Large random planar maps
and their scaling limits

Jean-François Le Gall

Abstract

We discuss scaling limits of random planar maps chosen uniformly at random in a certain class. This leads to a universal limiting space called the Brownian map, which is viewed as a random compact metric space. The Brownian map can be obtained as a quotient of the continuous random tree called the CRT, for an equivalence relation which is defined in terms of Brownian labels assigned to the vertices of the CRT. We discuss the known properties of the Brownian map. In particular, we give a complete description of the geodesics starting from the distinguished point called the root. We also discuss applications to various properties of large random planar maps.

1 Introduction

The main purpose of the present article is to survey recent developments about scaling limits of large planar maps chosen uniformly at random in a suitable class. Recall that planar maps are just (finite) graphs embedded in the plane. A planar map is thus the kind of object one would draw on a sheet of a paper if asked to give an example of a graph.

To explain what a scaling limit is, consider a combinatorial object, such as a path, a tree or a graph, and suppose that it is chosen at random in the class of all objects of size n. Often the resulting random object can be rescaled as $n \to \infty$ in such a way that it becomes close to a continuous model. For instance, one may consider all discrete paths with length n starting from the origin on the integer lattice $\mathbb{Z}^d$. If one chooses uniformly at random a path in this collection, then modulo a suitable rescaling (essentially by the factor $1/\sqrt{n}$) it will become close to a continuous Brownian path. More precisely, for any set A in the path space, satisfying certain regularity asssumptions, the probability that the rescaled discrete path of length n belongs to A will converge to the probability that the Brownian path belongs to A as $n \to \infty$

Studying such scaling limits is all the more interesting as they are universal, meaning that the same continuous model corresponds to the limit of many different classes of discrete objects. A fundamental example of this universality property is Brownian motion, which is well known to be the scaling limit of many different classes of random paths. The study of scaling limits is motivated by at least two important reasons:

- Often the continuous model is of interest in its own. For instance, Brownian motion has numerous applications, independently of the fact that it is the scaling limit of random walks.
- Knowing the continuous model gives insight into the properties of the large discrete objects. Lots of interesting distributional asymptotics for long random paths can be derived from explicit calculations on Brownian motion.

In the present work, we discuss scaling limits first for random trees and then for random planar maps. The reason for considering random trees first comes from our specific approach, which involves bijections between planar maps and certain classes of decorated trees. The scaling limits of trees and maps both lead to remarkable probabilistic objects. In the case of trees, the scaling limit is the CRT (Continuum Random Tree), which has been introduced and studied by Aldous [A1, A2] in the early nineties. The scaling limit of random planar maps, which we call the Brownian map, is then described as the quotient of the CRT for a certain (random) equivalence relation. The Brownian map may be thought of as the relevant probabilistic model for a random surface in the same sense as Brownian motion is the right model for a purely random continuous path. Indeed, one conjectures that the Brownian map appears as the continuous limit of many classes of planar maps, which are natural discretizations of surfaces.

Let us recall some basic definitions. A planar map is a proper embedding of a finite connected graph in the two-dimensional sphere $\mathbb{S}^2$. Loops and multiple edges are a priori allowed. The faces of the map are the connected components of the complement of the union of edges. A planar map is rooted if it has a distinguished oriented edge called the root edge, whose origin is called the root vertex. In what follows, we consider only rooted planar maps, even if this is not mentioned explicitly. Rooting maps avoids certain technical difficulties and is believed to have no influence on the problems we will be addressing.

Two rooted planar maps are said to be equivalent if the second one is the image of the first one under an orientation-preserving homeomorphism of the sphere, which also preserves the root edges. Two equivalent rooted planar maps will always be identified.

Given an integer $p \geq 3$, a p-angulation is a planar map where each face has degree p, that is p adjacent edges. One should count edge sides, so that if an edge lies entirely inside a face it is counted twice: For instance, the face in the upper right corner of Fig.1 below has degree 4, although it seems to be adjacent to only 3 edges. We denote by $\mathbb{M}_n^p$ the set of all rooted p-angulations with n faces. Thanks to the preceding identification, the set $\mathbb{M}_n^p$ is finite. A 3-angulation is called a triangulation, and a 4-angulation is called a quadrangulation. Fig.1 below shows a quadrangulation with 7 faces.

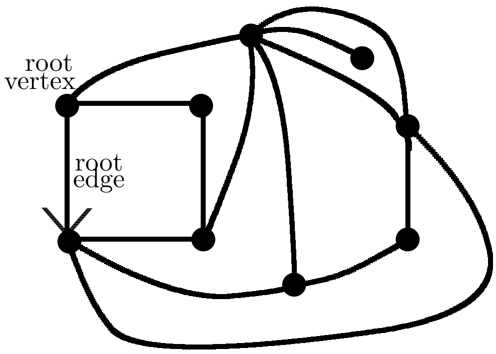

Figure 1.

Consider a planar map M. Let $V(M)$ denote the vertex set of M. A path in M with length k is a finite sequence $(a_0, a_1, \ldots, a_k)$ in $V(M)$ such that a_i and a_{i-1} are connected by an edge of the map, for every $i \in \{1, \ldots, k\}$. The graph distance $d_{gr}(a, a')$ beween two vertices a and a' is the minimal k such that there exists a path $\gamma = (a_0, a_1, \ldots, a_k)$ with $a_0 = a$ and $a_k = a'$. A path $\gamma = (a_0, a_1, \ldots, a_k)$ is called a discrete geodesic (from a_0 to a_k) if $k = d_{gr}(a_0, a_k)$. The set $V(M)$ equipped with the metric d_{gr} is a (finite) metric space. Clearly, the map M is not determined by the metric space $(V(M), d_{gr})$. Nonetheless, much information is contained in this metric space, and in what follows we will concentrate on the study of metric properties of planar maps.

Fix an integer $p \geq 3$ and, for every integer $n \geq 2$, let M_n be a random planar map chosen uniformly at random in the space $\mathbb{M}_n^p$. Following our initial discussion of scaling limits, one would like to prove that for a suitable choice of the positive constant α, the rescaled random metric spaces

$$(V(M_n), n^{-\alpha} d_{gr}) \tag{1}$$

converge in some appropriate sense towards a (non-degenerate) limiting random compact metric space. Moreover the limiting space is believed to be independent of p, up to trivial scaling factors. This corresponds to the universality property mentioned above.

The rescaling factor $n^{-\alpha}$ in (1) is needed if we want to get a "continuous" limit and to stay within the framework of compact metric spaces. It also makes sense to study the limit of the spaces $(V(M_n), d_{gr})$ without rescaling, and this gives rise to infinite random graphs (see Angel [An] and Angel and Schramm [AS] for the case of infinite triangulations, and Chassaing and Durhuus [CS] and Krikun [Kr] for infinite quadrangulations of the plane).

As stated above, the problem of the scaling limit for planar maps requires an adequate notion of the convergence of a sequence of compact metric spaces. Such a notion is provided by the Gromov-Hausdorff distance (Gromov [Gr], Burago, Burago and Ivanov [BBI]). Let (E_1, d_1) and (E_2, d_2) be two compact metric spaces. The Gromov-Hausdorff distance between (E_1, d_1) and (E_2, d_2) is

$$d_{GH}(E_1, E_2) = \inf \Big(d_{Haus}(\varphi_1(E_1), \varphi_2(E_2)) \Big),$$

where the infimum is over all isometric embeddings $\varphi_1 : E_1 \longrightarrow E$ and $\varphi_2 : E_2 \longrightarrow E$ of E_1 and E_2 into the same metric space (E, d), and d_{Haus} stands for the usual Hausdorff distance between compact subsets of E. If $\mathbb{K}$ denotes the space of all isometry classes of compact metric spaces, then d_{GH} is a distance on $\mathbb{K}$, and moreover the metric space $(\mathbb{K}, d_{GH})$ is Polish, that is separable and complete (see Chapter 7 of Burago, Burago and Ivanov [BBI] for a thorough discussion of the Gromov-Hausdorff distance).

Thanks to the previous discussion, it makes sense to study the convergence in distribution of the random metric spaces (1) as random variables with values in the Polish space $(\mathbb{K}, d_{GH})$. This problem was stated in this form for triangulations by Schramm [Sc]. The general idea of finding a continuous limit for large random planar maps had appeared earlier, especially in the pioneering paper of Chassaing and Schaeffer [CS]. The latter paper proves a limit theorem showing that the radius, or maximal distance from the root, of a quadrangulation with n faces chosen at random, rescaled by the factor $n^{-1/4}$, converges in distribution towards a nondegenerate limit (see Corollary 3.4 below). This gives evidence of the fact that the proper value of the constant α in (1) should be $\alpha = 1/4$.

For reasons that will be explained below, it turns out to be easier to handle bipartite planar maps: A planar map is bipartite if and only if all its faces have even degree. In the remaining part of this introduction, we thus restrict our attention to the case when p is even.

