

Geodesics in Brownian surfaces (Brownian maps)

Jérémie Bettinelli

École Polytechnique Laboratoire d'Informatique (LIX), 91128 Palaiseau Cedex, France. E-mail: jeremie.bettinelli@normalesup.org;
url: <http://www.normalesup.org/~bettineli>

Received 8 October 2014; accepted 19 December 2014

Abstract. We define a class of metric spaces we call *Brownian surfaces*, arising as the scaling limits of random maps on general orientable surfaces with a boundary and we study the geodesics from a uniformly chosen random point. These metric spaces generalize the well-known *Brownian map* and our results generalize the properties shown by Le Gall on geodesics in the latter space. We use a different approach based on two ingredients: we first study typical geodesics and then all geodesics by an “entrapment” strategy. In particular, we give geometrical characterizations of some subsets of interest, in terms of geodesics, boundary points and concatenations of geodesics forming a loop that is not homotopic to 0.

Résumé. On définit une classe d'espaces métriques aléatoires que nous appelons *surfaces browniennes* : ces objets apparaissent comme limites d'échelle de cartes aléatoires sur des surfaces orientables à bord générales. Dans un second temps, on étudie les géodésiques émanant d'un point choisi uniformément au hasard. Les surfaces browniennes généralisent la fameuse *carte brownienne* et nos résultats généralisent les propriétés obtenues par Le Gall sur les géodésiques dans cet espace. On utilise une approche différente reposant sur deux ingrédients : on étudie d'abord les géodésiques aux points typiques et on attrape ensuite les autres points en les « encerclant » par de telles géodésiques. En particulier, on obtient des caractérisations géométriques de certains sous-ensembles d'intérêt en termes de géodésiques, points du bord et concaténations des géodésiques formant une boucle non homotope à 0.

MSC: 60F17; 60C05; 60D05; 57N05; 05C80; 05C12

Keywords: Brownian surfaces; Brownian map; Geodesics; Random maps; Scaling limits; Gromov–Hausdorff topology; Random metric spaces; Bijections

1. Introduction

The purpose of this work is twofold. We first construct a class of metric spaces we call *Brownian surfaces*, which arise as the scaling limit of random discrete maps on general surfaces. They naturally generalize the famous *Brownian map* [25,31], which corresponds to the case where the surface in question is the sphere.

In a second step, we study all the geodesics starting from a “uniformly” chosen random point in a fixed Brownian surface. Our method allows to recover the main result of [24] and provides analogous results for general Brownian surfaces. As the topology of the spaces we consider here is richer than the topology of the sphere, we find new geometrical characterizations in terms of boundary points and pairs of geodesics whose concatenation forms a loop that is not homotopic to 0.

Recall that a *surface with a boundary* is a nonempty Hausdorff topological space in which every point has an open neighborhood homeomorphic to some open subset of $\mathbb{R} \times \mathbb{R}_+$. Its *boundary* is the 1-dimensional manifold consisting of the points having a neighborhood homeomorphic to a neighborhood of $(0, 0)$ in $\mathbb{R} \times \mathbb{R}_+$. Note that, in particular, a surface without boundary is a surface with a boundary whose boundary is empty. In this work, we will only consider compact connected orientable surfaces with a boundary. By the classification theorem, they are characterized up to homeomorphisms by two nonnegative integers, the genus g and the number p of connected components of the




Fig. 1. The surface with a boundary $\Sigma_{1,3}^\partial$.

boundary. We denote by $\Sigma_{g,p}^\partial$ the unique (up to homeomorphism) compact orientable surface of genus g with p boundary components; it can be obtained from the compact orientable surface of genus g by removing p disjoint open disks whose boundaries are pairwise disjoint circles. See Figure 1.

A map is a cellular embedding of a finite graph (possibly with multiple edges and loops) into a compact connected orientable surface without boundary, considered up to orientation-preserving homeomorphisms. Cellular means that the faces of the map (the connected components of the complement of edges) are open 2-cells, that is, homeomorphic to 2-dimensional open disks. We will consider in this work *maps with a boundary*, which means that some distinguished faces of the map are considered as “holes.” Given such a map, we may remove from every distinguished face an open disk whose boundary is a circle and obtain a surface with a boundary. It would probably be more satisfactory to directly define these objects as embedded into a surface with a boundary in such a way that the distinguished faces are homeomorphic to disks with an open disk removed and all other faces are 2-cells but the previous definition happens to be more convenient to our purpose.

The natural problem of *scaling limits* of random maps has generated many studies in the last decade. The most natural setting is the following. We consider maps as metric spaces, endowed with their natural graph metric. We choose uniformly at random a map of “size” n in some class, rescale the metric by the proper factor, and look at the limit in the sense of the Gromov–Hausdorff topology [21]. The size considered is often the number of faces of the map. From this point of view, the most studied class is the class of planar quadrangulations. The pioneering work of Chassaing and Schaeffer [17] revealed that the proper scaling factor in this case is $n^{-1/4}$. The problem was first addressed by Marckert and Mokkadem [28], who constructed a candidate limiting space called the *Brownian map*, and showed the convergence toward it in another sense. Le Gall [23] then showed the relative compactness of this sequence of metric spaces and that any of its accumulation points was almost surely of Hausdorff dimension 4. It is only recently that the solution of the problem was completed independently by Miermont [31] and Le Gall [25], who showed that the scaling limit is indeed the Brownian map. This last step, however, is not mandatory in order to identify the topology of the limit: Le Gall and Paulin [27], and later Miermont [29], showed that any possible limit is homeomorphic to the 2-dimensional sphere.

This line of reasoning lead the way to several extensions. The first kind of extension is to consider other classes of planar maps. Actually, Le Gall already considered in [23] the classes of κ -angulations, for even $\kappa \geq 4$. In [25], he considered the classes of κ -angulations for $\kappa = 3$ and for even $\kappa \geq 4$ as well as the case of Boltzmann distributions on bipartite planar maps, conditioned on their number of vertices. Another extension is due to Addario-Berry and Albenque [2]; in this work, they consider simple triangulations and simple quadrangulations, that is, triangulations and quadrangulations without loops and multiple edges. Together with Jacob and Miermont [9], we later added the case of maps conditioned on their number of edges, and Abraham [1] considered the case of bipartite maps conditioned on their number of edges. In all these cases, the limiting space is always the same Brownian map (up to a multiplicative constant): we say that the Brownian map is *universal* and we expect it to arise as the scaling limit of a lot more of natural classes of maps. A peculiar extension is due to Le Gall and Miermont [26] who consider maps with large faces, forcing the limit to fall out of this universality class: they obtain so-called *stable maps*, which are related to stable processes.

Another kind of extension is to consider quadrangulations on a fixed surface that is no longer the sphere. The case of orientable surfaces of positive genus was the focus of [6,7] and the case of the disk was considered in [8]. In both

these cases, the last step is still missing at the moment: the convergence is only known to hold along subsequences. In the present work, we complete this picture by considering quadrangulations on arbitrary surfaces with a boundary.

The starting point of these problems is a powerful bijective encoding of the maps in the studied class by simpler objects. In the case of planar quadrangulations, the bijection in question is the so-called Cori–Vauquelin–Schaeffer bijection [17,18,35] and the simpler objects are trees whose vertices carry integer labels satisfying some conditions. In the other cases, variants of this bijection are used [3,13,16,33] and the encoding objects usually have a more intricate combinatorial structure.

At the present time, little is known about the limiting spaces. The principal result about the metric properties of the Brownian map is due to Le Gall [24]. In this reference, Le Gall studies all the geodesics from a uniformly chosen random point in the Brownian map. In particular, he shows that, toward a typical point, there is only one such geodesic and that for particular points there are up to three distinct such geodesics. Moreover, if we consider the unique geodesics toward two typical points, it is shown that they share a common part of positive length. We believe that Le Gall’s approach could be adapted to our framework of general Brownian surfaces. We also think that, in the particular case $p = 0$, the results could be derived from Le Gall’s, by using a continuous analog to Chapuy’s bijection [15] (see also [7], Section 6, for such an application). In the present work, however, we choose to use a third approach.

Let $\rho^\bullet, x_1, x_2, \dots$ be fixed uniform points in a Brownian surface. Our strategy relies on the following facts, which hold almost surely.

- (a) There is only one geodesic from ρ^\bullet to x_i (Proposition 19).
- (b) It is possible to “entrap” all the geodesics from ρ^\bullet to any given point with geodesics from ρ^\bullet to x_i ’s (Section 6).

Combining these facts, we are able to identify all the geodesics starting from ρ^\bullet . Fact (a) will be shown by a technique inspired from [30]. In this reference, a bijection between quadrangulations with q distinguished vertices and maps with q faces is provided. Very roughly speaking, a proper use of the parameters of this bijection in the particular case $q = 2$ gives a parametrization of the geodesics between the two distinguished points. Fact (b) is the core of our approach and will be explained in more details in the next section, as we need some more notation.

Throughout this work, we fix an integer $g \geq 0$ and work in fixed genus g . In order to ease the reading, we provide in page 645 a recapitulation of our most used notation. We will see during Section 4 that a Brownian surface may be constructed both as a quotient of $[0, 1]$ and as a quotient of a more complicated underlying structure. As often as possible, r, s, t will denote numbers in $[0, 1]$, a, b, c will denote points in the underlying structure and x, y, z will denote points in the Brownian surface.

2. Main results

2.1. Brownian surfaces

Let \mathfrak{m} be a map. We let $V(\mathfrak{m})$ denote its set of vertices, $E(\mathfrak{m})$ its set of edges, and $\vec{E}(\mathfrak{m})$ its set of half-edges (oriented edges). We say that a face f is *incident* to a half-edge e (or that e is incident to f) if e is included in the boundary of f and is oriented in such a way that f lies to its left. The number of half-edges incident to a face is called its *degree*. For any half-edge e , we denote by \bar{e} its reverse, as well as e^- and e^+ its origin and end. We denote by $d_{\mathfrak{m}}$ the graph metric on \mathfrak{m} defined as follows: for any $v, v' \in V(\mathfrak{m})$, the distance $d_{\mathfrak{m}}(v, v')$ is the number of edges of any shortest path linking v to v' .

