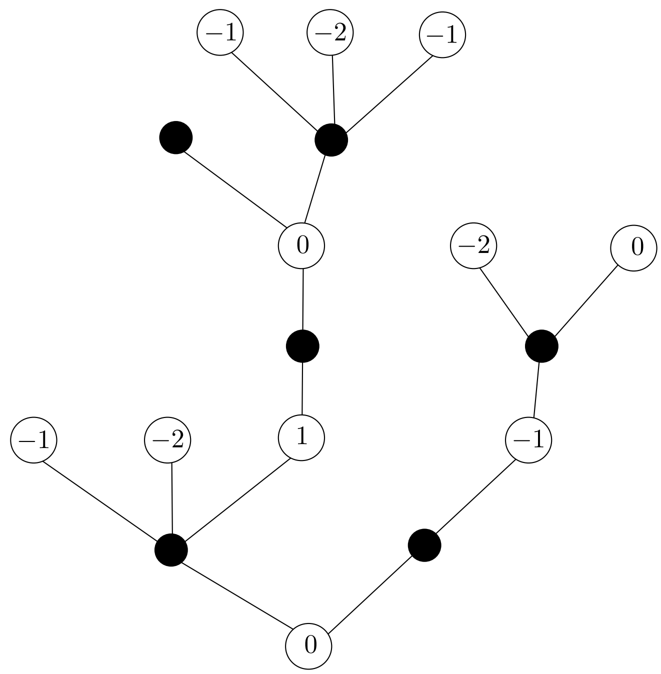

Figure 8.1: A labeled mobile

Note that for a given tree $\mathbf{t}$, there is only a finite number of mobile-admissible labelings. Indeed, with the same notation as above, the sequence $(\ell(u_{(i)}) - \ell(u_{(i-1)})), 1 \leq i \leq p)$, is an element of $\{-1, 0, 1, 2, 3, \ldots\}$ with total sum 0. It is a simple exercise to see that the number of such sequences is $\binom{2p-1}{p}$. Together with the constraint that the root vertex has label 0, it is then simple to see by induction that the number of possible mobile-admissible labelings of a given rooted tree $\mathbf{t}$ is

$$\prod_{u \in V_\bullet(\mathbf{t})} \binom{2\deg(u) - 1}{\deg(u)} = \prod_{u \in V_\bullet(\mathbf{t})} \binom{2k_u + 1}{k_u}.$$

where $k_u = \deg(u) - 1$ is the number of children of u in $\mathbf{t}$.

In particular, if $k_u = 1$ for every $u \in V_\bullet(\mathbf{t})$, we see that there are 3^n possible mobile-admissible labelings, where n is the cardinality of $V_\bullet(\mathbf{t})$: every

$v \in V_\circ(\mathbf{t}) \setminus \{u_0\}$ receives label $\ell - 1, \ell$ or $\ell + 1$ where ℓ is the label of the grandparent of v. We see that removing the vertices of $V_\bullet(\mathbf{t})$ yields an element of $\mathbb{T}_n$, and modulo this modification, the construction below boils down to the CVS bijection in this particular case.

The construction of the BDG bijection, going from a labeled mobile $(\mathbf{t}, \ell)$ to a map, now goes almost exactly as the CVS bijection. Let $e_0, e_1, \dots, e_{n-1}$ be the corners of $\mathbf{t}$ incident to *white* vertices: if $e'_0 = e_0, e'_1, \dots e'_{2n-1}$ is the usual contour exploration of $\mathbf{t}$, then $e_i = e'_{2i}, 0 \leq i \leq n-1$. We draw an arc from each white corner e_i to its successor $s(e_i)$, which is the next available white corner in (cyclic) contour order around $\mathbf{t}$ with label $\ell(e_i) - 1$. If there is no such vertex, we draw an arc from e_i to a distinguished vertex v_* added outside the support of $\mathbf{t}$. We root the resulting map $\mathbf{m}$, formed by the arcs thus drawn, at the arc from e_0 to $s(e_0)$, oriented in this way or the other depending on an external choice $\epsilon \in \{-1, 1\}$. The map $\mathbf{m}$ is naturally pointed at v_*, and we let $(\mathbf{m}, v_*) = \Phi((\mathbf{t}, \ell), \epsilon)$. See Figure 8.2 for an example.