In order to explain our main result about scaling limits of planar maps, we need to introduce some notation. Aldous' Continuum Random Tree (the CRT), viewed as a random compact metric space, is denoted by $(\mathcal{T}_{\mathbf{e}}, d_{\mathbf{e}})$. Its root ρ is a distinguished point of $\mathcal{T}_{\mathbf{e}}$. The reason for the notation $\mathcal{T}_{\mathbf{e}}$ comes from the fact that the CRT can be coded by a normalized Brownian excursion $\mathbf{e}$, as will be explained in Section 3 below. This coding makes it possible to introduce a lexicographical order on the tree $\mathcal{T}_{\mathbf{e}}$: If $a, b \in \mathcal{T}_{\mathbf{e}}$, one may consider the "lexicographical" interval $[a, b]$ which is informally defined as the subset of $\mathcal{T}_{\mathbf{e}}$ consisting of all points that are visited when going from a to b around the tree in clockwise order (see Section 4 for more rigorous definitions). Next, conditionally given $(\mathcal{T}_{\mathbf{e}}, d_{\mathbf{e}})$, we consider a centered Gaussian process $(Z_a)_{a \in \mathcal{T}_{\mathbf{e}}}$ such that $Z_\rho = 0$ and

$$E[(Z_a - Z_b)^2] = d_{\mathbf{e}}(a, b)$$

for every $a, b \in \mathcal{T}_{\mathbf{e}}$ (again this definition is slightly informal, as we consider a random process indexed by a random set – see Section 4 for a more rigorous presentation). The process Z should be understood as Brownian motion indexed by the tree $\mathcal{T}_{\mathbf{e}}$: Z_a is a "label" assigned to vertex a, and this label evolves as linear Brownian motion when varying a along a line segment of the tree. Finally, we define a random equivalence relation $\approx$ on $\mathcal{T}_{\mathbf{e}}$ by setting

$$a \approx b \qquad \text{iff} \qquad Z_a = Z_b = \min_{c \in [a,b]} Z_c \quad \text{or} \quad Z_a = Z_b = \min_{c \in [b,a]} Z_c.$$

Then, Theorem 4.1 below, taken from [L2], states that, from any sequence of values of n converging to $+\infty$, we can extract a subsequence along which we have the convergence in distribution

$$(V(M_n), n^{-1/4} d_{gr}) \longrightarrow (\mathcal{T}_{\mathbf{e}}/\approx, D) \tag{2}$$

where D is a metric on the quotient $\mathcal{T}_{\mathbf{e}}/\approx$, which induces the quotient topology on that space. The limiting space $(\mathcal{T}_{\mathbf{e}}/\approx, D)$ is called the Brownian map (to be more precise, we should say that we use the name Brownian map for any of the limiting random metric spaces that can arise in (2) when we vary p and the subsequence). This terminology comes from Marckert and Mokkadem [MMo], who discussed limits of rescaled random quadrangulations, however in a different sense than the Gromov-Hausdorff convergence. Our terminology slightly differs from that in [MMo], where the Brownian map is defined as the space $\mathcal{T}_{\mathbf{e}}/\approx$ with a specified metric which may or may not coincide with D.

The need for a subsequence in (2) comes from the fact that the limiting random metric D has not been fully characterized, and so there might be different metrics D corresponding to different subsequences. Still one believes that it should not be necessary to take a subsequence, and that the limiting metric space should be the same independently of p (even if p is odd), thus confirming the universality property of the Brownian map. The recent results of Marckert and Miermont [MMi], Miermont [Mi1] and Miermont and Weill [MW] strongly support this conjecture.

Even though the distribution of the Brownian map has not been fully characterized, many of its properties can be investigated in detail. In Section 5 below, we give two theorems showing on one hand that the Hausdorff dimension of the Brownian map is a.s. equal to 4, and on the other hand that the Brownian map is a.s. homeomorphic to the two-dimensional sphere. The last result is maybe not surprising since we started from graphs drawn on the sphere. Still it implies that typical large p-angulations will not have "small bottlenecks" (see Corollary 5.3 for a precise statement). In Section 6 we present recent results taken from [L3] about the structure of geodesics in the Brownian map. Here again, we provide applications to properties of large discrete planar maps, in the spirit of the observations made at the beginning of this introduction.

One may ask why the scaling limit of random planar maps should be related to the CRT. This can be understood from the existence of bijections between the sets $\mathbb{M}_n^p$ and various classes of labeled trees. In the particular case of quadrangulations, such bijections were discovered by Cori and Vauquelin [CS] and then studied extensively by Schaeffer [S]. More recently, Bouttier, Di Francesco and Guitter [BDG] provided a nice simple extension of the Cori-Vauquelin-Schaeffer bijection to bipartite planar maps (see Section 2 below). This result partly explains why we restrict our attention to bipartite planar maps: The bijections in the general case seem more difficult to use for technical reasons (see however Miermont [Mi1]). The scaling limit of the discrete trees that arise in the bijections with planar maps turns out to be given by the CRT (see Section 3). Since in the discrete setting vertices of the map are in one-to one correspondence with vertices of the associated tree, it is not surprising that the Brownian map can be constructed from the CRT. However, the correct definition requires identifying certain pairs of points in the CRT, via the introduction of the equivalence relation $\approx$. This is so because, already in the discrete setting, certain pairs of vertices that are far away from each other in the tree can be very close in the associated map. The principal difficulty in the proof of (2) is in fact to determine precisely those pairs of points that need to be identified in the continuous limit.

To conclude this introduction, let us briefly comment on the motivations for studying planar maps and their scaling limits. Planar maps were first studied by Tutte [Tu] in connection with his work on the four color theorem, and since then they have been studied extensively in combinatorics. Planar maps also have algebraic and geometric applications: See the book of Lando and Zvonkin [LZ] for more on this matter. Because of their relations with Feynman diagrams, planar maps soon attracted the attention of specialists of theoretical physics. The pioneering papers by 't Hooft [tH] and Brézin, Itzykson, Parisi and Zuber [BIP] related enumeration problems for planar maps with asymptotics of matrix integrals. The interest for random planar maps in theoretical physics grew significantly when these combinatorial objects were interpreted as models of random surfaces, especially in the setting of the theory of quantum gravity (see in particular the book of Ambjørn, Durhuus and Jonsson [ADJ]). Bouttier's thesis [Bo] describes applications of planar maps to the statistical physics of random surfaces. The recent papers [BG1, BG2, BG3] by Bouttier and Guitter address questions closely related to those of the present work from the perspective of theoretical physics. From the probabilistic point of view, the Brownian map appears to be a fascinating model of a random fractal surface, even if its properties are still far from being completely understood.

2 Bijections betwen maps and trees

Throughout the remaining part of this work, we fix an integer $p \geq 2$ and we deal with the set $\mathbb{M}_n^{2p}$ of all rooted $2p$-angulations with n faces. We will present a bijection between $\mathbb{M}_n^{2p}$ and and a certain set of labeled trees.

By definition, a plane tree τ is a finite subset of the set

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

of all finite sequences of positive integers (including the empty sequence $\emptyset$), which satisfies three obvious conditions: First $\emptyset \in \tau$, then, for every $v = (u_1, \ldots, u_k) \in \tau$ with $k \geq 1$, the sequence $(u_1, \ldots, u_{k-1})$ (the "parent" of v) also belongs to τ, and finally for every $v = (u_1, \ldots, u_k) \in \tau$ there exists an integer $k_v(\tau) \geq 0$ (the "number of children" of v) such that the vertex $vj := (u_1, \ldots, u_k, j)$ belongs to τ if and only if $1 \leq j \leq k_v(\tau)$. The generation of $v = (u_1, \ldots, u_k)$ is denoted by $|v| = k$.

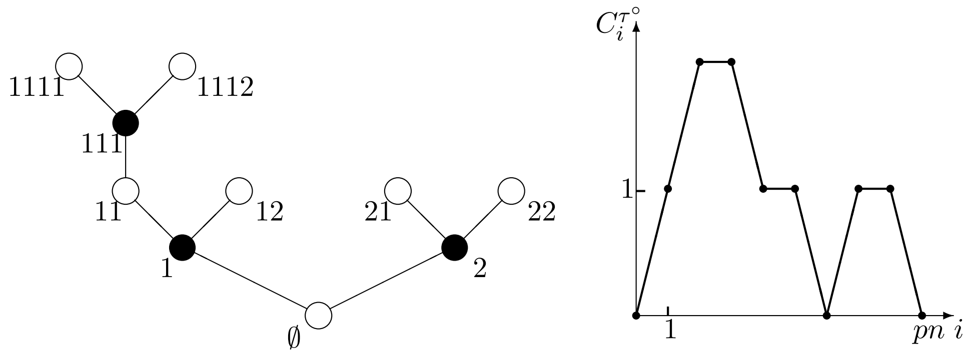

Figure 2. A 3-tree τ and the associated contour function C^{τ° of τ°.

A p-tree is a plane tree τ that satisfies the following additional property: For every $v \in \tau$ such that $|v|$ is odd, $k_v(\tau) = p - 1$.

If τ is a p-tree, vertices v of τ such that $|v|$ is even are called white vertices, and vertices v of τ such that $|v|$ is odd are called black vertices. We denote by τ° the set of all white vertices of τ and by $\tau^\bullet$ the set of all black vertices. See the left side of Fig.2 for an example of a 3-tree.

A labeled p-tree is a pair $\theta = (\tau, (\ell_v)_{v\in\tau^\circ})$ that consists of a p-tree τ and a collection of integer labels (taking values in $\mathbb{Z}$) assigned to the white vertices of τ, such that the following properties hold:

(a) $\ell_\emptyset = 1$.