For combinatorial reasons, we need to restrict ourselves to bipartite maps: a map is called *bipartite* if its vertex set can be partitioned into two subsets such that no edge links two vertices of the same subset. For $p \geq 0$, a *quadrangulation with p boundary components* is a bipartite map having p distinguished faces h_1, \dots, h_p and whose other faces are all of degree 4. The distinguished faces will be called *external faces* or *holes*. The other faces will be called *internal faces*. For technical reasons, the maps we consider will always implicitly be rooted, in the sense that they come with a distinguished half-edge, usually denoted by e_* . The *genus* of a map is defined as the genus of the surface into which it is embedded. For $n \in \mathbb{Z}_+$ and $\sigma = (\sigma^1, \dots, \sigma^p) \in \mathbb{N}^p$ (with the convention that $\mathbb{N}^0 := \{\emptyset\}$), we define the set $\mathcal{Q}_{n,\sigma}$ of all genus g quadrangulations with p boundary components having n internal faces and such that h_i is of degree $2\sigma^i$, for $1 \leq i \leq p$. See Figure 2 for an example and Figure 3 for a computer simulation of a large random quadrangulation.

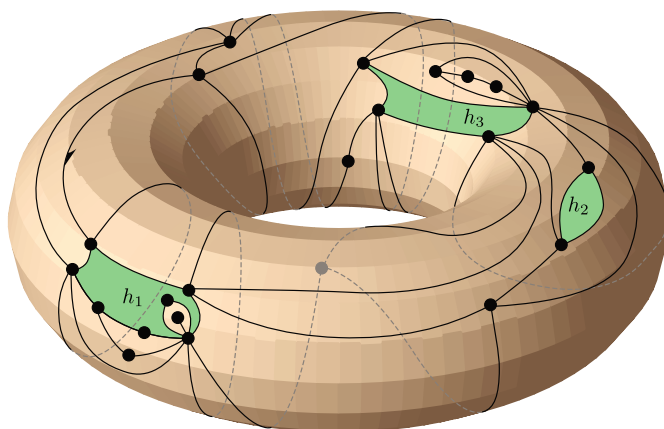


Fig. 2. A quadrangulation from $\mathcal{Q}_{19,(4,1,2)}$ in genus 1. The half arrowhead symbolizes the root.




Fig. 3. Simulation of a uniformly sampled quadrangulation from $\mathcal{Q}_{50000,(200,300)}$ in genus 2. The 2 boundary components are highlighted. Go to www.normalesup.org/~bettinel in order to zoom or rotate the map.

The Gromov–Hausdorff distance between two compact metric spaces (\mathcal{X}, δ) and (\mathcal{X}', δ') is defined by

$$d_{\text{GH}}((\mathcal{X}, \delta), (\mathcal{X}', \delta')) := \inf\{\delta_{\mathcal{H}}(\varphi(\mathcal{X}), \varphi'(\mathcal{X}'))\},$$

where the infimum is taken over all isometric embeddings $\varphi: \mathcal{X} \rightarrow \mathcal{X}''$ and $\varphi': \mathcal{X}' \rightarrow \mathcal{X}''$ of \mathcal{X} and \mathcal{X}' into the same metric space $(\mathcal{X}'', \delta'')$, and $\delta_{\mathcal{H}}$ stands for the usual Hausdorff distance between compact subsets of \mathcal{X}'' . This defines a metric on the set of isometry classes of compact metric spaces ([14], Theorem 7.3.30), making it a Polish space (this is a simple consequence of [14], Theorem 7.4.15).

The following result unifies and generalizes the main results from [6–8,23,27,29].

Theorem 1. *Let us consider an integer $p \geq 0$, positive real numbers $\sigma_{\infty}^1, \dots, \sigma_{\infty}^p > 0$ and a sequence of p -uples $\sigma_n = (\sigma_n^1, \dots, \sigma_n^p) \in \mathbb{N}^p$ such that $\sigma_n^i/\sqrt{2n} \rightarrow \sigma_{\infty}^i$, for $1 \leq i \leq p$. Let \mathfrak{q}_n be uniformly distributed over $\mathcal{Q}_{n, \sigma_n}$. Then, from any increasing sequence of integers, we may extract a subsequence $(n_k)_{k \geq 0}$ such that there exists a random metric space $(\mathfrak{q}_{\infty}^{\sigma}, d_{\infty}^{\sigma})$ satisfying*

$$\left(V(\mathfrak{q}_{n_k}), \left(\frac{9}{8n_k} \right)^{1/4} d_{\mathfrak{q}_{n_k}} \right) \xrightarrow[k \rightarrow \infty]{(d)} (\mathfrak{q}_{\infty}^{\sigma}, d_{\infty}^{\sigma})$$

in the sense of the Gromov–Hausdorff topology.

Moreover, regardless of the choice of the sequence of integers, the limiting space $(\mathfrak{q}_{\infty}^{\sigma}, d_{\infty}^{\sigma})$ is almost surely homeomorphic to $\Sigma_{g,p}^{\partial}$, has Hausdorff dimension 4, and every of the p connected components of its boundary has Hausdorff dimension 2.

Remark. *The constant $(8/9)^{1/4}$ is irrelevant in this statement. We chose to let it figure for the sake of consistency with previous works and because of following definitions. This constant is inherent to the case of quadrangulations: we believe that the same statement should hold with the same limiting spaces for other classes of maps embedded in the same surface and satisfying mild conditions, up to modifying this constant.*

In the spherical case $((g, p) = (0, 0))$, these results were shown by Le Gall [23] and the topology was identified by Le Gall and Paulin [27] (and also later by Miermont [29]). Using two different approaches, Le Gall [25] and Miermont [31] later substantially improved this result by showing that the extraction is not necessary in this case. It is strongly believed that it is not necessary in the general case either; this is the main focus of the works in progress [10,11]. Theorem 1 in the particular case $p = 0, g > 0$ can be found in [6,7] and the case $p = 1, g = 0$ is treated in [8].

Notice that, although it seems reasonable that the limiting space will have genus at most g and at most p holes, it is not clear a priori that it will be homeomorphic to $\Sigma_{g,p}^{\partial}$. We could imagine that some handles “disappear” or that some holes “merge” into a single hole. This does not happen; loosely speaking, this means that a uniform quadrangulation is sufficiently well spread over the surface, taking a macroscopic amount of space inbetween the holes and on every handle. Another noticeable fact is that the boundary of every hole is homeomorphic to a circle whereas, in the discrete picture, the holes do not in general have a simple curve as a boundary.

We believe that it is possible to extend this result to the case where some σ_{∞}^i 's are equal to 0. More precisely, making the extra hypothesis that the extraction in Theorem 1 is not necessary, we conjecture the following.

Conjecture 1. *Let $0 \leq p' \leq p$ be integers, let $\sigma_{\infty}^1, \dots, \sigma_{\infty}^{p'} > 0$ be positive real numbers and set $\sigma_{\infty}^{p'+1} = 0, \dots, \sigma_{\infty}^p = 0$. Let also $\sigma_n = (\sigma_n^1, \dots, \sigma_n^p) \in \mathbb{N}^p$ be p -uples such that $\sigma_n^i/\sqrt{2n} \rightarrow \sigma_{\infty}^i$, for $1 \leq i \leq p$ and consider a random variable \mathfrak{q}_n uniformly distributed over $\mathcal{Q}_{n, \sigma_n}$. Then, the following convergence holds in the sense of the Gromov–Hausdorff topology:*

$$\left(V(\mathfrak{q}_n), \left(\frac{9}{8n} \right)^{1/4} d_{\mathfrak{q}_n} \right) \xrightarrow[n \rightarrow \infty]{(d)} (\mathfrak{q}_{\infty}^{\sigma'}, d_{\infty}^{\sigma'}),$$

where $(\mathfrak{q}_{\infty}^{\sigma'}, d_{\infty}^{\sigma'})$ is the limiting space of Theorem 1 corresponding to $\sigma_{\infty}^1, \dots, \sigma_{\infty}^{p'}$.

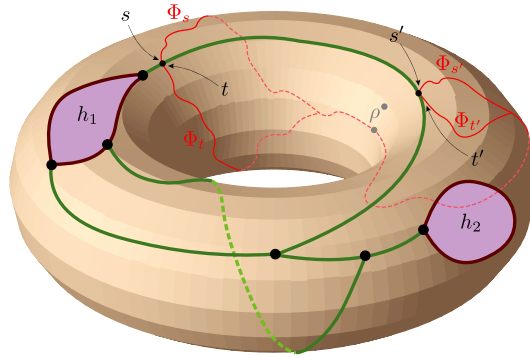


Fig. 4. Description of a Brownian torus with 2 holes. The backbone is the thick green and burgundy structure. It is shaped like a map with 2 holes and one face but we should think of the edges as having some “length.” Imagine that, on both sides of the edges, independent Brownian forests are embedded, except inside the 2 holes. In other words, Brownian CRT’s are grafted according to some Poissonian distribution. The union of these forests and the backbone constitutes the underlying structure \mathcal{M} .

We proved this conjecture in [8] in the case $g = 0$, $p' = 0$ and $p = 1$ (where it is known [25,31] that the extraction is not necessary) and we think that with some substantial work, our method can be adapted to this framework, provided the uniqueness of the limiting space. We decided, however, not to pursue this in the present paper as we prefer to concentrate on geodesics in the limiting spaces.

2.2. Geodesics in Brownian surfaces

We now turn to the geodesics in a fixed Brownian surface. We fix a subsequence $(n_k)_{k \geq 0}$ along which the convergence of Theorem 1 holds and we consider the corresponding Brownian surface $(q_\infty^\sigma, d_\infty^\sigma)$. Recall that, in a compact metric space (\mathcal{X}, δ) , a *geodesic* from a point $x \in \mathcal{X}$ to a point $y \in \mathcal{X}$ is a continuous path $\wp : [0, \delta(x, y)] \rightarrow \mathcal{X}$ such that $\wp(0) = x$, $\wp(\delta(x, y)) = y$ and $\delta(\wp(s), \wp(t)) = |t - s|$ for every $s, t \in [0, \delta(x, y)]$. The space (\mathcal{X}, δ) is called a *geodesic space* if any two points are connected by at least one geodesic. Because the metric space associated with a map is at d_{GH} -distance at most $1/2$ from a geodesic space and because the Gromov–Hausdorff limit of a sequence of geodesic spaces is also a geodesic space (see [14], Theorem 7.5.1), we see from Theorem 1 that any Brownian surface is a geodesic space. Although we do not know how to characterize all the geodesics in a Brownian surface, we are able to describe all the geodesics starting from a distinguished point. Before stating our results, we still need to introduce some more notions.