An important aspect of the bijection is that, like the CVS bijection, various properties of $\mathbf{m}$ are easily readable from $(\mathbf{t}, \ell)$. For instance, every face of $\mathbf{m}$ with degree $2k$ corresponds to exactly one vertex of $V_\bullet(\mathbf{t})$ with degree k (i.e. with $k-1$ children). This can be seen by a simple adaptation of the argument of Section 2.3. Again, we see that all faces are quadrangles if and only if $k_u = 2$ for every $u \in V_\bullet(\mathbf{t})$. Moreover, the set of vertices of $\mathbf{m}$ except v_* is exactly $V_\circ(\mathbf{t})$, and the labels have the exact same interpretation as in (2.2), which we recall now:

$$d_{\mathbf{m}}(v, v_*) = \ell(v) - \min \ell + 1 , \qquad v \in V_\circ(\mathbf{t}) .$$

8.3 Uniform $2p$-angulations and random labeled mobiles

We now fix an integer $p \geq 2$ once and for all, and for $n \geq 1$ we let (T_n, ℓ_n) be a uniform labeled p-mobile with n black vertices, that is, a labeled mobile in which every vertex $u \in V_\bullet(T_n)$ has degree p (i.e. $p-1$ children). We also let ϵ be a uniform random variable in $\{-1, 1\}$, independent of the rest. The results of the previous section show that $(M_n, v_*) = \Phi((T_n, \ell_n), \epsilon)$ is a uniform random rooted and pointed $2p$-angulation with n faces.

From there on, it is natural to try and generalize the results we have obtained in the previous chapters. As mentioned in the introduction of this

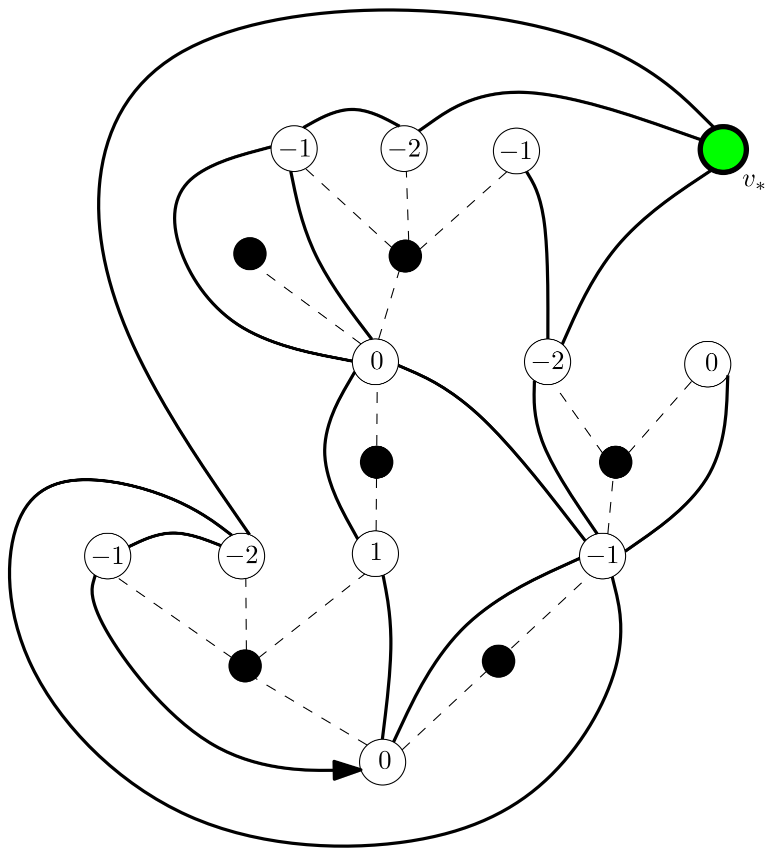

Figure 8.2: The Bouttier-Di Francesco-Guitter bijection performed on the labeled mobile of Figure 8.1, with ϵ chosen so that the root edge of the map points towards the root of the mobile

chapter, a nice re-rooting argument of [56] shows that not so much is needed: in fact, we need only generalize in an appropriate way certain results of Section 4.2.

Let us be more specific. A simple application of the Euler formula shows that the $2p$-angulation M_n has pn edges and $(p-1)n+2$ vertices. In particular, the mobile T_n has n black vertices and $(p-1)n+1$ white vertices, for a total of pn edges, or $2pn$ oriented edges. Thus, the contour exploration of the *white* corners has pn distinct terms, $e_0^n, e_1^n, \ldots, e_{pn-1}^n, e_{pn}^n = e_0^n$. We let $u_i^n = (e_i^n)^-$ for $0 \le i \le pn$, and let

$$C_n(i) = \frac{1}{2} d_{T_n}(u_i^n, u_0^n)\,, \qquad L_n(i) = \ell_n(u_i^n)\,, \qquad 0 \le i \le pn\,.$$

As usual, these two processes are extended to $[0, 2n]$ by linear interpolation between integer times. These are the natural analogs of the contour and label processes considered before, the division by 2 in the definition of the contour process being motivated by the fact that white vertices have even heights. We have the following generalization of Theorem 3.4.1. For $0 \le t \le 1$ set

$$C_{(n)}(t) = \left(\frac{p}{4(p-1)n}\right)^{1/2} C_n(pnt)\,, \quad L_{(n)}(t) = \left(\frac{9}{4p(p-1)n}\right)^{1/4} L_n(pnt)\,.$$

Theorem 8.3.1. *We have the following convergence*

$$(C_{(n)}, L_{(n)}) \xrightarrow[n\to\infty]{(d)} (\mathbb{e}, Z)\,.$$

in distribution in $\mathcal{C}([0,1], \mathbb{R})^2$.