(b) Let $v \in \tau^\bullet$, let $v_{(0)}$ be the parent of v and let $v_{(j)} = vj$, $1 \le j \le p-1$, be the children of v. Then for every $j \in \{0, 1, \dots, p-1\}$, $\ell_{v_{(j+1)}} \ge \ell_{v_{(j)}} - 1$, where by convention $v_{(p)} = v_{(0)}$.

A labeled p-tree is called a p-mobile if the labels satisfy the following additional condition:

(c) $\ell_v \ge 1$ for each $v \in \tau^\circ$.

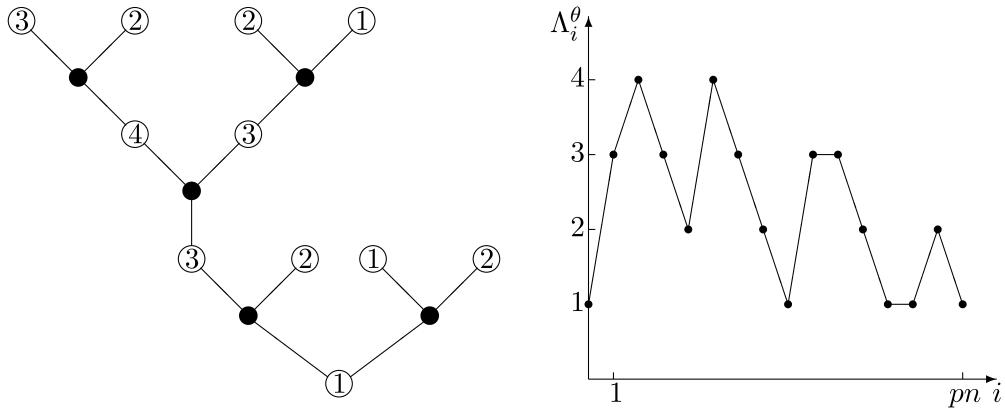

Figure 3. A 3-mobile θ with 5 black vertices and the associated spatial contour function.

The left side of Fig.3 gives an example of a p-mobile with $p = 3$. Condition (b) above means that if one lists the white vertices adjacent to a given black vertex in clockwise order, the labels of these vertices can decrease by at most one at each step.

Let τ be a p-tree with n black vertices and let $k = \#\tau - 1 = pn$. The depth-first search sequence of τ is the sequence $w_0, w_1, \dots, w_{2k}$ of vertices of τ which is obtained by induction as follows. First $w_0 = \emptyset$, and then for every $i \in \{0, \dots, 2k-1\}$, w_{i+1} is either the first child of w_i that has not yet appeared in the sequence $w_0, \dots, w_i$, or the parent of w_i if all children of w_i already appear in the sequence $w_0, \dots, w_i$. It is easy to verify that $w_{2k} = \emptyset$ and that all vertices of τ appear in the sequence $w_0, w_1, \dots, w_{2k}$ (of course some of them appear more than once).

Vertices w_i are white when i is even and black when i is odd. The contour sequence of τ° is by definition the sequence $v_0, \dots, v_k$ defined by $v_i = w_{2i}$ for every $i \in \{0, 1, \dots, k\}$.

Now let $\theta = (\tau, (\ell_v)_{v\in\tau^\circ})$ be a p-mobile with n black vertices. As previously, denote the contour sequence of τ° by $v_0, v_1, \dots, v_{pn}$. Suppose that the tree τ is drawn in the plane as pictured on Fig.4 and add an extra vertex ∂. We associate with θ a rooted $2p$-angulation M with n faces, whose set of vertices is

$$V(M) = \tau^\circ \cup \{\partial\}$$

and whose edges are obtained by the following device: For every $i \in \{0, 1, \dots, pn-1\}$,

- if $\ell_{v_i} = 1$, draw an edge between v_i and ∂ ;
- if $\ell_{v_i} \ge 2$, draw an edge between v_i and v_j, where j is the first index in the sequence $i+1, i+2, \dots, pn$ such that $\ell_{v_j} = \ell_{v_i} - 1$.

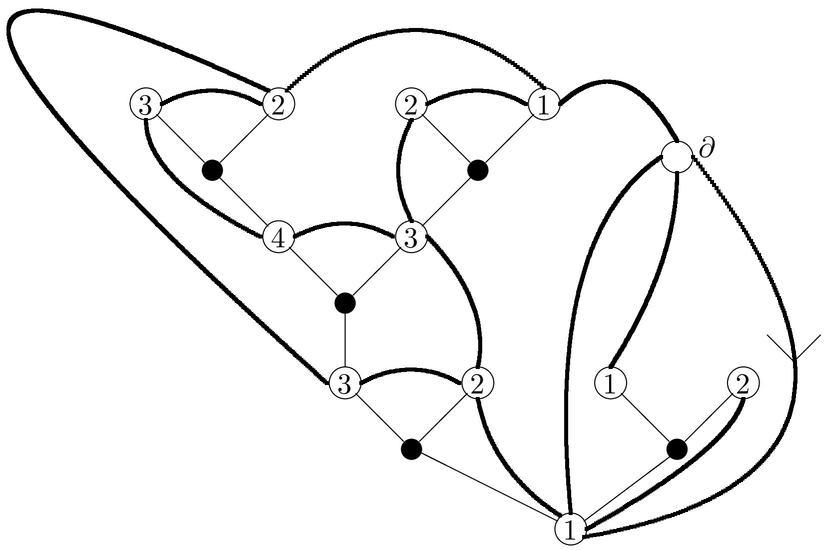

Figure 4. The Bouttier-Di Francesco-Guitter bijection: A rooted 3-mobile with 5 black vertices and the associated rooted 6-angulation with 5 faces. The root edge of the map is the oriented edge at the right end of the figure.

Notice that $v_{pn} = v_0 = \emptyset$ and $\ell_\emptyset = 1$, and that condition (b) in the definition of a p-tree entails that $\ell_{v_{i+1}} \geq \ell_{v_i} - 1$ for every $i \in \{0, 1, \ldots, pn-1\}$. This ensures that whenever $\ell_{v_i} \geq 2$ there is at least one vertex among $v_{i+1}, v_{i+2}, \ldots, v_{pn}$ with label $\ell_{v_i} - 1$. The construction can be made in such a way that edges do not intersect, except possibly at their endpoints: For every vertex v, each index i such that $v_i = v$ corresponds to a "corner" of v, and the associated edge starts from this corner. We refer to Section 2 of Bouttier et al [BDG] for a more detailed description.

The resulting planar map M is a $2p$-angulation, which is rooted at the oriented edge between ∂ and $v_0 = \emptyset$, corresponding to $i = 0$ in the previous construction. Each black vertex of τ is associated with a face of the map M. See Fig.4 for the 6-angulation associated with the 3-mobile of Fig.3.

The preceding construction yields a bijection between the set $\mathbb{T}_n^p$ of all p-mobiles with n black vertices and the set $\mathbb{M}_n^{2p}$. This is the Bouttier-Di Francesco-Guitter bijection [BDG], called the BDG bijection in what follows.

Furthermore, this bijection enjoys the following remarkable property, which is crucial for our purposes: The graph distance in M between the root vertex ∂ and another vertex $v \in \tau^\circ$ is equal to ℓ_v. Hence knowing the labels in the tree θ already gives a lot of information about distances in the map M.

In view of our applications, it will be convenient to code a p-mobile, or more generally a labeled p-tree, by a pair a discrete functions. The *contour function* of τ° (or of θ) is the discrete sequence $C_0^{\tau^\circ}, C_1^{\tau^\circ}, \ldots, C_{pn}^{\tau^\circ}$ defined by

$$C_i^{\tau^\circ} = \frac{1}{2}|v_i| \ , \text{ for every } 0 \leq i \leq pn.$$

See Fig.2 for an example with $p = n = 3$. It is easy to verify that the contour function determines τ°, which in turn determines the p-tree τ uniquely. We also introduce the *spatial contour function* of $\theta = (\tau, (\ell_v)_{v \in \tau^\circ})$, which is the discrete sequence $(\Lambda_0^\theta, \Lambda_1^\theta, \ldots, \Lambda_{pn}^\theta)$ defined by

$$\Lambda_i^\theta = \ell_{v_i} \ , \text{ for every } 0 \leq i \leq pn.$$

From property (b) of the labels and the definition of the contour sequence, it is clear that $\Lambda_{i+1}^\theta \geq \Lambda_i^\theta - 1$ for every $0 \leq i \leq pn - 1$ (cf Fig.3). The pair $(C^{\tau^\circ}, \Lambda^\theta)$ determines the labeled p-tree θ uniquely.

3 Scaling limits of trees

3.1 Plane trees

Our goal is to study the scaling limits of the labeled trees that appeared in the bijections with maps. We will start with the simpler problem of obtaining the scaling limit of plane trees. We first need to recall the definition of an $\mathbb{R}$-tree.

A metric space $(\mathcal{T}, d)$ is an $\mathbb{R}$-tree if the following two properties hold for every $a, b \in \mathcal{T}$.

(a) There is a unique isometric map $f_{a,b}$ from $[0, d(a,b)]$ into $\mathcal{T}$ such that $f_{a,b}(0) = a$ and $f_{a,b}(d(a,b)) = b$.