We will see later that the Brownian surface $(q_\infty^\sigma, d_\infty^\sigma)$ may be constructed as a quotient of an underlying structure \mathcal{M} . This structure consists of a *backbone*, which roughly captures its homotopy type, on which Brownian forests are glued. It generalizes Aldous’s CRT, which is the analog in the spherical case. The rigorous definition of this object will be given in Section 4.3.2, after some grounds will have been established. At this stage, we only give an informal description and refer to Figure 4 for a visual support. For the moment, do not pay attention to s, t, s', t' , and the corresponding paths Φ on the figure; these elements will be of interest in a short while.

It is also important to know at this point that \mathcal{M} comes with a natural distinguished point noted ρ^\bullet . This is the base point for the geodesics we will study. In general, there may exist simple loops in $(q_\infty^\sigma, d_\infty^\sigma)$ that are not homotopic to 0. This motivates the following definition. We let $\mathcal{N}(\rho^\bullet, q_\infty^\sigma)$ be the set of points $x \in q_\infty^\sigma$ for which there exists at least a pair $\{\wp, \wp'\}$ of geodesics from ρ^\bullet to x such that the concatenation of \wp with the reverse of \wp' is not homotopic to 0. The Hausdorff dimension of this set can be computed:

Theorem 2. *The set $\mathcal{N}(\rho^\bullet, q_\infty^\sigma)$ is a.s. empty if $g = 0$ and $p \in \{0, 1\}$; it a.s. has Hausdorff dimension 2 otherwise.*

We define the *order* of a point $a \in \mathcal{M}$ as the maximal number of connected components of $\mathcal{M} \setminus \{a\}$ restricted to the neighborhoods of a . Basic properties of Brownian forests show that this number is at most 3. The backbone $\mathcal{B} \subseteq \mathcal{M}$ is the set of points a either of order 2 and such that no connected components of $\mathcal{M} \setminus \{a\}$ are real trees or of order 3 and such that at most one connected component of $\mathcal{M} \setminus \{a\}$ is a real tree. For instance, note that $\mathcal{B} = \emptyset$ in

the case $(g, p) = (0, 0)$ of the sphere, as \mathcal{M} is a real tree in this case; in fact, this is the only case in which $\mathcal{B} = \emptyset$. Proposition 17 shows in particular that only points of order 1 can be identified in the quotient, so that we may define the order of a point in q_∞^σ as the order of any of its representatives in \mathcal{M} . Note that this definition a priori depends on \mathcal{M} . We denote by ∂q_∞^σ the boundary of $(q_\infty^\sigma, d_\infty^\sigma)$. The following theorem generalizes the main result of [24]:

Theorem 3. *The following properties hold almost surely.*

- (i) *For all $x \in q_\infty^\sigma$, the number of distinct geodesics from ρ^\bullet to x is equal to the order of x minus $\mathbb{1}_{\{x \in \partial q_\infty^\sigma\}}$. In particular, this number is typically 1, at most 2 for the boundary points and at most 3 for the interior points.*
- (ii) *The canonical projection of \mathcal{B} in q_∞^σ is the union of ∂q_∞^σ and the set $\mathcal{N}(\rho^\bullet, q_\infty^\sigma)$.*
- (iii) *The set of points reachable by 3 distinct geodesics and for which every pair $\{\wp, \wp'\}$ of geodesics is such that the concatenation of \wp with the reverse of \wp' is not homotopic to 0 is finite: its cardinality H is*
 - *equal to 0 if $g = 0$ and $p \in \{0, 1\}$;*
 - *equal to $4g - 2$ if $g \geq 1$ and $p = 0$;*
 - *a random variable whose distribution only depends on g , p and σ (it can even be computed) and whose support is $\{0, 1, \dots, 4g + p - 2\}$ otherwise.*
- (iv) *We suppose that $p \neq 0$ in order to avoid trivialities. The set $\partial q_\infty^\sigma \cap \mathcal{N}(\rho^\bullet, q_\infty^\sigma)$ is finite: its cardinality is equal to $4g + 2p - 2 - H$.*

Let us first informally describe in more details what happens on the example of Figure 4. The boundary ∂q_∞^σ corresponds to the burgundy edges and the set $\mathcal{N}(\rho^\bullet, q_\infty^\sigma)$ corresponds to the green edges. The projection in q_∞^σ of the vertices of the 3-face map corresponding to the backbone constitutes a finite set of cardinality $4g + 2p - 2$ (by a simple application of the Euler characteristic formula). Among these points, H are not incident to any hole and $4g + 2p - 2 - H$ are incident to some hole ($H = 4$ on the picture). As the 3-face map of the picture is actually a random variable, we see that H is also random and the exact distribution can be computed once we know the distribution of this 3-face map. Moreover, these $4g + 2p - 2$ points are precisely the ones described in (iii) or (iv), depending on whether they belong to ∂q_∞^σ or not.

We see that, in particular, our notion of order for a point in q_∞^σ only depends on ρ^\bullet . It is possible to define a uniform measure on q_∞^σ , as the projection of the Lebesgue measure on $[0, 1]$ (it can also be seen as the limit in some sense of the uniform measure on $V(q_n)$). Typical in (i) means that, if a point is distributed according to this measure, the number for this point is a.s. 1. We can also show that ρ^\bullet is uniformly distributed over q_∞^σ and the “invariance under uniform rerooting” property (see [24], Section 8, for precise statements in the spherical case, which can easily be extended to general Brownian surfaces) yields a similar result if we replace ρ^\bullet with another uniformly chosen random point. Another important fact to notice is that the distribution of H does not depend on the subsequence $(n_k)_{k \geq 0}$. This reinforces the conjecture that the extraction in Theorem 1 is actually not necessary.

Using basic properties of \mathcal{M} , we can classify the points of q_∞^σ as follows. The set of points reachable from ρ^\bullet by a unique geodesic a.s. has Hausdorff dimension 4. The set of points reachable from ρ^\bullet by at least two distinct geodesics is a dense subset of q_∞^σ , which is relatively small (we can show that its Hausdorff dimension is a.s. 2). An important part of these points moreover belong to $\mathcal{N}(\rho^\bullet, q_\infty^\sigma)$, which is also a set of dimension a.s. 2. Finally, the set of points that can be reached by three geodesics is countable and only a finite number of these points are such that the concatenation of any two geodesics is not homotopic to 0.

Note also that, as a random map on a surface is locally planar (at least far from the boundary), it is to be expected that, in the vicinity of ρ^\bullet , every Brownian surface should look alike. In this regard, it is probable that (i) in the latter vicinity can also be directly derived from Le Gall’s results [24] by an absolute continuity argument.

Let us comment a little more about the idea behind the proof of Theorem 3. It comes from a complete description of the geodesics from ρ^\bullet : we show that, in the language of [24], all the geodesics from ρ^\bullet are *simple* (see Proposition 23). We can picture \mathcal{M} as a continuous analog to a discrete map with 1 internal face and p holes. As for a real tree (which is the analog to a discrete tree, that is, a map with 1 face in genus 0), we will see that it is possible to define a contour order around the “internal face” of \mathcal{M} . The order of a point $a \in \mathcal{M}$ is then the number of times a is visited in this contour order plus 1 if a is incident to a hole (because the holes are not visited). To each time a is visited in the contour order correspond a simple geodesic. Furthermore, Lemma 25 gives a criterion to decide whether the concatenation

of two simple geodesics aiming at the same point are homotopic to 0 or not, in terms of the visiting times of the corresponding point in \mathcal{M} .

Going back to Figure 4, we denote by Φ_s the simple geodesic corresponding to the time s . While performing the contour order, at time s , we are visiting some forest (even possibly its floor) grafted on a half-edge of the 3-face map: we denote this half-edge by \tilde{e}_s . We consider a point in \mathcal{M} that is visited both at times s and t . Then the criterion of Lemma 25 states that the concatenation of Φ_s with the reverse of Φ_t is homotopic to 0 if and only if $\tilde{e}_s = \tilde{e}_t$. For instance, on the picture of Figure 4, $\tilde{e}_s \neq \tilde{e}_t$ so that Φ_s and Φ_t constitute a loop that is not homotopic to 0. On the contrary, $\Phi_{s'}$ and $\Phi_{t'}$ constitute a loop that is homotopic to 0. This explains the correspondence between the sets considered in Theorem 3 and the elements of the 3-face map. Finally, the $-\mathbb{1}_{\{x \in \partial q_\infty^\sigma\}}$ in (i) simply comes from the fact that there are no geodesics inside the holes.

We can now give the intuition of Fact (b) from the introduction. For simplicity, let us suppose that a is a point of order 1 and let s be the time at which it is visited in the contour order. By density, we can find times s_i and s_j as close as wanted to s such that $s_i < s < s_j$ and such that the corresponding points x_i and x_j in q_∞^σ are from our fixed sequence of uniform points. By Fact (a), the geodesics from x_i and x_j to ρ^\bullet are necessarily simple geodesics. Using the description of simple geodesics, we can show that these geodesics have to merge within the vicinity of the point x corresponding to a in q_∞^σ . As a result, every geodesic from x to ρ^\bullet is “trapped” between these two geodesics and can only escape through the segment of \mathcal{M} connecting the points corresponding to s_i and s_j . Using fine properties of the labels on \mathcal{M} , we show that it is possible to select s_i and s_j in such a way that the geodesic does not escape through this segment. Hence it has to coincide with the other two geodesics, at least after they merge.

To complete this presentation, we also recover the same “confluence property” as the one obtained by Le Gall in the spherical case ([24], Corollary 7.7).

Proposition 4 (Corollary 24). *Almost surely, for every $\varepsilon > 0$, there exists $\eta \in (0, \varepsilon)$ such that all the geodesics from ρ^\bullet to points outside of the ball of radius ε centered at ρ^\bullet share a common initial part of length η .*

To end this section, let us mention that, as observed by Le Gall in the spherical case, our results share surprising similarities with the results of Myers [32] in the context of differential geometry, where the surfaces considered possess by far more regularities. In the latter reference, it is shown that the locus of points reachable from a fixed point by at least two geodesics in a surface without boundary is a linear graph. Moreover, the subtraction of this locus from the surface yields a unique open 2-cell and the number of distinct geodesics to a point of the locus is equal to the order of this point in the locus (where the definition of order matches our definition of order in \mathcal{M}).

2.3. Geodesics in large quadrangulations

Theorem 3 allows us to derive asymptotic results for large quadrangulations. We adapt the presentation and technique of Le Gall [24] to obtain them. Note that, in this section, we do not restrict our attention to subsequences along which the convergence of Theorem 1 holds.