We are not going to prove this result, but content ourselves with explaining where the scaling factors come from. The convergence of $C_{(n)}$ to $\mathbb{e}$ is in fact a particular case of a rather general scaling limit result for Bienaymé-Galton-Watson (BGW) random trees. Let us explain how such trees enter the discussion. Let $\mu_\circ$ be a geometric law on $\mathbb{Z}_+$ with parameter $(p-1)/p$, so that its mean is $1/(p-1)$, and its variance is $p/(p-1)^2$. Consider a branching process starting from a single individual at generation 0, and such that

- each individual at even generations produces a random number of children at the next generation, with distribution $\mu_\circ$
- each individual at odd generations produces $p-1$ children at the next generation.

Here, we assume that the offspring of the different individuals involved are all independent. This is a simple instance of a two-type branching process, in which the types of individuals alternate between generations. If we decide to skip the odd generations, we see that this process boils down to a single-type branching process in which the offspring distribution μ of each individual is the law of $(p-1)G$, where G has distribution $\mu_\circ$. We see that μ has mean 1 and variance p, and therefore the branching process is critical and ends a.s. in finite time. This means that the genealogy of the branching process considered here is a.s. a finite tree T, in which the individuals are naturally partitioned as in mobiles, depending on their generations. The probability of a particular tree $\mathbf{t}$ is then

$$P(T=\mathbf{t}) = \prod_{v\in V_\circ(\mathbf{t})} \frac{p-1}{p^{k_v(\mathbf{t})+1}}\,,$$

whenever $\mathbf{t}$ is a p-mobile, and 0 otherwise. Since every black vertex in $\mathbf{t}$ is a child of a white vertex, and every white vertex except the root is a child of a black vertex, we see that

$$\sum_{v\in V_\circ(\mathbf{t})} k_v(\mathbf{t}) = \#V_\bullet(\mathbf{t})\,, \qquad (p-1)\#V_\bullet(\mathbf{t}) = \#V_\circ(\mathbf{t})-1\,,$$

so that this probability can be rewritten as

$$P(T=\mathbf{t}) = \left(\frac{p-1}{p}\right)^{\#V_\circ(\mathbf{t})} \frac{1}{p^{\#V_\bullet(\mathbf{t})}} = \frac{p-1}{p}\left(\frac{(p-1)^{p-1}}{p^p}\right)^{\#V_\bullet(\mathbf{t})}\,.$$

In particular, it depends on $\mathbf{t}$ only *via* $\#V_\bullet(\mathbf{t})$. Therefore, conditioning T on the event $\{\#V_\bullet(T)=n\}$ produces a uniform random rooted p-mobile.

Due to this discussion, we can view T_n as a two-type BGW tree conditioned on its total number of vertices at odd heights being n, or equivalently, conditioned on having $n(p-1)+1$ vertices at even heights (this last fact is due to the very particular form of the offspring distribution of black vertices). Let T° be the tree T in which odd generations are skipped, so that the vertices of T° are the elements of $V_\circ(T)$, and v is the parent of u in T° if and only if v is the grandparent of u in T. As mentioned above, T° is a usual, single-type BGW tree, with critical offspring distribution μ (the law of $(p-1)G$, where G is geometrically distributed with parameter $(p-1)/p$). Moreover, due to the discussion above, conditioning T on having n black vertices boils down to conditioning T° on having $n(p-1)+1$ vertices, and this has the same law as the tree T_n°, which is the tree T_n in which odd generations are skipped. Now, note that C_n is none other than the contour process associated with the tree T_n°. At this point, we can apply standard results for convergence of conditioned BGW trees [3], showing that

$$\left(\frac{C_n(pnt)}{\sqrt{(p-1)n}}\right)_{0\leq t\leq 1} \xrightarrow[n\to\infty]{(d)} \frac{2}{\sqrt{p}}\mathbb{e}\,, \tag{8.2}$$

where the normalization factor $(p-1)n$ on the left-hand side is asymptotically equivalent to the total number of vertices in the tree T_n°, and the factor $\sqrt{p}$ on the right-hand side is the standard deviation of μ. This explains the first marginal convergence in Theorem 8.3.1.