(b) If q is a continuous injective map from $[0,1]$ into $\mathcal{T}$, such that $q(0) = a$ and $q(1) = b$, we have

$$q([0,1]) = f_{a,b}([0, d(a,b)]).$$

A rooted $\mathbb{R}$-ree is an $\mathbb{R}$-tree $(\mathcal{T}, d)$ with a distinguished vertex $\rho = \rho(\mathcal{T})$ called the root.

Informally, one should think of a (compact) $\mathbb{R}$-tree as a connected union of line segments in the plane with no loops. For any two points a and b in the tree, there is a unique arc going from a to b in the tree, which is isometric to a line segment.

The multiplicity of a point a of $\mathcal{T}$ is the number of connected components of $\mathcal{T}\backslash\{a\}$. The point a is called a leaf if its multiplicity is one, and a branching point if its multiplicity is at least 3. We will be interested in compact $\mathbb{R}$-trees. Even for such trees, there can be (countably) infinitely many branching points and uncountably many leaves. This will indeed be the case for the random $\mathbb{R}$-trees that we will introduce.

We turn to the construction of (rooted) $\mathbb{R}$-trees from their contour functions. This is a continuous analogue of the well-known coding of plane trees by Dyck paths. Let $g : [0,1] \longrightarrow \mathbb{R}_+$ be a nonnegative continuous function such that $g(0) = g(1) = 0$. We will explain how to associate with g a compact $\mathbb{R}$-tree $(\mathcal{T}_g, d_g)$.

For every $s, t \in [0,1]$, we set

$$m_g(s,t) = \inf_{r \in [s \wedge t, s \vee t]} g(r),$$

and

$$d_g(s,t) = g(s) + g(t) - 2m_g(s,t).$$

It is easy to verify that d_g is a pseudo-metric on $[0,1]$. As usual, we introduce the equivalence relation $s \sim_g t$ if and only if $d_g(s,t) = 0$ (or equivalently if and only if $g(s) = g(t) = m_g(s,t)$). The function d_g induces a distance on the quotient space $\mathcal{T}_g := [0,1] \,/ \sim_g$, and we keep the notation d_g for this distance. We denote by $p_g : [0,1] \longrightarrow \mathcal{T}_g$ the canonical projection. Clearly p_g is continuous (when $[0,1]$ is equipped with the Euclidean metric and $\mathcal{T}_g$ with the metric d_g), and therefore $\mathcal{T}_g = p_g([0,1])$ is a compact metric space.

By Theorem 2.1 of [DL], the metric space $(\mathcal{T}_g, d_g)$ is a compact $\mathbb{R}$-tree. Furthermore the mapping $g \longrightarrow \mathcal{T}_g$ is continuous with respect to the Gromov-Hausdorff distance, if the set of continuous functions g is equipped with the supremum distance. We will always view $(\mathcal{T}_g, d_g)$ as a rooted $\mathbb{R}$-tree with root $\rho_g = p_g(0) = p_g(1)$. Note that $d_g(\rho_g, a) = g(s)$ if $a = p_g(s)$.

It is important to observe that the tree $\mathcal{T}_g$ inherits a "lexicographical order" from its coding by the function g. If $a, b \in \mathcal{T}_g$, the vertex a comes before b in lexicographical order if the smallest representative of a in $[0,1]$ is smaller than any representative of b in $[0,1]$.

By definition, the CRT is the random compact $\mathbb{R}$-tree $(\mathcal{T}_{\mathbf{e}}, d_{\mathbf{e}})$ coded in the previous sense by a normalized Brownian excursion $\mathbf{e} = (\mathbf{e}_t)_{0 \le t \le 1}$ (recall that a normalized Brownian excursion is a linear

Brownian motion over the time interval $[0,1]$, conditioned to start and to end at the origin, and to remain positive over the interval $(0,1)$). The CRT appears as the scaling limit of plane trees, as shown by the following theorem, which is a reformulation of a result in Aldous [A2].

Theorem 3.1 *For every $n \geq 1$, let τ_n be a random tree that is uniformly distributed over the set of all plane trees with n edges, and denote the graph distance on τ_n by d_{gr}. Then, the rescaled trees*

$$(\tau_n, (2n)^{-1/2} d_{gr})$$

converge in distribution towards the CRT, in the Gromov-Hausdorff sense.

There are in fact many other classes of random discrete trees for which the scaling limit is the CRT. For instance, it is not hard to see that this holds for random trees that are uniformly distributed over the set of all p-trees with n edges (considering only those values of n for which this set is not empty). The latter fact is an immediate consequence of the convergence of first components in Proposition 3.2 below.

3.2 Labeled trees and mobiles

In view of our applications to random planar maps, we need to understand the scaling limit of the p-mobiles of Section 2. We start with the simpler case of labeled p-trees.

For technical reasons, it is more convenient to deal with convergence of the coding functions rather than with convergence of the trees themselves. We first introduce the random functions that will appear in the limit. Let g be as above a continuous function from $[0,1]$ into $\mathbb{R}_+$ such that $g(0) = g(1) = 0$. We can consider the centered Gaussian process $(W^g_t)_{t\in[0,1]}$ whose distribution is characterized by the covariance function

$$\mathrm{cov}(W^g_s, W^g_t) = m_g(s,t),$$

for every $s,t \in [0,1]$. Note that $E[(W^g_s - W^g_t)^2] = d_g(s,t)$. The process $(W^g_s)_{s\in[0,1]}$ is called the Brownian snake driven by g. In the usual terminology, it is in fact the "head of the snake" rather than the snake itself – See [L1] for more information about Brownian snakes.

Under mild regularity assumptions on g, which will be satisfied in our applications, one can construct $(W^g_s)_{s\in[0,1]}$ so that it has continuous sample paths. Then the property $E[(W^g_s - W^g_t)^2] = d_g(s,t)$ implies that a.s. for every $s \in [0,1]$, W^g_s only depends on the equivalence class of s in the quotient $\mathcal{T}_g = [0,1]/\sim_g$. So we can find a process $Z^g = (Z^g_a)_{a\in\mathcal{T}_g}$ such that $Z^g_a = W^g_t$ whenever $a = p_g(t)$. The process Z^g should be interpreted as Brownian motion indexed by $\mathcal{T}_g$, which was briefly discussed in Section 1.

As in the previous subsection, we then randomize g. Precisely, we let $\mathbf{e} = (\mathbf{e}_s)_{s\in[0,1]}$ be as above a normalized Brownian excursion and we define a random process $(W_s)_{s\in[0,1]}$ such that, conditionally given $\mathbf{e}$, $(W_s)_{s\in[0,1]}$ is distributed as the Brownian snake driven by $\mathbf{e}$. We may again define "labels" $(Z_a)_{a\in\mathcal{T}_\mathbf{e}}$ by requiring that $Z_a = W_t$ whenever $a = p_\mathbf{e}(t)$.

We can now state a first result corresponding to the scaling limit of labeled p-trees. To simplify notation, we set

$$\lambda_p = \frac{1}{2}\sqrt{\frac{p}{p-1}}\;, \quad \kappa_p = \Big(\frac{9}{4p(p-1)}\Big)^{1/4}.$$

Proposition 3.2 *For every $n \geq 1$, let $(\tau_n, (\ell^n_v)_{v\in\tau^\circ_n})$ be uniformly distributed over the set of all labeled p-trees with n edges, and let C^n and Λ^n be respectively the contour function and the spatial contour function of $(\tau_n, (\ell^n_v)_{v\in\tau^\circ_n})$. Then,*

$$\Big(\lambda_p\, n^{-1/2}\, C^n_{[pnt]}, \kappa_p\, n^{-1/4} \Lambda^n_{[pnt]}\Big)_{0\leq t\leq 1} \xrightarrow[n\to\infty]{(d)} (\mathbf{e}_t, W_t)_{0\leq t\leq 1}.$$

This proposition is a special case of results proved in [MMi]. The convergence of the processes $(n^{-1/2}\,C^n_{[pnt]})_{t\in[0,1]}$ towards the Brownian excursion is essentially a variant of Theorem 3.1 (or rather of the formulation of this theorem in terms of contour functions, as in [A2]). The convergence of labels is related to general results about convergence of "discrete snakes" towards the Brownian snake, which are proved in [JM].

Of course the previous proposition is not sufficient for our purposes, since we are interested in p-mobiles and not in labelled p-trees. This means that we need to take the positivity constraint of labels (property (c) of the definition) into account. At an intuitive level, one may guess that this positivity constraint leads to considering the limiting pair of Proposition 3.2 conditioned on the event $\{W_s \geq 0$ for every $s \in [0,1]\}$. This conditioning however requires some care, since the conditioning event clearly has probability zero.