We define the distance between a path \wp going successively through the vertices v_1, \dots, v_k and a path \wp' going successively through the vertices $v'_1, \dots, v'_{k'}$ in the same map \mathfrak{m} by

$$d_{\text{path}}(\wp, \wp') = \max_{i \geq 0} \{d_{\mathfrak{m}}(v_{i \wedge k}, v'_{i \wedge k'})\}.$$

For $\varepsilon > 0$ and $v, v' \in V(q_n)$, we let $\text{Geod}_n^\varepsilon(v, v')$ be the set of paths from v to v' of length at most $(1 + \varepsilon)d_{q_n}(v, v')$. For $\delta > 0$, we also let $\text{Mult}_{n,\delta}^\varepsilon(v, v')$ be the maximal integer r such that there exist r paths $\wp_1, \dots, \wp_r \in \text{Geod}_n^\varepsilon(v, v')$ satisfying $d_{\text{path}}(\wp_i, \wp_j) \geq \delta n^{1/4}$ whenever $i \neq j$.

Let $(\varepsilon_n)_n$ be a fixed deterministic sequence of nonnegative real numbers such that $\varepsilon_n \rightarrow 0$ as $n \rightarrow \infty$ and, conditionally given q_n , let v_n^\bullet be a uniformly distributed vertex in $V(q_n)$. We denote by $\partial V(q_n)$ the set of vertices in $V(q_n)$ that are incident to a hole. We start with the translation of (i) in terms of large discrete maps.

Proposition 5. *Let v_n denote a uniformly distributed vertex in $V(q_n)$, independent from v_n^\bullet . Then, for every $\delta > 0$, $\mathbb{P}(\text{Mult}_{n,\delta}^{\varepsilon_n}(v_n^\bullet, v_n) = 1) \rightarrow 1$ as $n \rightarrow \infty$.*

Moreover, for every $\delta > 0$ and for every sequence $(\eta_n)_n$ of positive numbers such that $\eta_n \rightarrow 0$ as $n \rightarrow \infty$, $\mathbb{P}(\exists v \in V(q_n): d_{q_n}(v, \partial V(q_n)) \leq n^{1/4}\eta_n, \text{Mult}_{n,\delta}^{\varepsilon_n}(v_n^\bullet, v) \geq 3) \rightarrow 0$ as $n \rightarrow \infty$.

Finally, for every $\delta > 0$, $\mathbb{P}(\exists v \in V(q_n): \text{Mult}_{n,\delta}^{\varepsilon_n}(v_n^\bullet, v) \geq 4) \rightarrow 0$ as $n \rightarrow \infty$ and

$$\lim_{\delta \rightarrow 0} \liminf_{n \rightarrow \infty} \mathbb{P}(\exists v \in V(q_n): \text{Mult}_{n,\delta}^0(v_n^\bullet, v) = 3) = 1.$$

Note that we chose to state the result with geodesics in the displayed equation, instead of “approximate” geodesics because it is actually stronger. Of course, the weaker result with ε_n in place of 0 also holds. The first part of Proposition 5 roughly means that, in a large uniform quadrangulation, there is essentially a unique macroscopic geodesic from v_n^\bullet to a uniformly chosen vertex and this remains true even if we allow some “slack” to the geodesics. The second part states that the points asymptotically close to the boundary cannot be reached from v_n^\bullet by more than two distinct macroscopic geodesics, even if “slack” is allowed. The last part says that asymptotically, there exist points that can be reached from v_n^\bullet by three distinct macroscopic geodesics but no points can be reached by four distinct macroscopic geodesics, even if “slack” is allowed.

We now complete the interpretation. For $\varepsilon > 0$ and $v, v' \in V(q_n)$, we let $\text{HMult}_n^\varepsilon(v, v')$ be the maximal integer r such that there exist r paths $\wp_1, \dots, \wp_r \in \text{Geod}_n^\varepsilon(v, v')$ such that, whenever $i \neq j$, the concatenation of \wp_i with the reverse of \wp_j is not homotopic to 0. Here, homotopic to 0 means homotopic to 0 in the surface with a boundary corresponding to the map, that is, the surface in which the map is embedded, with open disks removed from the external faces. Note that, as the topology is preserved at the limit (handles and holes do not vanish), we can find a (random) δ small enough such that, for large n , $\text{HMult}_n^\varepsilon(v, v') \leq \text{Mult}_{n,\delta}^\varepsilon(v, v')$. With a little bit of work, this yields the following statement:

Proposition 6. We have $\mathbb{P}(\exists v \in V(q_n): \text{HMult}_n^{\varepsilon_n}(v_n^\bullet, v) \geq 4) \rightarrow 0$ as $n \rightarrow \infty$.

In order to give a discrete version of (ii), we need to know that the encoding bijection gives rise to a particular subset $\mathcal{B}(q_n, v_n^\bullet) \subseteq V(q_n)$, which is the discrete counterpart of \mathcal{B} . For the reader already familiar with this kind of bijections, it is the 2-core of the encoding object, that is, the set of vertices remaining after iteratively removing from it all the vertices of degree 1. See the end of Section 3.1 for a precise definition.

It is not hard to see that all the points $v \in \mathcal{B}(q_n, v_n^\bullet) \setminus \partial V(q_n)$ satisfy $\text{HMult}_n^0(v_n^\bullet, v) \geq 2$. The following proposition roughly states that all such points should be close to this set.

Proposition 7. For every $\delta > 0$,

$$\mathbb{P}(\exists v \in V(q_n): d_{q_n}(v, \mathcal{B}(q_n, v_n^\bullet) \setminus \partial V(q_n)) \geq \delta n^{1/4}, \text{HMult}_n^{\varepsilon_n}(v_n^\bullet, v) \geq 2) \xrightarrow[n \rightarrow \infty]{} 0.$$

The following statement means that the vertices $v \in V(q_n)$ such that $\text{HMult}_n^{\varepsilon_n}(v_n^\bullet, v) = 3$ are essentially dispatched in small separate “zones” and that the number of these zones is asymptotically the random variable H .

Proposition 8. Let $A_{n,\delta}^{\varepsilon_n}(j)$ denote the event that there exist j vertices $v_1, \dots, v_j \in V(q_n)$ such that $\text{HMult}_n^{\varepsilon_n}(v_n^\bullet, v_i) = 3$ for $1 \leq i \leq j$ and $d_{q_n}(v_i, v_{i'}) \geq \delta n^{1/4}$ whenever $i \neq i'$. Then

$$\lim_{\delta \rightarrow 0} \liminf_{n \rightarrow \infty} \mathbb{P}(A_{n,\delta}^{\varepsilon_n}(j) \setminus A_{n,\delta}^{\varepsilon_n}(j+1)) = \lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} \mathbb{P}(A_{n,\delta}^{\varepsilon_n}(j) \setminus A_{n,\delta}^{\varepsilon_n}(j+1)) = \mathbb{P}(H = j),$$

where H is the random variable defined in (iii).

Finally, we have a similar result for the boundary vertices v for which $\text{HMult}_n^{\varepsilon_n}(v_n^\bullet, v) \geq 2$.

Proposition 9. We suppose that $p \neq 0$ and we denote by $B_{n,\delta}^{\varepsilon_n}(j)$ the event that there exist j vertices $v_1, \dots, v_j \in \partial V(q_n)$ such that $\text{HMult}_n^{\varepsilon_n}(v_n^\bullet, v_i) \geq 2$ for $1 \leq i \leq j$ and $d_{q_n}(v_i, v_{i'}) \geq \delta n^{1/4}$ whenever $i \neq i'$. Then

$$\lim_{\delta \rightarrow 0} \liminf_{n \rightarrow \infty} \mathbb{P}(B_{n,\delta}^{\varepsilon_n}(j) \setminus B_{n,\delta}^{\varepsilon_n}(j+1)) = \mathbb{P}(H = 4g + 2p - 2 - j)$$

and the same statement holds if we replace the \liminf with a \limsup .

2.4. Organization of the paper

In Section 3, we explain how to encode quadrangulations by simpler maps. We begin with the mapping for pointed quadrangulations, which constitutes the starting point of our study and we then present the two-point mapping, which will be crucially used in Section 5.1. Section 4 is devoted to scaling limits and the proof of Theorem 1. As it uses the general framework of [6–8], we only sketch it and point to the particularity of the cases considered here. We then study typical geodesics and show Fact (a) in Section 5 and finally, we study the general geodesics in Section 6 by the entrapment technique (Fact (b)). This leads to Proposition 23, which identifies all the geodesics from ρ^\bullet . Section 7 is devoted to the remaining proofs.

3. Encoding quadrangulations

Let us now present the bijections allowing to encode quadrangulations carrying one or two distinguished vertices. Our description is a slight reformulation of the Bouttier–Di Francesco–Guitter bijection [13] and the Miermont bijection [30]. We refer the reader to these references for proofs. Throughout this section, n and p denote nonnegative integers and $\sigma = (\sigma^1, \dots, \sigma^p)$ a p -uple of positive integers.

3.1. Quadrangulations carrying one distinguished vertex

Let $q \in \mathcal{Q}_{n,\sigma}$ be a quadrangulation and $v^\bullet \in V(q)$ one of its vertices. We assign labels to the vertices of q as follows: for every vertex $v \in V(q)$, we set $l(v) := d_q(v^\bullet, v)$. Because q is by definition bipartite, the labels of both ends of any edge differ by exactly 1. As a result, the internal faces can be of two types: the labels around the face are either $d, d + 1, d + 2, d + 1$, or $d, d + 1, d, d + 1$ for some d . We add a new edge inside every internal face following the convention depicted on the left part of Figure 5.

A *corner* is an angular sector delimited by two successive half-edges in the contour of a face. The vertex located at the end of the first half-edge is called the vertex *incident* to the corner. If c is a corner incident to a vertex v , we write $l(c) := l(v)$ with a slight abuse of notation. For each i , we let $c_i^0, c_i^1, \dots, c_i^{2\sigma^i - 1}$ be the corners of h_i read in clockwise order, starting at an arbitrary corner (and we use the convention $c_i^{2\sigma^i} := c_i^0$). We link together in a cycle the corners c_i^k 's such that $l(c_i^{k+1}) = l(c_i^k) - 1$, as shown on the right part of Figure 5. Note that, because $l(c_i^{k+1}) - l(c_i^k) \in \{-1, +1\}$, there are exactly σ^i such corners.