It remains to explain where the scaling factor in the second convergence comes from. Note that if (T_n, ℓ_n) is a uniformly chosen random labeled p-mobile with n black vertices, then T_n is a uniformly chosen p-mobile with n black vertices, and conditionally on T_n, ℓ_n is a uniform mobile-admissible labeling of T_n. Hence, the situation is analogous to the one encountered in the previous chapters. Namely, for $u \in V_\circ(T_n)$, we can write

$$\ell_n(u) = \sum_{\substack{v\in V_\circ(T_n)\prec u\\ v\neq u_0^n}} Y_v\,, \tag{8.3}$$

where $Y_v = \ell_n(v) - \ell_n(\neg\neg v)$ (note that we have to take the grandparent of v since $\neg v$ is a black vertex). Then, the different terms involved in this sum are independent. However, the main difference with the previous situation is that these are not identically distributed. Indeed, if v is the k-th child of its parent, with $k \in \{1, 2, \ldots, p-1\}$, then $\ell_n(v) - \ell_n(\neg\neg v)$ has the law of $X_1 + \ldots + X_k$ where $(X_1, \ldots, X_p)$ is uniform in the set

$$\{(x_1, \ldots, x_p) \in \{-1, 0, 1, 2, \ldots\} : x_1 + \ldots + x_p = 0\}\,.$$

Exercise: Show that $(X_1, \ldots, X_p)$ is an exchangeable sequence[1] of centered random variables, with

$$P(X_1 = l) = \frac{\binom{2p-l-3}{p-2}}{\binom{2p-1}{p}}\,, \qquad -1 \leq l \leq p-1\,.$$

[1] A sequence of random variables $(X_1, \ldots, X_p)$ is exchangeable if it has the same distribution as $(X_{\sigma(1)}, \ldots, X_{\sigma(p)})$ for every permutation σ of $\{1, 2, \ldots, p\}$.

Deduce that for $1 \leq k \leq p-1$, one has

$$\mathrm{Var}\,(X_1 + \ldots + X_k) = k\mathrm{Var}\,(X_1) + k(k-1)\mathrm{Cov}\,(X_1, X_2)\frac{2k(p-k)}{p+1}\,.$$

At this point, we see that the law of the random variable Y_v described above depends strongly on the *rank* of v among the $p-1$ children of $\neg v$. The intuition is that, in a typical branch of T_n, say the one going from u_0^n to $u_{\lfloor pnt \rfloor}^n$ for some $t \in (0,1)$, which contains about $h_n = \mathbb{e}_t\sqrt{4(p-1)n/p}$ white vertices according to (8.2), the number of times we see a white vertex being the k-th child of its parent is of order $h_n/(p-1)$, for every $k \in \{1, 2, \ldots, p-1\}$. Hence, the sum (8.3) for $u = u_{\lfloor pnt \rfloor}^n$ behaves as a sum of $p-1$ independent terms, each of which is a sum of $h_n/(p-1)$ identically distributed, centered random variables with respective variances $2k(p-k)/(p+1)$ for $1 \leq k \leq p-1$. The central limit theorem implies that this sum should be of order

$$N\sqrt{\frac{h_n}{p-1}\sum_{k=1}^{p-1}\frac{2k(p-k)}{p+1}} \sim N\mathbb{e}_t^{1/2}\left(\frac{4p(p-1)n}{9}\right)^{1/4}\,,$$

where N is a standard Gaussian random variable. We see that given $\mathbb{e}_t$, the random variable $\mathbb{e}_t^{1/2}N$ has the same marginal distribution as Z_t, and this explains the rescaling of the second marginal in Theorem 8.3.1.

Of course, turning these considerations into a rigorous argument requires some technicalities that we omit here.

8.4 The re-rooting argument

Starting from Theorem 8.3.1, it is not very difficult to see that the arguments of Section 4.2 carry over to this case almost *verbatim*. Namely, recall that the key tool in that section was the first upper-bound in Proposition 2.3.8, and it is not difficult to generalize it to the context of the BDG bijection, with a similar proof.

Therefore, one can easily obtain the following result. Define a pseudo-metric $d_n(i,j) = d_{M_n}(u_i^n, u_j^n)$ for $0 \leq i, j \leq pn$, and extend it to $[0, pn]$ by a similar interpolation as in (4.2). Then define

$$D_n(s,t) = \left(\frac{9}{4p(p-1)n}\right)^{1/4} d_n(pns, pnt)\,, \qquad 0 \leq s, t \leq 1\,.$$

Theorem 8.4.1. *The family of laws of* $(D_n(s,t), 0 \le s,t \le 1)$*, as* n *varies, is relatively compact for the weak topology on probability measures on* $\mathcal{C}([0,1]^2, \mathbb{R})$*. Moreover, every subsequential limit* $(\mathbb{e}, Z, D')$ *of* $(C_{(n)}, L_{(n)}, D_n)$ *satisfies the following properties a.s.:*