According to [LW], this conditioned pair, which we denote by $(\overline{\mathbf{e}}_t, \overline{W}_t)_{t\in[0,1]}$, can be constructed as follows. If s_* denotes the (almost surely unique) time in $[0,1]$ such that $W_{s_*} = \min\{W_s : 0 \leq s \leq 1\}$, we set for every $t \in [0,1]$,

- $\overline{\mathbf{e}}_t = \mathbf{e}_{s_*} + \mathbf{e}_{s_*\oplus t} - 2\,m_{\mathbf{e}}(s_*, s_*\oplus t)$;
- $\overline{W}_t = W_{s_*\oplus t} - W_{s_*}$.

where $s_*\oplus t = s_*+t$ if $s_*+t \leq 1$ and $s_*\oplus t = s_*+t-1$ if $s_*+t > 1$. This definition is better understood in terms of trees. First note that $\overline{W}_t$ only depends on the equivalence class of t in $\mathcal{T}_{\overline{\mathbf{e}}} = [0,1]/\sim_{\overline{\mathbf{e}}}$, and thus we may construct the labels $(\overline{Z}_a)_{a\in\mathcal{T}_{\overline{\mathbf{e}}}}$ such that $\overline{Z}_a = \overline{W}_t$ if $a = p_{\overline{\mathbf{e}}}(t)$. Then, the tree $\mathcal{T}_{\overline{\mathbf{e}}}$ is canonically identified with the tree $\mathcal{T}_{\mathbf{e}}$ re-rooted at the vertex $p_{\mathbf{e}}(s_*)$ with minimum label (see Lemma 2.2 in [DL]), and, modulo this identification, we have $\overline{Z}_a = Z_a - \min\{Z_c : c \in \mathcal{T}_{\mathbf{e}}\}$, meaning that the original labels are shifted to become nonnegative.

With the preceding notation we can now state the analogue of Proposition 3.2 for p-mobiles, which is proved in [We].

Theorem 3.3 *For every $n \geq 1$, let $(\overline{\tau}_n, (\overline{\ell}^n_v)_{v\in\overline{\tau}^\circ_n})$ be uniformly distributed over the set of all p-mobiles with n edges, and let $\overline{C}^n$ and $\overline{\Lambda}^n$ be respectively the contour function and the spatial contour function of $(\overline{\tau}_n, (\overline{\ell}^n_v)_{v\in\overline{\tau}^\circ_n})$. Then,*

$$\Big(\lambda_p\, n^{-1/2}\,\overline{C}^n_{[pnt]}, \kappa_p\, n^{-1/4}\overline{\Lambda}^n_{[pnt]}\Big)_{0\leq t\leq 1} \xrightarrow[n\to\infty]{(d)} (\overline{\mathbf{e}}_t, \overline{W}_t)_{0\leq t\leq 1}.$$

The following corollary was obtained in Chassaing and Schaeffer [CS] in the case $p = 2$ of quadrangulations. The general case can be found in [We], but the same result in a slightly different setting had been derived earlier by Marckert and Miermont [MMi]. See also [Mi1] for extensions to planar maps that are not bipartite.

Corollary 3.4 *For every integer $n \geq 2$, let M_n be a random planar map that is uniformly distributed over the set $\mathbb{M}^{2p}_n$ of all rooted $2p$-angulations with n faces. Denote by ∂ the root vertex of M_n and let $R(M_n) = \max\{d_{gr}(\partial, v) : v \in V(M_n)\}$ be the radius of the map. Then,*

$$\kappa_p n^{-1/4} R(M_n) \xrightarrow[n\to\infty]{(d)} \max_{0\leq t\leq 1} W_t - \min_{0\leq t\leq 1} W_t.$$

The proof is immediate from Theorem 3.3. Indeed, we know that M_n may be constructed as the image of $(\overline{\tau}_n, (\overline{\ell}^n_v)_{v\in\overline{\tau}^\circ_n})$ under the BDG bijection. Then, we have

$$R(M_n) = \max_{v\in\overline{\tau}^\circ_n} \overline{\ell}^n_v = \max_{0\leq k\leq pn} \overline{\Lambda}^n_k.$$

On the other hand, Theorem 3.3 implies that

$$\kappa_p n^{-1/4} \max_{0\le k\le pn} \overline{\Lambda}^n_k \xrightarrow[n\to\infty]{(d)} \max_{0\le t\le 1} \overline{W}_t$$

and from the definition of $\overline{W}$ in terms of W, we have also

$$\max_{0\le t\le 1} \overline{W}_t = \max_{0\le t\le 1} W_t - \min_{0\le t\le 1} W_t.$$

Remark. Detailed information about the limiting distribution in Corollary 3.4 can be found in [De].

4 Convergence towards the Brownian map

We now turn to the discussion of the convergence (2) in the case of uniformly distributed $2p$-angulations. Our results will involve the random pair $(\overline{\mathbf{e}}, \overline{W})$ which was introduced at the end of the previous section. This should not come as a surprise since this pair appears in the scaling limit of large p-mobiles (Theorem 3.3), and we know that $2p$-angulations are coded by p-mobiles. To simplify notation, we write $\overline{\mathcal{T}} = \mathcal{T}_{\overline{\mathbf{e}}}$ for the tree coded by $\overline{\mathbf{e}}$, and $\overline{\rho}$ for the root of $\overline{\mathcal{T}}$. Also recall that $(\overline{Z}_a)_{a\in\overline{\mathcal{T}}}$ are the labels induced on $\overline{\mathcal{T}}$ by the process $\overline{W}$.

In the discrete setting of $2p$-angulations, vertices of the map (except the root) are in one-to-one correspondence with vertices of the coding tree. A naive guess would be that a similar property holds in the continuous setting. It turns out that this is not correct and that one needs to identify certain vertices of the continuous random tree $\overline{\mathcal{T}}$, which plays the same role as a p-tree in the discrete setting.

Let $s, t \in [0, 1]$. By definition,

$$s \simeq t \qquad \text{if and only if} \qquad \overline{W}_{s\wedge t} = \overline{W}_{s\vee t} = \min_{s\wedge t\le r\le s\vee t} \overline{W}_r.$$

In this way we obtain a random equivalence relation on $[0, 1]$. For $a, b \in \overline{\mathcal{T}}$, we then say that $a \approx b$ if and only if there exist a representative s of a in $[0, 1]$ and a representative t of b in $[0, 1]$ such that $s \simeq t$. It turns out that $\approx$ is also an equivalence relation on $\overline{\mathcal{T}}$, a.s. Informally, $a \approx b$ if and only if a and b have the same label ($\overline{Z}_a = \overline{Z}_b$), and when going from a to b in lexicographical order (or in reverse lexicographical order) around the tree, one encounters only vertices with greater label. The preceding definition of the equivalence relation $\approx$ can be seen to be equivalent to the more informal one given in Section 1, modulo the identification of the trees $\mathcal{T}_{\mathbf{e}}$ and $\overline{\mathcal{T}}$ up to re-rooting.

It is easy to understand why the equivalence relation $\approx$ should be relevant to our description of the scaling limit of random maps. Indeed consider two white vertices u and u' in a p-mobile $(\tau, (\ell_v)_{v\in\tau^\circ})$, and recall our notation $(v_0, v_1, \ldots)$ for the contour sequence of τ°. Then u and u' will be connected by an edge of the associated map if and only if we can write $\{u, u'\} = \{v_i, v_j\}$, with $i < j$, in such a way that

(a) $\ell_{v_j} = \ell_{v_i} - 1\,,$

(b) $\ell_{v_k} \ge \ell_{v_i}$ for all $k \in \{i, i+1, \ldots, j-1\}$.

Note that this may occur for vertices that are far away from each other in the tree, and that such pairs of vertices should be identified in the scaling limit of maps. Recalling that the process $\overline{W}$ is the scaling limit of the spatial contour sequence of p-mobiles (Theorem 3.3), we see that our definition of the equivalence relation $\approx$ is just a continuous analogue of properties (a) and (b).

We denote by $\mathbf{m}_\infty$ the quotient space $\mathbf{m}_\infty = \overline{\mathcal{T}}/\approx$. Notice that $\overline{Z}_a = \overline{Z}_b$ if $a \approx b$, so that the labels $\overline{Z}_x$ can be defined with no ambiguity for every $x \in \mathbf{m}_\infty$.

The following theorem is the main result of [L2].

Theorem 4.1 *For every integer $n \geq 2$, let M_n be a random planar map that is uniformly distributed over the set $\mathbb{M}_n^{2p}$ of all rooted $2p$-angulations with n faces. From every strictly increasing sequence of positive integers, we can extract a subsequence along which the following convergence holds:*

$$\left(V(M_n), \kappa_p n^{-1/4} d_{gr}\right) \xrightarrow[n\to\infty]{(d)} (\mathbf{m}_\infty, D),$$

where D is a random distance on $\mathbf{m}_\infty$, that induces the quotient topology on this space. Furthermore, for every $x \in \mathbf{m}_\infty$,

$$D(\overline{\rho}, x) = \overline{Z}_x. \tag{3}$$

Remark. In (3), the root $\overline{\rho}$ of $\overline{\mathcal{T}}$ is identified with its equivalence class in $\mathbf{m}_\infty$, which is a singleton. We will do this identification systematically, and $\overline{\rho}$ thus appears as a distinguished point of $\mathbf{m}_\infty$. The property of invariance under uniform re-rooting (Theorem 8.1 in [L3]) however shows that this distinguished point plays no special role.