We then only keep the new edges we added and the vertices in $V(q) \setminus \{v^\bullet\}$. The object we obtain is a labeled map m of genus g with $p + 1$ faces. There is an obvious correspondence between the external faces of q and p of the faces of m . Let h_1, \dots, h_p also denote these faces in m . Note that, by construction, these faces all have a simple boundary and are of degrees $\sigma^1, \dots, \sigma^p$. Remark also that v^\bullet lies within the remaining face of m , which we denote by f^\bullet . See Figure 6.

We root the map m as follows. Let e be the only half-edge among the root of q and its reverse such that $l(e^+) = l(e^-) + 1$, and let f be the face of q that is incident to e . If f is an internal face, the root of m is the half-edge corresponding to the edge we added in f , directed from e^+ . If f is an external face, there are two new half-edges inside f starting from e^+ ; the root of m is the one incident to f^\bullet . See Figure 7.




Fig. 5. *Left.* Adding a new edge inside an internal face. *Right.* Adding σ^i new edges inside the external face h_i . On this example, $\sigma^i = 6$.




Fig. 6. The construction for a map in $\mathcal{Q}_{12,(2,4,1)}$ in the case $g = 0$.




Fig. 7. Rooting the map m . On the picture, the two possible roots for q yielding the same root for m are shown.

For each i , let $\vec{h}_i^1, \vec{h}_i^2, \dots, \vec{h}_i^{\sigma^i}$ be the half-edges incident to h_i in m , read in counterclockwise order around it. The labels of m satisfy the following:

- for $1 \leq i \leq p$ and $1 \leq j \leq \sigma^i$, we have $l(\vec{h}_i^{j-}) - l(\vec{h}_i^{j+}) \geq -1$;
- for any half-edge $e \in \vec{E}(m)$ such that neither e nor its reverse \bar{e} is incident to an h_i , we have $|l(e^+) - l(e^-)| \leq 1$.

We will consider the labels of m up to an additive constant: we write

$$[\mathbb{I}] := \{v \in V(m) \mapsto l(v) + a : a \in \mathbb{Z}\}$$

the class of l for this equivalence relation. We say that two faces are *adjacent* if there exist a half-edge incident to one and whose reverse is incident to the other. Let $\mathcal{M}_{n,\sigma}^1$ denote the set of genus g maps having $n + |\sigma|$ edges (where we write $|\sigma| := \sum_{i=1}^p \sigma^i$) and $p + 1$ faces denoted by $h_1, \dots, h_p, f^\bullet$ such that, for all i , h_i has a simple boundary, is of degree σ^i and is not adjacent to any other h_j , and such that the root is not incident to any hole h_j . Note that any edge is forbidden to be incident to two different holes, but there may exist vertices that are incident to two or more holes. We also denote by $\mathfrak{M}_{n,\sigma}^1$ the set of pairs whose first coordinate lies in $\mathcal{M}_{n,\sigma}^1$ and whose second coordinate is an equivalence class of labeling functions on the vertices of the map satisfying the two previous itemized conditions.

Proposition 10. *The mapping $(q, v^\bullet) \mapsto (m, [\mathbb{I}])$ is a two-to-one mapping from the set of quadrangulations in $\mathcal{Q}_{n,\sigma}$ carrying one distinguished vertex to the set $\mathfrak{M}_{n,\sigma}^1$.*

At this point, we may properly define the set $\mathcal{B}(q, v^\bullet)$ appearing in Section 2.3. Let $(m, [\mathbb{I}])$ denote the map corresponding to (q, v^\bullet) by the previous mapping. Then $\mathcal{B}(q, v^\bullet)$ is the 2-core of m , that is, the set of vertices of m remaining after iteratively removing from m all its vertices of degree 1.

3.2. Quadrangulations carrying two distinguished vertices

In the case of doubly-pointed quadrangulations, the construction is very similar. We consider a quadrangulation $q \in \mathcal{Q}_{n,\sigma}$ and $v^\bullet, v^{\bullet\bullet} \in V(q)$ two of its vertices such that $d := d_q(v^\bullet, v^{\bullet\bullet}) \geq 2$. We also need an integer $\lambda \in \{1, 2, \dots, d - 1\}$. We assign labels to the vertices of q in a similar fashion as in the last section but we now take into account the extra vertex $v^{\bullet\bullet}$: we define

$$l(v) := \min(d_q(v^\bullet, v), d_q(v^{\bullet\bullet}, v) + 2\lambda - d), \quad v \in V(q). \tag{1}$$




Fig. 8. On this example, $d = 3$ and $\lambda = 1$. It can be checked for example that the vertices corresponding in q to the three vertices labeled 1 on the interface between f^\bullet and $f^{\bullet\bullet}$ are at distance 1 from v^\bullet and at distance 2 from $v^{\bullet\bullet}$.

A simple way to picture this function is to imagine water flowing through the edges of q at rate 1, starting from two sources: one located at the vertex v^\bullet opened at time 0 and one located at the vertex $v^{\bullet\bullet}$ opened at time $2\lambda - d$.

We then follow rigorously the same procedure as in last section, and we obtain a class of labeled maps $(m, [\Gamma])$ lying in $\mathfrak{M}_{n,\sigma}^2$, which is defined exactly as $\mathfrak{M}_{n,\sigma}^1$, with the exception that the map now has $p + 2$ faces $h_1, \dots, h_p, f^\bullet, f^{\bullet\bullet}$ instead of $p + 1$. It turns out that v^\bullet and $v^{\bullet\bullet}$ do not belong to the same face of m and we let f^\bullet be the face containing v^\bullet and $f^{\bullet\bullet}$ be the one containing $v^{\bullet\bullet}$. We also define the analog $\mathcal{M}_{n,\sigma}^2$ of $\mathcal{M}_{n,\sigma}^1$ with $p + 2$ faces. Let us emphasize at this point that, because the holes have a simple boundary and are not adjacent with one another, every edge of m is incident to at least one of the faces f^\bullet or $f^{\bullet\bullet}$. See Figure 8.

Proposition 11. *This construction provides a two-to-one mapping from*

$$\{(q, (v^\bullet, v^{\bullet\bullet}), \lambda) : q \in \mathcal{Q}_{n,\sigma}, (v^\bullet, v^{\bullet\bullet}) \in V(q)^2 | d_q(v^\bullet, v^{\bullet\bullet}) \geq 2, 1 \leq \lambda \leq d_q(v^\bullet, v^{\bullet\bullet}) - 1\}$$

to $\mathfrak{M}_{n,\sigma}^2$.

Remark. In [30], Miermont considers what he calls delay vectors. In the particular case of 2-pointed quadrangulations, it is more natural to use the integer λ instead. The construction we give here corresponds to the delay vector $[0, 2\lambda - d]$ with his notation. It can moreover be observed that, if u is any of the vertices where the water from the two sources meet for the first time, then $\lambda = d_q(v^\bullet, u)$ and $d - \lambda = d_q(v^{\bullet\bullet}, u)$. In some sense, the integer λ parametrizes the geodesics between v^\bullet and $v^{\bullet\bullet}$. This observation will be the key point of the proof of Proposition 19.

3.3. Reverse mappings

In both cases, the reverse mapping uses the same procedure. We describe it briefly in this section. Let $(m, [\Gamma])$ be in either $\mathfrak{M}_{n,\sigma}^1$ or $\mathfrak{M}_{n,\sigma}^2$. First, we add inside f^\bullet a new vertex v^\bullet with label

$$l(v^\bullet) := \min_{u \in f^\bullet} l(u) - 1,$$

where the notation $u \in f$ means that u belongs to the boundary of f . Following the counterclockwise order around f^\bullet , we draw arcs linking every corner to the first subsequent corner that has a strictly smaller label. If no such corners exist, which means that the corner we are visiting has a label that is minimal on f^\bullet , we draw the arc from the corner to the extra vertex v^\bullet . It is possible to draw these arcs in such a way that they do not cross each other or the edges of m . In the case where $(m, [\Gamma]) \in \mathfrak{M}_{n,\sigma}^2$, we follow the same procedure for $f^{\bullet\bullet}$: we add a vertex $v^{\bullet\bullet}$ inside it with label $l(v^{\bullet\bullet}) := \min_{u \in f^{\bullet\bullet}} l(u) - 1$ and we draw arcs in the same way.

When removing the edges of m , we are left with a quadrangulation q . Each hole h_i naturally corresponds to an external face of q , which we also denote by h_i . The root of q is defined as the arc drawn from the corner preceding the root of m , oriented in one direction or the other.

If $(m, [\Gamma]) \in \mathfrak{M}_{n,\sigma}^1$, its two pre-images are the pairs (q, v^\bullet) , where q is rooted in the two possible ways. If $(m, [\Gamma]) \in \mathfrak{M}_{n,\sigma}^2$, its two pre-images are

$$(q, (v^\bullet, v^{\bullet\bullet}), \lambda) \quad \text{where } \lambda := \frac{1}{2}(l(v^{\bullet\bullet}) - l(v^\bullet) + d_q(v^\bullet, v^{\bullet\bullet})),$$

with the two possible rootings of q .

Lemma 12. *In the case of the two-point mapping, the following holds:*

- (i) if $v \in f^\bullet$, then $l(v) - l(v^\bullet) = d_q(v^\bullet, v)$ and $l(v) - l(v^{\bullet\bullet}) \leq d_q(v^{\bullet\bullet}, v)$;
- (ii) if $v \in f^{\bullet\bullet}$, then $l(v) - l(v^\bullet) \leq d_q(v^\bullet, v)$ and $l(v) - l(v^{\bullet\bullet}) = d_q(v^{\bullet\bullet}, v)$;
- (iii) if $v \in f^{\bullet\bullet}$, then $d_q(v^\bullet, v) \geq \lambda$, if $v \in f^\bullet$, then $d_q(v^{\bullet\bullet}, v) \geq d_q(v^{\bullet\bullet}, v^\bullet) - \lambda$, and

$$\lambda = \min_{v \in f^\bullet \cap f^{\bullet\bullet}} d_q(v^\bullet, v) = \min_{v \in f^\bullet \cap f^{\bullet\bullet}} l(v) - \min_{v \in f^\bullet} l(v) + 1.$$

Proof. (i) and (ii). By (1), for any vertex v , we have $l(v) - l(v^\bullet) \leq d_q(v^\bullet, v)$ and $l(v) - l(v^{\bullet\bullet}) \leq d_q(v^{\bullet\bullet}, v)$ (note that the representative of $[\Gamma]$ appearing in (1) is the one vanishing at v^\bullet). Moreover, if $v \in f^\bullet$, the successive arcs drawn from v to v^\bullet form a path in q of length $l(v) - l(v^\bullet)$, so that $d_q(v^\bullet, v) \leq l(v) - l(v^\bullet)$. This gives (i) and (ii) is proven in a similar fashion.