1. D' *is a pseudo-metric on* $[0,1]$*,*
2. *for every* $s,t \in [0,1]$*,* $d_{\mathbb{e}}(s,t) = 0$ *implies* $D'(s,t) = 0$
3. *for every* $s,t \in [0,1]$*,* $D'(s,t) \le d_Z(s,t)$
4. *for every* $s \in [0,1]$*,* $D'(s, s_*) = Z_s - \inf Z$
5. *if* U, V *are uniform random variables in* $[0,1]$*, independent of all other random variables considered, then* $D'(U,V)$ *and* $D'(s_*, U)$ *have the same distribution.*

We leave the proof as an exercise to the reader, and a good way to review the material presented in Chapter 4. Here is how one proves Theorem 8.1.1 starting from there.

PROOF. [Theorem 8.1.1] As in Theorem 8.4.1, let $(\mathbb{e}, Z, D')$ be a subsequential limit of $(C_{(n)}, L_{(n)}, D_n)$. It suffices to show that $D' = D = D^*$ is the Brownian map pseudo-metric on $[0,1]$ encoded by $(\mathbb{e}, Z)$.

From 1., 2. and 3. in Theorem 8.4.1, we immediately obtain $D' \le D$ by the maximality property of $D^* = D$. Let U, V be independent uniform random variables in $[0,1]$, independent of the other random variables considered. Then

$$E[D'(U,V)] = E[D'(s_*, U)] = E[Z_U - \inf Z] = E[D(s_*, U)] = E[D(U,V)]\,,$$

where we have used the re-rooting invariance of D. Since $D' \le D$, this means that $D'(U,V) = D(U,V)$ a.s. By Fubini's theorem, this means that a.s., $D'(u,v) = D(u,v)$ for Lebesgue-almost all $(u,v) \in [0,1]^2$. By continuity of D, D' and a density argument, we get $D = D'$, as wanted. □

This proof can appear extremely simple, to the point that the reader may ask her/himself why such an argument has not been applied straight away to quadrangulations in the earlier chapters of these notes. The reason is that we have used in a crucial way that the quotient distance $D = D^*$ is invariant under re-rooting, a property that we could deduce from the fact that it is the limiting distance for quadrangulations, that are obviously re-root-invariant.

In other words, if we knew from purely continuum arguments that D^* was invariant under re-rooting, then we would get another proof of the convergence of random quadrangulations to the Brownian map. However, such a proof is still missing.

Notes for Chapter 8

Understanding the extent and limits of the "universality class" of the Brownian map is in itself a natural line of research. Note that the convergence of uniform p-angulations to the Brownian map for odd $p \geq 5$ is strongly believed to hold, but not known. This is mainly due to the technicality of the BDG bijection in non-bipartite cases, although it simplifies notably in the case $p = 3$, and indeed the argument above was also performed in this case by Le Gall [56], hence solving the initial question raised by Schramm in his ICM lecture [77]. The papers by Le Gall and Beltran [12], and Addario-Berry and Albenque [2] have investigated scaling limits of random maps with additional topological constraints, the latter focusing in particular on the case of simple triangulations: see [2] for a discussion of the motivation for this result in relation with circle packings. C. Abraham [1] and Bettinelli, Jacob and Miermont [19] have proved convergence of uniform random maps with n edges ([1] dealing with the bipartite case) to the Brownian map. Although the latter problem deals with a non-bipartite case, [19] circumvents the inherent difficulty of the BDG bijection by making use of an alternative bijection due to Ambjørn and Budd [4]. Hence, one obtains the following result, which is an exact analog of Theorem 3.3.4 in the world of maps instead of trees.

Theorem 8.4.2. *Let M_n be a uniform random rooted map with n edges. Then*

$$\left(V(M_n), \left(\frac{9}{8n}\right)^{1/4} d_{M_n}\right) \xrightarrow[n\to\infty]{(d)} (S, D)\,,$$

in distribution for the Gromov-Hausdorff topology, where (S, D) is the Brownian map.

A striking fact is that the scaling constant $(8/9)^{1/4}$ is the same as for quadrangulations. This is somehow consistent with the fact that these models are in direct relation through the "trivial" bijection of Section 2.3, however, it is not known so far if this bijection asymptotically preserves the graph distances. Another natural and related question would be to ask whether, if M_n^* is the dual map to M_n, it holds that M_n and M_n^* are asymptotically isometric, in the sense that

$$\mathrm{d_{GH}}((V(M_n), d_{M_n}), (V(M_n^*), d_{M_n^*})) = o(n^{1/4})$$

in probability as $n \to \infty$. Of course, one could ask the same question for any model of random map that is known to converge to the Brownian map, but in this case, M_n and M_n^* obviously have the same law. To our knowledge, it is not known whether $(V(Q_n^*), n^{-1/4} d_{Q_n^*})$ converges to some multiple of the Brownian map, where Q_n^* is the dual graph of the uniform quadrangulation with n faces (so Q_n^* is a uniform tetravalent map with n vertices).