The limiting random metric space $(\mathbf{m}_\infty, D)$ is called the Brownian map. This terminology is slightly abusive because, as we already explained in Section 1, the random distance D may depend on the choice of p and of the subsequence in the theorem. One conjectures that D does not depend on these choices and that the same limiting random metric space appears as the scaling limit of more general random planar maps, such as triangulations for instance. This conjecture justifies that the name Brownian map is used in this work to denote one of the possible limits arising in Theorem 4.1. The results that are stated in Sections 5 and 6 below hold for any of these limits.

Sketch of proof. The proof of Theorem 4.1 consists of two main steps. The first one is a compactness argument showing that sequential limits of $(V(M_n), \kappa_p n^{-1/4} d_{gr})$ exist, and that any such limit can be written as a quotient space of $\overline{\mathcal{T}}$. The second step, which is the hard part of the proof, is the identification of the equivalence relation corresponding to this quotient. Let us briefly sketch the compactness argument of the first step.

The random map M_n is the image under the BDG bijection of a p-mobile $(\overline{\tau}_n, (\overline{\ell}^n_v)_{v\in\overline{\tau}_n^\circ})$, and we can thus identify $V(M_n)$ with $\overline{\tau}_n^\circ \cup \{\partial\}$. We write $v_0^n, v_1^n, \ldots, v_{pn}^n$ for the contour sequence of the tree $\overline{\tau}_n^\circ$. As in Theorem 3.3, let $\overline{\Lambda}^n$ be the spatial contour function of $(\overline{\tau}_n, (\overline{\ell}^n_v)_{v\in\overline{\tau}_n^\circ})$, so that $\overline{\Lambda}^n_i = \overline{\ell}^n_{v_i^n}$ by definition. For every $i, j \in \{0, 1, \ldots, pn\}$, set

$$d_n(i,j) = d_{gr}(v_i^n, v_j^n).$$

Lemma 4.2 *For every $i, j \in \{0, 1, \ldots, pn\}$,*

$$d_n(i,j) \leq d_n^\circ(i,j) := \overline{\Lambda}^n_i + \overline{\Lambda}^n_j - 2 \min_{i\wedge j \leq k \leq i \vee j} \overline{\Lambda}^n_k + 2.$$

This lemma essentially follows from the properties of the BDG bijection. Note that we can construct a discrete geodesic from v_i^n to ∂ via the following procedure. We first look for the first index $i_1 > i$ such that the vertex $v_{i_1}^n$ has label $\overline{\ell}^n_i - 1$. By construction $d_n(i, i_1) = 1$. Similarly, we then look for the first index $i_2 > i_1$ such that $v_{i_2}^n$ has label $\overline{\ell}^n_i - 2$, and we have $d_n(i_1, i_2) = 1$. We continue this way until we arrive at a vertex with label 1, which is connected to ∂. We can similarly construct a discrete geodesic from v_j^n to ∂. However the two discrete geodesics we have obtained coincide for vertices whose distance from the root is less than or equal to

$$\min_{i\wedge j \leq k \leq i\vee j} \overline{\Lambda}^n_k - 1.$$

The bound of the lemma follows.

We extend the definition of $d_n(s,t)$ and $d_n^\circ(s,t)$ to noninteger values $s,t \in [0,pn]$ by linear interpolation. Next we use Theorem 3.3, which gives

$$\Big(\kappa_p n^{-1/4}\, d_n^\circ(pns,pnt)\Big)_{0\le s,t\le 1} \xrightarrow[n\to\infty]{(d)} (D^\circ(s,t))_{0\le s,t\le 1}$$

where

$$D^\circ(s,t) = \overline{W}_s + \overline{W}_t - 2 \min_{s\wedge t\le r\le s\vee t} \overline{W}_r.$$

This implies that we can find two sequences ε_k, δ_k of positive reals converging to 0, such that, with a probability tending to 1 as $k\to\infty$, we have for every $n\ge 2$, and $s,t\in[0,1]$,

$$|t-s| < \delta_k \Rightarrow n^{-1/4} d_n^\circ(pns,pnt) < \varepsilon_k \Rightarrow n^{-1/4} d_n(pns,pnt) < \varepsilon_k.$$

It follows that with probability close to one when k is large, each of the metric spaces $(V(M_n), n^{-1/4}d_n)$ can be covered by at most $[\frac{1}{\delta_k}]+1$ balls of radius ε_k. By standard compactness criteria for the Gromov-Hausdorff convergence, this gives the tightness of the sequence of distributions of the metric spaces $(V(M_n), \kappa_p n^{-1/4} d_{gr})$.

Remarks. (i) The preceding argument also yields a useful bound on the limiting distance D in Theorem 4.1. We denote by $\mathbf{p} = \Pi \circ p_{\overline{\mathbf{e}}}$ the composition of the projection $p_{\overline{\mathbf{e}}} : [0,1] \to \overline{\mathcal{T}}$ and the canonical projection $\Pi : \overline{\mathcal{T}} \to \mathbf{m}_\infty$. For every $x,y \in \mathbf{m}_\infty$, set

$$D^\circ(x,y) = \inf\{D^\circ(s,t) : s,t\in[0,1], \mathbf{p}(s)=a, \mathbf{p}(t)=b\}.$$

Then, for every $s,y\in\mathbf{m}_\infty$,

$$D(x,y) \le D^\circ(x,y). \tag{4}$$

This follows as a consequence of Lemma 4.2.
(ii) One may ask whether the quotient $\mathbf{m}_\infty = \overline{\mathcal{T}}/\approx$ involves identifying many pairs of points. In some sense, it does not: A typical equivalence class for $\approx$ is a singleton, and non-trivial equivalence classes can contain at most three points (there are only countably many classes containing three points). It is also true that if a is a point of $\overline{\mathcal{T}}$ that is not a leaf, then the equivalence class of a is a singleton. Thus only certain leaves of $\overline{\mathcal{T}}$ are identified with certain other leaves. In a sense, getting from the CRT to the Brownian map requires identifying relatively few pairs of points. Still these identifications drastically change the topology of the space, as we will see below (Theorem 5.2).

Theorem 4.1 leads to the obvious problem of characterizing the random distance D, which would imply that there is no need for taking a subsequence in the theorem. Provided that the characterization does not depend on p, this would also prove that the limiting space does not depend on the choice of p. Let us formulate a conjecture for D from [L2] (see also [MMo]).

Conjecture. For every $x,y\in\mathbf{m}_\infty$, $D(x,y) = \inf\{\sum_{i=1}^k D^\circ(x_{i-1},x_i)\}$ where the infimum is over all choices of the integer k and the sequence $x_0,x_1,\ldots,x_k\in\mathbf{m}_\infty$ such that $x_0=x$ and $x_k=y$.

Even if the preceding questions are still open, we will see in the next sections that much can be said about the Brownian map, and that the properties of this limiting space already have interesting consequences for large random planar maps.

5 Two theorems about the Brownian map

In this section and the next one, the Brownian map $(\mathbf{m}_\infty, D)$ is one of the possible limits arising in the convergence of Theorem 4.1.

Theorem 5.1 *The Hausdorff dimension of the Brownian map is*

$$\dim(\mathbf{m}_\infty, D) = 4 \qquad a.s.$$

The bound $\dim(\mathbf{m}_\infty, D) \leq 4$ is very easy to derive from our construction. Indeed, the bound (4) almost immediately implies that the projection $\mathbf{p} : [0,1] \longrightarrow \mathbf{m}_\infty$ is Hölder continuous with exponent $1/4$, which gives the desired upper bound. See [L2] for a proof of the corresponding lower bound.

Note that the topological type of the Brownian map is completely characterized in Theorem 4.1: The metric D induces the quotient topology on $\mathbf{m}_\infty$. The following theorem, which is the main result of [LP], identifies this topological type.

Theorem 5.2 *The space $(\mathbf{m}_\infty, D)$ is almost surely homeomorphic to the two-dimensional sphere $\mathbb{S}^2$.*

The proof of Theorem 5.2 is based on the expression of the Brownian map as a quotient space, and on a classical theorem of Moore giving sufficient conditions for a quotient space of the sphere to be still homeomorphic to the sphere. An alternative approach has been given by Miermont [Mi3].

Theorem 5.2 implies that with a probability close to one when n is large, a typical $2p$-angulation cannot have a separating cycle of length small in comparison with the diameter of the map, and such that both sides of the cycle have a "macroscopic" size. Indeed the existence of such "bottlenecks" in the map would imply that the scaling limit is a topological space which can be disconnected by removing a single point, and this is of course not true for the sphere. We state the previous observation more precisely, recalling that the diameter of a typical $2p$-angulation with n faces is of order $n^{1/4}$ (cf Corollary 3.4).

Corollary 5.3 *For every integer $n \geq 2$, let M_n be a random planar map that is uniformly distributed over the set $\mathbb{M}_n^{2p}$ of all rooted $2p$-angulations with n faces. Let $\alpha > 0$ and let $\psi : \mathbb{N} \longrightarrow \mathbb{R}_+$ be a function such that $\psi(n) = o(n^{1/4})$ as $n \to \infty$. Then, with a probability tending to 1 as $n \to \infty$, there exists no injective cycle C of the map M_n with length less than $\psi(n)$, such that the set of vertices that lie in either connected component of the complement of C in the sphere has diameter at least $\alpha n^{1/4}$.*

6 Geodesics in the Brownian map

Our goal in this section is to discuss geodesics in the Brownian map, and then to apply this discussion to asymptotic properties of large planar maps. We rely on the recent paper [L3]. See Miermont [Mi2] and Bouttier and Guitter [BG1, BG3] for other interesting results about geodesics in large random planar maps.