(iii). Let us fix $v \in f^{\bullet\bullet}$. By the triangle inequality and (ii),

$$d_q(v^\bullet, v^{\bullet\bullet}) \leq d_q(v^\bullet, v) + l(v) - l(v^{\bullet\bullet}) \leq 2d_q(v^\bullet, v) + l(v^\bullet) - l(v^{\bullet\bullet}).$$

The first part of the statement follows and the second part is obtained by the same argument. Finally, if $v \in f^\bullet \cap f^{\bullet\bullet}$ is on a geodesic linking v^\bullet to $v^{\bullet\bullet}$, all the previous inequalities are equalities and $d_q(v^\bullet, v) = \lambda$. The first equality follows and the second is obtained thanks to (i). Note also that the condition of nonadjacency of the holes makes it clear that $f^\bullet \cap f^{\bullet\bullet} \neq \emptyset$. \square

3.4. Further decomposition

A labeled map $(m, [\Gamma])$ from $\mathfrak{M}_{n,\sigma}^1$ or $\mathfrak{M}_{n,\sigma}^2$ can be decomposed into simpler objects: a *scheme*, which in some sense accounts for the homotopy type of m , and a collection of forests indexed by some half-edges of the scheme. The labeling function naturally gives rise to labels on the vertices of these forests as well as to bridges recording the labels on the cycles of the map m .

Remark 1. *The case $(m, [\Gamma]) \in \mathfrak{M}_{n,\sigma}^1$ in genus 0 is somehow degenerate. Indeed, in this case, m is merely a plane tree and cannot be further decomposed in our sense. Until further notice, we suppose that we are not in this case.*

Remark 2. *Up to a slight difference caused by the root, a scheme is sometimes called kernel in graph theory. We chose to stick with the terminology of scheme, which seems more common in the context of maps.*

3.4.1. Decomposition of the map

Let us first focus on the map m and keep the labels for later. We refer to Figure 11 below for visual support. We iteratively remove from m all its vertices of degree 1 that are not extremities of its root e_* . The set of vertices remaining at this point is called the *floor* of m . Among these vertices, some are called *nodes*: all vertices of degree 3 or more are nodes and,

- if e_*^- is of degree 1, then e_*^- is a node;
- if e_*^+ is of degree 1, then e_*^+ is a node;
- if neither e_*^- nor e_*^+ has degree 1, then e_*^- is a node.

On this map, the vertices that are not nodes are of degree 2 and are arranged into chains joining nodes. We define the map \mathfrak{s} by replacing each of these chains by a single edge. The root of \mathfrak{s} is defined as the edge replacing the chain that contains e_* , oriented in the same direction as e_* . The map \mathfrak{s} is a scheme in the following sense.

Definition 1. A scheme of type $(p, 1)$ is a genus g map with $p + 1$ faces denoted by $h_1, \dots, h_p, f^\bullet$, whose root is not incident to any h_j , and that satisfies the following conditions. For all i , h_i has a simple boundary and is not adjacent to any h_j . There may only be one vertex with degree 1 or 2: if it has degree 1, then it is an extremity of the root; if it has degree 2, then it is the origin of the root.

A scheme of type $(p, 2)$ is a genus g map with $p + 2$ faces denoted by $h_1, \dots, h_p, f^\bullet, f^{\bullet\bullet}$ and satisfying the same conditions.

Remark. A more conventional definition would be to forbid vertices of degree less than 2. Our choice of “keeping the root present” in the scheme is done for combinatorial reasons. It is only in Section 7 that we will need to get rid of this “root,” which does not play any geometrical part. Compare the definition of the backbone $\mathcal{B}(q, v^\bullet)$ used in the geometrical statements of Section 2.3 with the definition of the floor used for combinatorial purposes.

Let $\mathfrak{S}_{p,1}$ and $\mathfrak{S}_{p,2}$ be the finite sets of schemes of type $(p, 1)$ and $(p, 2)$. For example, in genus $g = 0$, the set $\mathfrak{S}_{1,1}$ contains the three maps represented on Figure 9.

We will use the following formalism for forests.

Definition 2. A forest of length $\xi \geq 1$ and mass $m \geq 0$ is an ordered family of ξ trees (which can be defined as planar one-faced maps) with total number of edges equal to m . Let \mathcal{F}_ξ^m denote the set of these forests.

It will be convenient to systematically add a $\xi + 1$ th tree to a forest of \mathcal{F}_ξ^m consisting of one single vertex. The floor of a forest is the set of the root vertices of its trees, with the extra vertex-tree included. In the drawings, we will also add extra edges linking the elements of the floor (as on Figure 11).

The half-edges of the scheme \mathfrak{s} are of several different types. Let $\vec{H}_i(\mathfrak{s})$ be the set of half-edges incident to the hole h_i and let $\vec{F}(\mathfrak{s})$ be the set of half-edges that are not incident to any holes. We break the set $\vec{F}(\mathfrak{s})$ into two subsets: let $\vec{I}(\mathfrak{s}) := \{e \in \vec{F}(\mathfrak{s}) : \bar{e} \in \vec{F}(\mathfrak{s})\}$ and $\vec{B}(\mathfrak{s}) := \{e \in \vec{F}(\mathfrak{s}) : \bar{e} \notin \vec{F}(\mathfrak{s})\}$. See Figure 10. The letter I stands for “internal half-edges” and B stands for “boundary half-edges.”

Every half-edge $e \in \vec{F}(\mathfrak{s})$ naturally corresponds to a forest \mathfrak{f}^e defined as follows. Let $f \in \{f^\bullet, f^{\bullet\bullet}\}$ be the face incident to e and let $e' \in \vec{F}(\mathfrak{s})$ be the half-edge preceding e in the contour order of f . By definition, the half-edges e and e' correspond to chains of half-edges in \mathfrak{m} : let \hat{e} and \hat{e}' be the last half-edges of these chains. The forest \mathfrak{f}^e




Fig. 9. The three elements of $\mathfrak{S}_{1,1}$ in genus $g = 0$.




Fig. 10. Notation for the different types of half-edges of a scheme. The set $\vec{F}(\mathfrak{s})$ is the union of $\vec{B}(\mathfrak{s})$ and $\vec{I}(\mathfrak{s})$.




Fig. 11. Decomposition of a labeled map from $\mathfrak{M}_{33,(3,4)}^2$ in genus 0 into a scheme $s \in \mathfrak{S}_{2,2}$, a collection of labeled forests $(f^e, l^e)_{e \in \bar{F}(s)}$ and a collection of bridges $(b^e)_{e \in \bar{F}(s)}$. The bridges have different colors, depending on whether they are interior bridges or boundary bridges. On the left, the floor and the 6 nodes of m are represented in red and with thicker outlines.

corresponds to the set of half-edges in m visited between \hat{e}' and \hat{e} (\hat{e}' excluded, \hat{e} included) in the contour order of f . See Figure 11.

For $e \in \bar{F}(s)$, we denote by $\xi^e \geq 1$ and $m^e \geq 0$ the length and mass of the forest f^e : $f^e \in \mathcal{F}_{\xi^e}^{m^e}$.

Proposition 13. For $q \in \{1, 2\}$, the above decomposition provides a bijection between the set $\mathcal{M}_{n,\sigma}^q$ and the set of all pairs

$$(s, (f^e)_{e \in \bar{F}(s)}),$$

where $s \in \mathfrak{S}_{p,q}$ and such that there exist a collection of positive integers $(\xi^e)_{e \in \bar{E}(s)}$ and a collection of nonnegative integers $(m^e)_{e \in \bar{F}(s)}$ satisfying the following:

- for all $e \in \bar{F}(s)$, $f^e \in \mathcal{F}_{\xi^e}^{m^e}$;
- for all $e \in \bar{E}(s)$, $\xi^e = \xi^{e'}$;
- for $0 \leq i \leq p$, $\sum_{e \in \bar{H}_i(s)} \xi^e = \sigma^i$;
- $\sum_{e \in \bar{F}(s)} m^e + \frac{1}{2} \sum_{e \in \bar{E}(s)} \xi^e = n + |\sigma|$.

3.4.2. Decomposition of the labeled map

Let us now turn to the labels. The terminology we use here should become clear in a moment.

Definition 3. We call interior bridge of length $\xi \geq 1$ from $a \in \mathbb{Z}$ to $b \in \mathbb{Z}$ a sequence of integers $(b(0), \dots, b(\xi))$ such that $b(0) = a$, $b(\xi) = b$ and, for all $0 \leq i \leq \xi - 1$, we have $b(i + 1) - b(i) \in \{-1, 0, 1\}$. We write $\mathcal{I}_\xi(a, b)$ the set of these interior bridges.

We call boundary bridge of length $\xi \geq 1$ from $a \in \mathbb{Z}$ to $b \in \mathbb{Z}$ a sequence $(b(0), \dots, b(\xi))$ of integers such that $b(0) = a$, $b(\xi) = b$ and, for all $0 \leq i \leq \xi - 1$, we have $b(i + 1) - b(i) \geq -1$. We write $\mathcal{B}_\xi(a, b)$ the set of these boundary bridges.

Definition 4. A labeled forest is a pair (f, ℓ) where f is a forest and $\ell: V(f) \rightarrow \mathbb{Z}$ is a function satisfying the following:

- for all u lying in the floor of f , $\ell(u) = 0$;
- if u and v are linked by an edge, then $|\ell(u) - \ell(v)| \leq 1$.

We denote by \mathfrak{F}_ξ^m the set of labeled forests of length ξ and mass m .