The common feature of all models considered so far is that the largest face degrees in these models is small compared to the total diameter of the maps. It is however possible to force faces with large degrees to appear for suitable choices of the weights involved in the definition of Boltzmann maps. For these choices, one can prove scaling limit results, but the limit is not the Brownian map anymore. See Le Gall and Miermont [57]. This problem is partly motivated by the study of statistical physics models on random maps, see e.g. [23, 22, 21].

Bibliography

[1] C. Abraham. Rescaled bipartite planar maps converge to the Brownian map. *Preprint, http://arxiv.org/abs/1312.5959*, 2013.

[2] L. Addario-Berry and M. Albenque. The scaling limit of random simple triangulations and random simple quadrangulations. *Preprint, http://arxiv.org/abs/1306.5227*, 2013.

[3] D. J. Aldous. The continuum random tree. III. *Ann. Probab.*, 21(1):248–289, 1993.

[4] J. Ambjørn and T. G. Budd. Trees and spatial topology change in causal dynamical triangulations. *J. Phys. A*, 46(31):315201, 33, 2013.

[5] J. Ambjørn, B. Durhuus, and T. Jonsson. *Quantum geometry. A statistical field theory approach.* Cambridge Monographs on Mathematical Physics. Cambridge University Press, Cambridge, 1997.

[6] J. Ambjørn and Y. Watabiki. Scaling in quantum gravity. *Nuclear Phys. B*, 445(1):129–142, 1995.

[7] S. Andres and N. Kajino. Continuity and estimates of the liouville heat kernel with applications to spectral dimensions. 2014. arXiv:1407.3240.

[8] O. Angel. Growth and percolation on the uniform infinite planar triangulation. *Geom. Funct. Anal.*, 13(5):935–974, 2003.

[9] O. Angel and N. Curien. Percolations on random maps I: Half-plane models. *Ann. Inst. H. Poincaré (B)*, 2014+. To appear.

[10] D. Arquès. Les hypercartes planaires sont des arbres très bien étiquetés. *Discrete Math.*, 58(1):11–24, 1986.

[11] E. G. Begle. Regular convergence. *Duke Math. J.*, 11:441–450, 1944.

[12] J. Beltran and J.-F. Le Gall. Quadrangulations with no pendant vertices. *Bernoulli*, 19(4):1150–1175, 2013.

[13] E. A. Bender and E. R. Canfield. The asymptotic number of rooted maps on a surface. *J. Combin. Theory Ser. A*, 43(2):244–257, 1986.

[14] I. Benjamini and N. Curien. Simple random walk on the uniform infinite planar quadrangulation: subdiffusivity via pioneer points. *Geom. Funct. Anal.*, 23(2):501–531, 2013.

[15] N. Berestycki. Diffusion in planar Liouville quantum gravity. 2013. arXiv:1301.3356.

[16] N. Berestycki, C. Garban, R. Rhodes, and V. Vargas. KPZ formula derived from Liouville heat kernel. 2014. arXiv:1406.7280.

[17] J. Bettinelli. The topology of scaling limits of positive genus random quadrangulations. *Ann. Probab.*, 40:no. 5, 1897–1944, 2012.

[18] J. Bettinelli. Geodesics in Brownian surfaces (Brownian maps). 2014. arXiv:1401.3602.

[19] J. Bettinelli, E. Jacob, and G. Miermont. The scaling limit of uniform random plane maps, via the Ambjørn-Budd bijection. *Electr. J. Probab.*, 2013. To appear.

[20] P. Biane. Relations entre pont et excursion du mouvement brownien réel. *Ann. Inst. Henri Poincaré Probab. Stat.*, 22(1):1–7, 1986.

[21] G. Borot, J. Bouttier, and E. Guitter. Loop models on random maps via nested loops: the case of domain symmetry breaking and application to the Potts model. *J. Phys. A*, 45(49):494017, 35, 2012.

[22] G. Borot, J. Bouttier, and E. Guitter. More on the $O(n)$ model on random maps via nested loops: loops with bending energy. *J. Phys. A*, 45(27):275206, 32, 2012.

[23] G. Borot, J. Bouttier, and E. Guitter. A recursive approach to the $O(n)$ model on random maps via nested loops. *J. Phys. A*, 45(4):045002, 38, 2012.