We start by recalling a general definition. If (E, δ) is a compact metric space and $x, y \in E$, a geodesic or shortest path from x to y is a continuous path $\gamma = (\gamma(t))_{0 \leq t \leq \delta(x,y)}$ such that $\gamma(0) = x$, $\gamma(\delta(x,y)) = y$ and $\delta(\gamma(s), \gamma(t)) = |t - s|$ for every $s, t \in [0, \delta(x,y)]$. The space (E, δ) is then called a geodesic space if any two points in E are connected by (at least) one geodesic. From the fact that Gromov-Hausdorff limits of geodesic spaces are geodesic spaces (see [BBI], Theorem 7.5.1), one gets that $(\mathbf{m}_\infty, D)$ is almost surely a geodesic space. We will determine explicitly the geodesics between the root $\overline{\rho}$ of $\mathbf{m}_\infty$ and an arbitrary point of $\mathbf{m}_\infty$.

We define the skeleton $\mathrm{Sk}(\overline{\mathcal{T}})$ as the set of all points of the tree $\overline{\mathcal{T}}$ that are not leaves (equivalently these are the points whose removal disconnects the tree). One can verify that the restriction of the projection $\Pi : \overline{\mathcal{T}} \longrightarrow \mathbf{m}_\infty$ to $\mathrm{Sk}(\overline{\mathcal{T}})$ is a homeomorphism. Moreover, since Π is Hölder continuous with exponent $1/2 - \varepsilon$ for every $\varepsilon > 0$ (essentially by the bound (4)), and $\mathrm{Sk}(\overline{\mathcal{T}})$ has dimension one, the Hausdorff dimension of $\Pi(\mathrm{Sk}(\overline{\mathcal{T}}))$ is less than or equal to 2. One can indeed prove that $\dim \Pi(\mathrm{Sk}(\overline{\mathcal{T}})) = 2$.

We write $\mathrm{Skel}_\infty = \Pi(\mathrm{Sk}(\overline{\mathcal{T}}))$ to simplify notation. Since the Hausdorff dimension of $\mathbf{m}_\infty$ is equal to 4 almost surely (Theorem 5.1), the set Skel_∞ is a relatively small subset of $\mathbf{m}_\infty$. The set Skel_∞ is dense in $\mathbf{m}_\infty$ and from the previous observations it is homeomorphic to a non-compact $\mathbb{R}$-tree. Moreover, for every $x \in \mathrm{Skel}_\infty$, the set $\mathrm{Skel}_\infty \backslash \{x\}$ is not connected.

The following theorem provides a nice geometric interpretation of the set Skel_∞.

Theorem 6.1 *The following properties hold almost surely. For every $x \in \mathbf{m}_\infty \backslash \mathrm{Skel}_\infty$, there is a unique geodesic from $\overline{\rho}$ to x. On the other hand, for every $x \in \mathrm{Skel}_\infty$, the number of distinct geodesics from $\overline{\rho}$ to x is equal to the number of connected components of $\mathrm{Skel}_\infty \backslash \{x\}$. In particular, the maximal number of distinct geodesics from $\overline{\rho}$ to a point of $\mathbf{m}_\infty$ is equal to 3, and there are countably many points for which this number is attained.*

Remark. The invariance of the distribution of the Brownian map under uniform re-rooting (see Section 8 in [L3]) shows that results analogous to Theorem 6.1 hold if one replaces the root $\overline{\rho}$ by a point z distributed uniformly over $\mathbf{m}_\infty$. Here the word "uniformly" refers to the volume measure λ on $\mathbf{m}_\infty$, which is the image of Lebesgue measure on $[0,1]$ under the projection $\mathbf{p} = \Pi \circ p_{\overline{\mathbf{e}}}$.

Theorem 6.1 opens a new perspective on our construction of the Brownian map $(\mathbf{m}_\infty, D)$ as a quotient space of the random tree $\overline{\mathcal{T}}$ (at first, this construction may appear artificial, even though it is a continuous counterpart of the BDG bijection). Indeed, Theorem 6.1 shows that the skeleton of $\overline{\mathcal{T}}$, or rather its homeomorphic image under the canonical projection Π, has an intrinsic geometric meaning: It exactly corresponds to the cut locus of $\mathbf{m}_\infty$ relative to the root $\overline{\rho}$, provided we define this cut locus as the set of all points that are connected to $\overline{\rho}$ by at least two distinct geodesics (this definition of the cut locus is slightly different from the one that appears in Riemannian geometry). Remarkably enough, the assertions of Theorem 6.1 parallel the known results in the setting of differential geometry, which go back to Poincaré [Po] and Myers [My].

To give a hint of the proof of Theorem 6.1, let us introduce the notion of a simple geodesic. Let $x \in \mathbf{m}_\infty$, let $a \in \overline{\mathcal{T}}$ be such that $\Pi(a) = x$, and let $t \in [0,1]$ be such that $p_{\overline{\mathbf{e}}}(t) = a$. Recall that we have $D(\overline{\rho}, x) = \overline{Z}_x = \overline{Z}_a = \overline{W}_t$. For every $r \in [0, D(\overline{\rho}, x)]$, set

$$\gamma_t(r) = \sup\{s \in [0,t] : \overline{W}_s = r\}.$$

By a continuity argument, $\gamma_t(r)$ is well defined and $\overline{W}_{\gamma_t(r)} = r$. Set $\Gamma_t(r) = \mathbf{p}(\gamma_t(r))$ for every $r \in [0, D(\overline{\rho}, x)]$. We have

$$D(\overline{\rho}, \Gamma_t(r)) = \overline{W}_{\gamma_t(r)} = r.$$

On the other hand, if $0 \le r \le r' \le t$,

$$\min_{\gamma_t(r) \le s \le \gamma_t(r')} \overline{W}_s = r$$

by the definition of $\gamma_t(r)$. The bound (4) now gives

$$D(\Gamma_t(r), \Gamma_t(r')) \le r' - r.$$

Since the reverse bound is just the triangle inequality, we have obtained that

$$D(\Gamma_t(r), \Gamma_t(r')) = r' - r$$

for every $0 \le r \le r' \le D(\overline{\rho}, x)$. Clearly $\Gamma_t(0) = \overline{\rho}$ and $\Gamma_t(D(\overline{\rho}, x)) = \mathbf{p}(t) = x$. Thus we have proved that the path $(\Gamma_t(r))_{0 \le r \le D(\overline{\rho}, x)}$ is a geodesic from $\overline{\rho}$ to x. Such a geodesic is called a *simple geodesic*.

Remark. The preceding construction of simple geodesics is just a continuous analogue of the construction of discrete geodesics that was outlined in the proof of Lemma 4.2.

The main difficulty in the proof of Theorem 6.1 is to check that *all* geodesics from the root are simple geodesics. From this, the various statements of Theorem 6.1 follow by counting how many simple geodesics can exist for a given point $x \in \mathbf{m}_\infty$. In order that there exist more than one, two situations can occur:

- There exist several values of a such that $\Pi(a) = x$ (these values thus lie in the same equivalence class for $\approx$, and by a previous remark they are all leaves of $\overline{\mathcal{T}}$). However, essentially from the definition of $\approx$, one can check that the simple geodesics corresponding to these different values of a are the same.

- There is only one value of a such that $\Pi(a) = x$, but there are several values of $t \in [0,1]$ such that $p_{\overline{\mathbf{e}}}(t) = a$. This means that a belongs to the skeleton of $\overline{\mathcal{T}}$, and the number of values of t such that $p_{\overline{\mathbf{e}}}(t) = a$ is the multiplicity of a in $\overline{\mathcal{T}}$. In that case, one easily checks that the simple geodesics Γ_t, for all t such that $p_{\overline{\mathbf{e}}}(t) = a$, are distinct.

The statement of Theorem 6.1 is a consequence of this discussion. Note that the number of connected components of $\mathrm{Skel}_\infty \backslash \{x\}$ is at most 3 because $\overline{\mathcal{T}}$, or equivalently the CRT, has only binary branching points, as a consequence of the fact that Brownian minima are distinct.

The next corollary gives a surprising confluence property for geodesics starting from the root.

Corollary 6.2 *Almost surely, for every $\eta > 0$, there exists $\alpha \in]0, \eta[$ such that the following holds. Let $x, x' \in \mathbf{m}_\infty$ such that $D(\overline{\rho}, x) \geq \eta$ and $D(\overline{\rho}, x') \geq \eta$, and let ω, respectively ω', be a geodesic from $\overline{\rho}$ to x, resp. from $\overline{\rho}$ to x'. Then, $\omega(t) = \omega'(t)$ for every $t \in [0, \alpha]$.*

Since we know that all geodesics from the root are simple geodesics, this corollary easily follows from the fact that two simple geodesics must coincide near the root. We indeed used a similar property in the discrete setting in the proof of Lemma 4.2.

To conclude this section, let us give two applications of the previous results to geodesics in large planar maps. In the discrete setting, there is of course no hope to establish the uniqueness of geodesics between two vertices (see [BG1, BG3] for asymptotic results about the number of geodesics). Still it makes sense to deal with macroscopic uniqueness, meaning that any two geodesics will be close at an order that is small in comparison with the diameter of the map.