There is a trivial one-to-one correspondence between the nodes of \mathfrak{m} and the vertices of \mathfrak{s} so that l naturally gives a canonical labeling of the vertices of \mathfrak{s} as follows. Let $v_* \in V(\mathfrak{s})$ be the origin of the root in \mathfrak{s} and let $v \in V(\mathfrak{s})$. We denote by v'_* and $v' \in V(\mathfrak{m})$ the corresponding nodes in \mathfrak{m} and we set $l^v := l(v') - l(v'_*)$. Let $e \in \vec{F}(\mathfrak{s})$. It naturally corresponds to a chain $e_1, e_2, \dots, e_{\xi^e}$ of half-edges in \mathfrak{m} . We define the bridge

$$b^e := (l(e_1^-) - l(v'_*), l(e_2^-) - l(v'_*), \dots, l(e_{\xi^e}^-) - l(v'_*), l(e_{\xi^e}^+) - l(v'_*)).$$

The constraints on l show that, if $e \in \vec{I}(\mathfrak{s})$ then $b^e \in \mathfrak{J}_{\xi^e}(l^{e^-}, l^{e^+})$ and, if $e \in \vec{B}(\mathfrak{s})$ then $b^e \in \mathfrak{B}_{\xi^e}(l^{e^-}, l^{e^+})$. Moreover, the forest \mathfrak{f}^e from last section naturally inherits from l a labeling function l^e defined as follows. Let $u \in V(\mathfrak{f}^e)$ and let $\rho \in V(\mathfrak{f}^e)$ be the root of the tree to which u belongs. These vertices correspond to two vertices u' and $\rho' \in V(\mathfrak{m})$. We set $l^e(u) := l(u') - l(\rho')$. See Figure 11.

Proposition 14. For $q \in \{1, 2\}$, the above decomposition provides a bijection between the set $\mathfrak{M}_{n,\sigma}^q$ and the set of all triples

$$(\mathfrak{s}, (\mathfrak{f}^e, l^e)_{e \in \vec{F}(\mathfrak{s})}, (b^e)_{e \in \vec{F}(\mathfrak{s})}),$$

where $\mathfrak{s} \in \mathfrak{S}_{p,q}$ and such that there exist a collection of positive integers $(\xi^e)_{e \in \vec{E}(\mathfrak{s})}$, a collection of nonnegative integers $(m^e)_{e \in \vec{F}(\mathfrak{s})}$ and a collection of integers $(l^v)_{v \in V(\mathfrak{s})}$ satisfying the following:

- for all $e \in \vec{F}(\mathfrak{s})$, $(\mathfrak{f}^e, l^e) \in \mathfrak{F}_{\xi^e}^{m^e}$;
- for all $e \in \vec{E}(\mathfrak{s})$, $\xi^e = \xi^{e'}$;
- $l^{v_*} = 0$, where v_* is the origin of the root of \mathfrak{s} ;
- for all $e \in \vec{I}(\mathfrak{s})$, $b^e \in \mathfrak{J}_{\xi^e}(l^{e^-}, l^{e^+})$ and $b^e = (b^e(\xi^e), \dots, b^e(0))$;
- for all $e \in \vec{B}(\mathfrak{s})$, $b^e \in \mathfrak{B}_{\xi^e}(l^{e^-}, l^{e^+})$;
- for $0 \leq i \leq p$, $\sum_{e \in \vec{H}_i(\mathfrak{s})} \xi^e = \sigma^i$;
- $\sum_{e \in \vec{F}(\mathfrak{s})} m^e + \frac{1}{2} \sum_{e \in \vec{E}(\mathfrak{s})} \xi^e = n + |\sigma|$.

Note that the three collections of integers are entirely determined by the triple of scheme, forests and bridges.

3.4.3. Encoding by real-valued functions

For $e \in \vec{F}(\mathfrak{s})$ we encode the labeled forest (\mathfrak{f}^e, l^e) by its so-called *contour pair* (C^e, L^e) defined as follows. We see \mathfrak{f}^e as a planar one-faced map with $2m^e + \xi^e$ edges (recall that we add a vertex-tree at the end and join by edges the elements of the floor). First, let $\mathfrak{f}^e(0), \mathfrak{f}^e(1), \dots, \mathfrak{f}^e(2m^e + \xi^e)$ be the vertices of \mathfrak{f}^e read in counterclockwise order around the face, starting at the first corner of the first tree. The *contour function* $C^e : [0, 2m^e + \xi^e] \rightarrow \mathbb{R}_+$ and the *label function* $L^e : [0, 2m^e + \xi^e] \rightarrow \mathbb{R}$ are defined by

$$C^e(i) := d_{\mathfrak{f}^e}(\mathfrak{f}^e(i), \mathfrak{f}^e(2m^e + \xi^e)) \quad \text{and} \quad L^e(i) := l^e(\mathfrak{f}^e(i)), \quad 0 \leq i \leq 2m^e + \xi^e,$$

and linearly interpolated between integer values (see Figure 12).

We also define the function $B^e : [0, \xi^e] \rightarrow \mathbb{R}$ by

$$B^e(i) := b^e(i), \quad 0 \leq i \leq \xi^e,$$




Fig. 12. The contour pair of a labeled forest from $\mathfrak{F}_{7,2}^{20}$. On the right, the paths are dashed on the intervals corresponding to edges linking elements of the floor.

and we linearly interpolate it between integer values. We will use the standard notation

$$\underline{X}(s) := \inf_{0 \leq t \leq s} X(t)$$

for the past infimum of any process X . Remark that the function

$$s \in [0, 2m^e + \xi^e] \mapsto L^e(s) + B^e(\xi^e - \underline{C}^e(s))$$

records the labels up to an additive constant of the part of m that corresponds to f^e . This will become useful in Section 4.3.2.

4. Scaling limit

From now on, we fix an integer $p \geq 0$ and a sequence $\sigma_n = (\sigma_n^1, \dots, \sigma_n^p) \in \mathbb{N}^p$ of p -uples such that

$$\sigma_{(n)}^i := \frac{\sigma_n^i}{\sqrt{2n}} \rightarrow \sigma_\infty^i \in (0, \infty), \quad 1 \leq i \leq p.$$

We set $\sigma_\infty := (\sigma_\infty^1, \dots, \sigma_\infty^p)$ and $\sigma_{(n)} := (\sigma_{(n)}^1, \dots, \sigma_{(n)}^p)$ the rescaled version of σ_n . Recall also that the genus $g \geq 0$ is fixed. We furthermore suppose in this section that $(g, p) \neq (0, 0)$ as this case requires special care.

Remark. Throughout this paper, the notation with a parenthesized n will always refer to suitably rescaled objects and the notation with an ∞ symbol will refer to limiting objects, as in the definitions above.

Let q_n be a random variable uniformly distributed over the set $\mathcal{Q}_{n, \sigma_n}$ and let $v_n^\bullet \in V(q_n)$ be one of its vertices chosen uniformly at random. Let $(m_n, [l_n]) \in \mathfrak{M}_{n, \sigma_n}^1$ be the image of (q_n, v_n^\bullet) through the two-to-one mapping of Proposition 10 and let

$$(\mathfrak{s}_n, (f_n^e, l_n^e)_{e \in \bar{F}(\mathfrak{s}_n)}, (b_n^e)_{e \in \bar{F}(\mathfrak{s}_n)})$$

be the decomposition of $(m_n, [l_n])$ appearing in Proposition 14. We let $(\xi_n^e)_{e \in \bar{E}(\mathfrak{s}_n)}$, $(m_n^e)_{e \in \bar{F}(\mathfrak{s}_n)}$ and $(l_n^v)_{v \in V(\mathfrak{s}_n)}$ be the three collections of integers from Proposition 14. For all e , we also denote by (C_n^e, L_n^e) the contour pair of (f_n^e, l_n^e) as well as B_n^e the interpolation of b_n^e . Finally, we define the rescaled versions of these objects

$$m_{(n)}^e := \frac{2m_n^e + \xi_n^e}{2n}, \quad \xi_{(n)}^e := \frac{\xi_n^e}{\sqrt{2n}}, \quad l_{(n)}^v := \frac{l_n^v}{(8n/9)^{1/4}},$$

$$C_{(n)}^e := \left(\frac{C_n^e(2ns)}{\sqrt{2n}} \right)_{0 \leq s \leq m_{(n)}^e}, \quad L_{(n)}^e := \left(\frac{L_n^e(2ns)}{(8n/9)^{1/4}} \right)_{0 \leq s \leq m_{(n)}^e}, \quad B_{(n)}^e := \left(\frac{B_n^e(\sqrt{2ns})}{(8n/9)^{1/4}} \right)_{0 \leq s \leq \xi_{(n)}^e}.$$

The goal of this section is to give the limit of the joint distribution of these processes. We first need to introduce the limiting processes.

4.1. Brownian bridges, first-passage Brownian bridges, and Brownian snake

We will work on the space $\mathcal{K} := \bigcup_{x \in \mathbb{R}_+} \mathcal{C}([0, x], \mathbb{R})$, endowed with the metric

$$d_{\mathcal{K}}(f, g) := |\zeta(f) - \zeta(g)| + \sup_{y \geq 0} |f(y \wedge \zeta(f)) - g(y \wedge \zeta(g))|,$$

where $\zeta(f)$ denotes the only x such that $f \in \mathcal{C}([0, x], \mathbb{R})$.

We call *Brownian bridge* of length ξ from a to b a standard Brownian motion on $[0, \xi]$ started at a , conditioned on being at b at time ξ (see for example [5,6,12,34]). We also call *first-passage Brownian bridge* of length m from a

to $b < a$ a standard Brownian motion on $[0, m]$ started at a , and conditioned on hitting b for the first time at time m . We refer the reader to [6] for a proper definition of this conditioning, as well as for some convergence results of the discrete analogs.

The so-called *Brownian snake's head* driven by a process $X \in \mathcal{C}([0, x], \mathbb{R})$ may be defined as the process $(X(s), Z(s))_{0 \leq s \leq x}$, where, conditionally given X , the process Z is a centered Gaussian process with covariance function

$$\text{cov}(Z(s), Z(s')) = \inf_{s \wedge s' \leq t \leq s \vee s'} (X(t) - \underline{X}(t)).$$

We refer to [6,20,22] for more details about this process.

4.2. Convergence of the encoding elements

In the limit, we will see that only the schemes that maximize the cardinal of \vec{E} remain.

Definition 5. A scheme is dominant if it has one vertex of degree exactly 1 and if all its other vertices are of degree exactly 3. Let $\mathfrak{S}_{p,1}^*$ and $\mathfrak{S}_{p,2}^*$ denote the sets of dominant schemes of type $(p, 1)$ and $(p, 2)$.

For instance, the two right-most schemes of Figure 9 are the two elements of $\mathfrak{S}_{1,1}^*$ in genus 0. Note that the degree 1-vertex in the previous definition is necessarily an extremity of the root and that the dominant schemes are the ones whose number of edges is maximal (this is a simple consequence of Euler characteristic formula). Moreover, the condition on the degrees implies that every vertex of a dominant scheme is incident to at most one hole. In particular, the holes of a dominant scheme are well “separated” in the sense that their boundaries are pairwise disjoint simple loops, which are connected by some edges.