[24] M. Bousquet-Mélou and A. Jehanne. Polynomial equations with one catalytic variable, algebraic series and map enumeration. *J. Combin. Theory Ser. B*, 96(5):623–672, 2006.

[25] J. Bouttier, P. Di Francesco, and E. Guitter. Planar maps as labeled mobiles. *Electron. J. Combin.*, 11:Research Paper 69, 27 pp. (electronic), 2004.

[26] J. Bouttier and E. Guitter. The three-point function of planar quadrangulations. *J. Stat. Mech. Theory Exp.*, 7:P07020, 39, 2008.

[27] E. Brézin, C. Itzykson, G. Parisi, and J. B. Zuber. Planar diagrams. *Comm. Math. Phys.*, 59(1):35–51, 1978.

[28] D. Burago, Y. Burago, and S. Ivanov. *A course in metric geometry*, volume 33 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2001.

[29] G. Chapuy, M. Marcus, and G. Schaeffer. A bijection for rooted maps on orientable surfaces. *SIAM J. Discrete Math.*, 23(3):1587–1611, 2009.

[30] P. Chassaing and B. Durhuus. Local limit of labeled trees and expected volume growth in a random quadrangulation. *Ann. Probab.*, 34(3):879–917, 2006.

[31] P. Chassaing and G. Schaeffer. Random planar lattices and integrated superBrownian excursion. *Probab. Theory Related Fields*, 128(2):161–212, 2004.

[32] I. Chiswell. *Introduction to* Λ*-trees*. World Scientific Publishing Co., Inc., River Edge, NJ, 2001.

[33] R. Cori and B. Vauquelin. Planar maps are well labeled trees. *Canad. J. Math.*, 33(5):1023–1042, 1981.

[34] N. Curien. A glimpse of the conformal structure of random planar maps. 2013. arXiv:1308.1807.

[35] N. Curien and J.-F. Le Gall. The Brownian plane. *J. Theor. Probab.*, 2014+. To appear.

[36] N. Curien, L. Ménard, and G. Miermont. A view from infinity of the uniform infinite planar quadrangulation. *ALEA Lat. Am. J. Probab. Math. Stat.*, 10(1):45–88, 2013.

[37] J.-F. Delmas. Computation of moments for the length of the one dimensional ISE support. *Electron. J. Probab.*, 8:no. 17, 15 pp. (electronic), 2003.

[38] B. Duplantier and S. Sheffield. Liouville quantum gravity and KPZ. *Invent. Math.*, 185(2):333–393, 2011.

[39] P. Flajolet and R. Sedgewick. *Analytic Combinatorics*. Cambridge University Press, Cambridge, 2009.

[40] C. Garban, R. Rhodes, and V. Vargas. Liouville Brownian motion. 2013. arXiv:1301.2876.

[41] C. Garban, R. Rhodes, and V. Vargas. On the heat kernel and the Dirichlet form of Liouville Brownian motion. 2013. arXiv:1302.6050.

[42] I. P. Goulden and D. M. Jackson. *Combinatorial enumeration*. Dover Publications Inc., Mineola, NY, 2004. With a foreword by Gian-Carlo Rota, Reprint of the 1983 original.

[43] I. P. Goulden and D. M. Jackson. The KP hierarchy, branched covers, and triangulations. *Adv. Math.*, 219(3):932–951, 2008.

[44] A. Greven, P. Pfaffelhuber, and A. Winter. Convergence in distribution of random metric measure spaces (Λ-coalescent measure trees). *Probab. Theory Related Fields*, 145(1-2):285–322, 2009.

[45] M. Gromov. *Metric structures for Riemannian and non-Riemannian spaces*, volume 152 of *Progress in Mathematics*. Birkhäuser Boston Inc., Boston, MA, 1999.

[46] O. Gurel-Gurevich and A. Nachmias. Recurrence of planar graph limits. *Ann. of Math. (2)*, 177(2):761–781, 2013.

[47] V. G. Knizhnik, A. M. Polyakov, and A. B. Zamolodchikov. Fractal structure of 2D-quantum gravity. *Modern Phys. Lett. A*, 3(8):819–826, 1988.

[48] M. Krikun. Local structure of random quadrangulations. 2005. Preprint.

[49] M. A. Krikun. A uniformly distributed infinite planar triangulation and a related branching process. *Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI)*, 307(Teor. Predst. Din. Sist. Komb. i Algoritm. Metody. 10):141–174, 282–283, 2004.

[50] S. K. Lando and A. K. Zvonkin. *Graphs on surfaces and their applications*, volume 141 of *Encyclopaedia of Mathematical Sciences*. Springer-Verlag, Berlin, 2004.

[51] J.-F. Le Gall. Random trees and applications. *Probab. Surv.*, 2:245–311 (electronic), 2005.