We recall that the random planar map M_n is uniform distributed over the set $\mathbb{M}_n^{2p}$ of all rooted $2p$-angulations with n faces, and that ∂ denotes the root vertex of M_n. For every $v \in V(M_n)$, we denote by $\mathrm{Geo}_n(\partial \to v)$ the set of all discrete geodesics from ∂ to v in the map M_n.

If γ, γ' are two discrete paths with the same length k, we set

$$d(\gamma, \gamma') = \max_{0 \leq i \leq k} d_{gr}(\gamma(i), \gamma'(i)).$$

Corollary 6.3 *Let $\varepsilon > 0$. Then,*

$$\frac{1}{n} \#\{v \in V(M_n) : \exists \gamma, \gamma' \in \mathrm{Geo}_n(\partial \to v),\ d(\gamma, \gamma') \geq \varepsilon n^{1/4}\} \underset{n\to\infty}{\longrightarrow} 0$$

in probability.

This means that for a typical vertex v in the map M_n, the discrete geodesic from ∂ to v is "macroscopically" unique. A stronger statement can be obtained by considering approximate geodesics, i.e. discrete paths from ∂ to v whose length is bounded above by $d_{gr}(\partial, v) + o(n^{1/4})$. Also note that a related uniqueness result has been obtained by Miermont in [Mi2].

Now what about exceptional vertices in the map M_n ? Does there exist vertices v such that there are several macroscopically different geodesics from ∂ to v ? The following corollary provides an answer to this question. Before giving the statement, we need to introduce another notation. For $v \in V(M_n)$, and $\varepsilon > 0$, we set

$$\mathrm{Mult}_\varepsilon(v) = \max\{k : \exists \gamma_1, \ldots, \gamma_k \in \mathrm{Geo}_n(\partial, v),\ d(\gamma_i, \gamma_j) \geq \varepsilon n^{1/4} \text{ if } i \neq j\}.$$

Corollary 6.4 *For every $\varepsilon > 0$,*

$$P[\exists v \in V(M_n) : \mathrm{Mult}_\varepsilon(v) \geq 4] \underset{n\to\infty}{\longrightarrow} 0\,.$$

However,

$$\lim_{\varepsilon\to 0} \Big(\liminf_{n\to\infty} P[\exists v \in V(M_n) : \mathrm{Mult}_\varepsilon(v) = 3] \Big) = 1\,.$$

Loosely speaking, there can be at most 3 macroscopically different geodesics from ∂ to an arbitrary vertex of M_n.

Remark. In the last two corollaries, the root vertex ∂ can be replaced by a vertex chosen uniformly at random in M_n.

References

[A1] Aldous, D. (1991) The continuum random tree I. *Ann. Probab.* **19**, 1-28.

[A2] Aldous, D. (1993) The continuum random tree III. *Ann. Probab.* **21**, 248-289.

[ADJ] Ambjørn, J., Durhuus, B., Jonsson, T. (1997) *Quantum geometry. A statistical field theory approach.* Cambridge Monographs on Mathematical Physics. Cambridge University Press, Cambridge.

[An] Angel, O. (2003) Growth and percolation on the uniform infinite planar triangulation. *Geom. Funct. Anal.* **3**, 935-974.

[AS] Angel, O., Schramm, O. (2003) Uniform infinite planar triangulations. *Comm. Math. Phys.* **241**, 191-213.

[Bo] Bouttier, J. (2005) *Physique statistique des surfaces aléatoires et combinatoire bijective des cartes planaires.* PhD thesis, Université Paris 6.

`http://tel.ccsd.cnrs.fr/documents/archives0/00/01/06/51/index.html`

[BDG] Bouttier, J., Di Francesco, P., Guitter, E. (2004) Planar maps as labeled mobiles. *Electronic J. Combinatorics* **11**, #R69.

[BG1] Bouttier, J., Guitter, E. (2008) Statistics of geodesics in large quadrangulations. *J. Phys. A* **41**, no. 14, 145001, 30pp.

[BG2] Bouttier, J., Guitter, E. (2008) The three-point function of planar quadrangulations. *J. Stat. Mech. Theory Exp.* no.7, P07020, 39pp.

[BG3] Bouttier, J., Guitter, E. (2008) Confluence of geodesic paths and separating loops in large planar quadrangulations. Preprint.

[BIP] Brézin, E., Itzykson, C., Parisi, G., Zuber, J.B. (1978) Planar diagrams. *Comm. Math. Phys.* **59**, 35-51.

[BBI] Burago, D., Burago, Y., Ivanov, S. (2001) *A course in metric geometry.* Graduate Studies in Mathematics, vol. 33. AMS, Boston.

[CS] Chassaing, P., Schaeffer, G. (2004) Random planar lattices and integrated superBrownian excursion. *Probab. Th. Rel. Fields* **128**, 161-212.

[CD] Chassaing, P., Durhuus, B. (2006) Local limit of labeled trees and expected volume growth in a random quadrangulation. *Ann. Probab.* **34**, 879-917.

[CV] Cori, R., Vauquelin, B. (1981) Planar maps are well labeled trees. *Canad. J. Math.* **33**, 1023-1042.

[De] Delmas, J.F. (2003) Computation of moments for the length of the one dimensional ISE support. *Electron. J. Probab.* **8**, no. 17, 15 pp.

[DL] Duquesne, T., Le Gall, J.F. (2005) Probabilistic and fractal aspects of Lévy trees. *Probab. Th. Rel. Fields* **131**, 553-603.

[Gr] Gromov, M. (2001) *Metric structures for Riemannian and non-Riemannian spaces.* Birkhäuser.

[tH] 't Hooft, G. (1974) A planar diagram theory for strong interactions. *Nucl. Phys. B* **72**, 461-473.

[JM] Janson, S., Marckert, J.F. (2005) Convergence of discrete snakes. *J. Theoret. Probability* **18**, 615-645.

[Kr] Krikun, M. (2006) Local structure of random quadrangulations. Preprint.
arXiv:math.PR/0512304

[LZ] Lando, S.K., Zvonkin, A.K. (2004) *Graphs on surfaces and their applications.* Vol. 141 of *Encyclopedia of Mathematical Sciences.* Springer, Berlin.

[L1] Le Gall, J.F. (1999) *Spatial branching processes, random snakes and partial differential equations.* Lectures in Mathematics ETH Zürich. Birkhäuser, Boston.

[L2] Le Gall, J.F. (2007) The topological structure of scaling limits of large planar maps. *Invent. Math.* **169**, 621-670

[L3] Le Gall, J.F. (2008) Geodesics in large planar maps and in the Brownian map. *Acta Math.*, to appear.

[LP] Le Gall, J.F., Paulin, F. (2008) Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere. *Geomet. Funct. Anal.* **18**, 893-918

[LW] Le Gall, J.F., Weill, M. (2006) Conditioned Brownian trees. *Ann. Inst. H. Poincaré, Probab. Stat.* **42**, 455-489.

[MMi] Marckert, J.F., Miermont, G. (2007) Invariance principles for random bipartite planar maps. *Ann. Probab.* **35**, 1642-1705.

[MMo] Marckert, J.F., Mokkadem, A. (2006) Limit of normalized quadrangulations. The Brownian map. *Ann. Probab.* **34**, 2144-2202.

[Mi1] Miermont, G. (2006) An invariance principle for random planar maps. In: Proceedings Fourth Colloquium on Mathematics and Computer Science CMCS'06 (Nancy, France), pp. 39-58 (electronic)

[Mi2] Miermont, G. (2007) Tessellations of random maps of arbitrary genus. *Ann. Sci. Ec. Norm. Supér.*, to appear.

[Mi3] Miermont, G. (2008) On the sphericity of scaling limits of random planar quadrangulations. *Electron. Comm. Probab.* **13**, 248-257.

[MW] Miermont, G., Weill, M. (2008) Radius and profile of random planar maps with faces of arbitrary degrees. *Electron. J. Probab.* **13**, 79-106.

[My] Myers, S.B. (1935) Connections between differential geometry and topology. I. Simply connected surfaces. *Duke Math. J.* **1**, 376-391.

[Po] Poincaré, H. (1905) Sur les lignes géodésiques des surfaces convexes. *Trans. Amer. Math. Soc.* **6**, 237-274.

[S] Schaeffer, G. (1998) *Conjugaison d'arbres et cartes combinatoires aléatoires.* PhD thesis, Université Bordeaux I.

`http://www.lix.polytechnique.fr/~schaeffe/Biblio/`

[Sc] Schramm, O. (2007) Conformally invariant scaling limits: an overview and a collection of problems. Proceedings of the International Congress of Mathematicians Madrid 2006, Vol.I, pp. 513-543.

[Tu] Tutte, W.T. (1963) A census of planar maps. *Canad. J. Math.* **15**, 249-271.

[We] Weill, M. (2007) Asymptotics for rooted planar maps and scaling limits of two-type Galton-Watson trees. *Electron. J. Probab.* **12**, 862-925.

Institut universitaire de France et Université Paris-Sud
Mathématiques, bâtiment 425
Centre d'Orsay, 91405 ORSAY Cedex, FRANCE