[52] J.-F. Le Gall. A conditional limit theorem for tree-indexed random walk. *Stochastic Process. Appl.*, 116(4):539–567, 2006.

[53] J.-F. Le Gall. Random real trees. *Ann. Fac. Sci. Toulouse Math. (6)*, 15(1):35–62, 2006.

[54] J.-F. Le Gall. The topological structure of scaling limits of large planar maps. *Invent. Math.*, 169(3):621–670, 2007.

[55] J.-F. Le Gall. Geodesics in large planar maps and in the Brownian map. *Acta Math.*, 205(2):287–360, 2010.

[56] J.-F. Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.*, 41(4):2880–2960, 2013.

[57] J.-F. Le Gall and G. Miermont. Scaling limits of random planar maps with large faces. *Ann. Probab.*, 39(1):1–69, 2011.

[58] J.-F. Le Gall and G. Miermont. Scaling limits of random trees and planar maps. In *Probability and statistical physics in two and more dimensions*, volume 15 of *Clay Math. Proc.*, pages 155–211. Amer. Math. Soc., Providence, RI, 2012.

[59] J.-F. Le Gall and F. Paulin. Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere. *Geom. Funct. Anal.*, 18(3):893–918, 2008.

[60] J.-F. Le Gall and M. Weill. Conditioned Brownian trees. *Ann. Inst. H. Poincaré Probab. Statist.*, 42(4):455–489, 2006.

[61] P. Maillard, R. Rhodes, V. Vargas, and O. Zeitouni. Liouville heat kernel: regularity and bounds. 2014. arXiv:1406.0491.

[62] J.-F. Marckert and G. Miermont. Invariance principles for random bipartite planar maps. *Ann. Probab.*, 35(5):1642–1705, 2007.

[63] J.-F. Marckert and A. Mokkadem. Limit of normalized random quadrangulations: the Brownian map. *Ann. Probab.*, 34(6):2144–2202, 2006.

[64] L. Ménard. The two uniform infinite quadrangulations of the plane have the same law. *Ann. Inst. Henri Poincaré Probab. Stat.*, 46(1):190–208, 2010.

[65] G. Miermont. An invariance principle for random planar maps. In *Fourth Colloquium on Mathematics and Computer Sciences CMCS'06*, Discrete Math. Theor. Comput. Sci. Proc., AG, pages 39–58 (electronic). Nancy, 2006.

[66] G. Miermont. On the sphericity of scaling limits of random planar quadrangulations. *Electron. Commun. Probab.*, 13:248–257, 2008.

[67] G. Miermont. Tessellations of random maps of arbitrary genus. *Ann. Sci. Éc. Norm. Supér. (4)*, 42(5):725–781, 2009.

[68] G. Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.*, 210(2):319–401, 2013.

[69] G. Miermont and M. Weill. Radius and profile of random planar maps with faces of arbitrary degrees. *Electron. J. Probab.*, 13:no. 4, 79–106, 2008.

[70] J. Miller and S. Sheffield. Quantum Loewner Evolution. 2013. arXiv:1312.5745.

[71] B. Mohar and C. Thomassen. *Graphs on surfaces*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, 2001.

[72] R. L. Moore. Concerning upper semi-continuous collections of continua. *Trans. Amer. Math. Soc.*, 27(4):416–428, 1925.

[73] A. Okounkov. Random matrices and random permutations. *Internat. Math. Res. Notices*, (20):1043–1095, 2000.

[74] J. Pitman. *Combinatorial stochastic processes*, volume 1875 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2006. Lectures from the 32nd Summer School on Probability Theory held in Saint-Flour, July 7–24, 2002, With a foreword by Jean Picard.

[75] A. M. Polyakov. Quantum geometry of bosonic strings. *Phys. Lett. B*, 103(3):207–210, 1981.

[76] G. Schaeffer. *Conjugaison d'arbres et cartes combinatoires aléatoires.* PhD thesis, Université Bordeaux I, 1998.

[77] O. Schramm. Conformally invariant scaling limits: an overview and a collection of problems. In *International Congress of Mathematicians. Vol. I*, pages 513–543. Eur. Math. Soc., Zürich, 2007.

[78] G. 't Hooft. A planar diagram theory for strong interactions. *Nucl. Phys. B*, 72:461–473, 1974.

[79] W. T. Tutte. A census of planar maps. *Canad. J. Math.*, 15:249–271, 1963.

[80] M. Weill. Asymptotics for rooted planar maps and scaling limits of two-type spatial trees. *Electron. J. Probab.*, 12:Paper no. 31, 862–925 (electronic), 2007.

[81] A. Zvonkin. Matrix integrals and map enumeration: an accessible introduction. *Math. Comput. Modelling*, 26(8-10):281–304, 1997. Combinatorics and physics (Marseilles, 1995).