The compatibility condition on the previous collections of integers lead us to define, for a scheme \mathfrak{s} with root e_* , the set $\mathcal{T}_{\mathfrak{s}}$ of triples

$$\left((m^e)_{e \in \vec{F}(\mathfrak{s})}, (\xi^e)_{e \in \vec{E}(\mathfrak{s})}, (l^v)_{v \in V(\mathfrak{s})} \right) \in \mathbb{R}_+^{\vec{F}(\mathfrak{s})} \times \mathbb{R}_+^{\vec{E}(\mathfrak{s})} \times \mathbb{R}^{V(\mathfrak{s})}$$

such that

- $\sum_{e \in \vec{F}(\mathfrak{s})} m^e = 1,$
- for all $e \in \vec{E}(\mathfrak{s}), \xi^{\bar{e}} = \xi^e,$
- for $0 \leq i \leq p, \sum_{e \in \vec{H}_i(\mathfrak{s})} \xi^e = \sigma_{\infty}^i,$
- $l^{e_*} = 0.$

We define the measure $\mathcal{L}_{\mathfrak{s}}$ on $\mathcal{T}_{\mathfrak{s}}$ as follows. For every $1 \leq i \leq p,$ we distinguish a half-edge $\vec{h}_i \in \vec{H}_i(\mathfrak{s}).$ We also consider an orientation $\check{I}(\mathfrak{s})$ of $\vec{I}(\mathfrak{s}),$ that is, a subset of $\vec{I}(\mathfrak{s})$ containing exactly one half-edge among $\{e, \bar{e}\}$ for every $e \in \vec{I}(\mathfrak{s}).$ There is then an obvious bijection between a part of

$$\mathbb{R}_+^{\vec{F}(\mathfrak{s}) \setminus \{e_*\}} \times \mathbb{R}_+^{\check{I}(\mathfrak{s}) \cup \bigcup_i \vec{H}_i(\mathfrak{s}) \setminus \{\vec{h}_i\}} \times \mathbb{R}^{V(\mathfrak{s}) \setminus \{e_*^-\}}$$

and $\mathcal{T}_{\mathfrak{s}}.$ The measure $\mathcal{L}_{\mathfrak{s}}$ is defined as the push-forward of the Lebesgue measure on this part. We denote by p_a the density of a centered Gaussian variable with variance $a > 0,$ as well as $-q_a$ its derivative:

$$p_a(x) := \frac{1}{\sqrt{2\pi a}} \exp\left(-\frac{x^2}{2a}\right) \quad \text{and} \quad q_a(x) = \frac{x}{a} p_a(x), \quad x \in \mathbb{R}.$$

We let μ be the probability measure on $\bigcup_{\mathfrak{s} \in \mathfrak{S}_{p,1}^*} \{\mathfrak{s}\} \times \mathcal{T}_{\mathfrak{s}}$ defined, for all measurable function $\psi,$ by

$$\mu(\psi) := \frac{1}{\Upsilon} \sum_{\mathfrak{s} \in \mathfrak{S}_{p,1}^*} \int_{\mathcal{T}_{\mathfrak{s}}} d\mathcal{L}_{\mathfrak{s}} \psi(\mathfrak{s}, ((m^e), (\xi^e), (l^v))) \prod_{e \in \vec{F}(\mathfrak{s})} q_{m^e}(\xi^e) \prod_{e \in \check{I}(\mathfrak{s}) \cup \vec{B}(\mathfrak{s})} p_{(\kappa^e)^2 \xi^e}(l^{e^+} - l^{e^-}),$$

where

$$\kappa^e := \begin{cases} 1 & \text{if } e \in \vec{I}(\mathfrak{s}_\infty), \\ \sqrt{3} & \text{if } e \in \vec{B}(\mathfrak{s}_\infty) \end{cases}$$

and

$$\gamma := \sum_{\mathfrak{s} \in \mathfrak{S}_{p,1}^*} \int_{\mathcal{T}_{\mathfrak{s}}} d\mathcal{L}_{\mathfrak{s}} \prod_{e \in \vec{F}(\mathfrak{s})} q_{m^e}(\xi^e) \prod_{e \in \check{I}(\mathfrak{s}) \cup \vec{B}(\mathfrak{s})} P_{(\kappa^e)^2 \xi^e}(l^{e^+} - l^{e^-})$$

is a normalization constant. The measure μ may also be described in terms of Gaussian Free Fields on the metric graphs corresponding to the schemes. We refer to the remark before Section 4.2 in [30] for further details about this fact as it will not be used in this work. In fact, the precise expression of μ will not be needed in what follows; we gave it for self-containment reasons and as it allows to compute the precise distribution of the random variable H appearing in Theorem 3.

Proposition 15. *The random vector*

$$(\mathfrak{s}_n, (m_{(n)}^e)_{e \in \vec{F}(\mathfrak{s}_n)}, (\xi_{(n)}^e)_{e \in \vec{E}(\mathfrak{s}_n)}, (l_{(n)}^v)_{v \in V(\mathfrak{s}_n)}, (C_{(n)}^e, L_{(n)}^e)_{e \in \vec{F}(\mathfrak{s}_n)}, (B_{(n)}^e)_{e \in \vec{B}(\mathfrak{s}_n)})$$

converges in distribution toward a random vector

$$(\mathfrak{s}_\infty, (m_\infty^e)_{e \in \vec{F}(\mathfrak{s}_\infty)}, (\xi_\infty^e)_{e \in \vec{E}(\mathfrak{s}_\infty)}, (l_\infty^v)_{v \in V(\mathfrak{s}_\infty)}, (C_\infty^e, L_\infty^e)_{e \in \vec{F}(\mathfrak{s}_\infty)}, (B_\infty^e)_{e \in \vec{B}(\mathfrak{s}_\infty)})$$

whose law is defined as follows:

- the vector $(\mathfrak{s}_\infty, ((m_\infty^e)_{e \in \vec{F}(\mathfrak{s}_\infty)}, (\xi_\infty^e)_{e \in \vec{E}(\mathfrak{s}_\infty)}, (l_\infty^v)_{v \in V(\mathfrak{s}_\infty)}))$ is distributed according to the probability measure μ ;
- conditionally given this vector,
 - the processes (C_∞^e, L_∞^e) , $e \in \vec{F}(\mathfrak{s}_\infty)$ and B_∞^e , $e \in \check{I}(\mathfrak{s}_\infty) \cup \vec{B}(\mathfrak{s}_\infty)$ are independent,
 - the process (C_∞^e, L_∞^e) has the law of a Brownian snake's head driven by a first-passage Brownian bridge of length m_∞^e from ξ_∞^e to 0,
 - the process B_∞^e has the law of a Brownian bridge of length ξ_∞^e from $l_\infty^{e^-}$ to $l_\infty^{e^+}$, multiplied by the factor κ^e ,
 - the bridges are linked through the relation $B_\infty^e(s) = B_\infty^e(\xi_\infty^e - s)$, $0 \leq s \leq \xi_\infty^e$, whenever $e \in \vec{I}(\mathfrak{s}_\infty)$.

The previous proposition is easily obtained by the method provided in [6] (see in particular Proposition 7 and Section 5, as well as [8], Proposition 7); we leave the details to the reader.

The factor κ^e accounts for the fact that the steps of boundary bridges have a larger variance than the steps of interior bridges. This seemingly harmless factor causes some difficulties for the technical estimates of [8]. Note also that this proposition is the reason why the factor $(8/9)^{1/4}$ appears.

Let us emphasize at this stage that the limiting scheme \mathfrak{s}_∞ is a.s. dominant and, as such, possesses the properties observed after Definition 5.

We obtain a similar statement when performing the two-point mapping instead of the one-point mapping, up to a bias induced by the factor λ of Section 3.2. We will not need this statement in full details and will come back to it during Step 2 of the proof of Proposition 19.

From now on, we apply Skorokhod's representation theorem and assume that the convergence of Proposition 15 holds almost surely. In particular, this entails that $\mathfrak{s}_n = \mathfrak{s}_\infty$ for n large enough.

4.3. Proof of Theorem 1

The whole proof of Theorem 1 is obtained by adding the arguments of [6–8]. The general strategy is borrowed from [23], as well as from [29] for the topology. We will not treat it here in full details and we refer the interested reader to these references for more information. We will only recall the main lines of reasoning and introduce the notions that will become useful in the next sections. We also exclude the case $(g, p) = (0, 0)$, which is degenerate in some sense and requires special definitions. This is the original case treated in [23,29].

4.3.1. *Convergence of the metric space*

The first assertion is proved by the method of Le Gall’s pioneering paper [23]. Recall that q_n is uniformly distributed over \mathcal{Q}_{n,σ_n} , that v_n^\bullet is uniformly distributed over $V(q_n)$ and that $(m_n, [l_n]) \in \mathfrak{M}_{n,\sigma_n}^1$ is the labeled map corresponding to (q_n, v_n^\bullet) . We arrange the corners of the internal face of m_n according to the contour order, starting from the corner preceding the root: this gives a natural (noninjective) ordering of the vertices of m_n , which we write $m_n(0), \dots, m_n(2n + |\sigma_n|)$ with a slight abuse of notation. As the vertex set of m_n corresponds to $V(q_n) \setminus \{v_n^\bullet\}$, this also provides an ordering $q_n(0), \dots, q_n(2n + |\sigma_n|)$ of the vertices of q_n . (The fact that v_n^\bullet is left out will not be important when we take the scaling limit.)

We may equip the set $[0, 1]$ with the pseudo-metric $d_{(n)}$ defined by

$$d_{(n)}(s, t) := \left(\frac{9}{8n}\right)^{1/4} d_{q_n}(q_n((2n + |\sigma_n|)s), q_n((2n + |\sigma_n|)t)), \quad s, t \in \left\{0, \frac{1}{2n + |\sigma_n|}, \dots, 1\right\},$$

and by linear interpolation between these values. It is easy to see that the metric space $[0, 1]_{\{d_{(n)}=0\}}$ equipped with the quotient metric is close in the Gromov–Hausdorff sense to the rescaled metric space corresponding to q_n . As the labels of m_n represent the distances in q_n to the point v_n^\bullet , it turns out that it is possible to bound this pseudo-metric by an explicit function of l_n , which converges toward an explicit process, thanks to Proposition 15. This entails that, up to extraction, $d_{(n)}$ converges toward a pseudo-metric d_∞^σ , the convergence holding jointly with the one of Proposition 15. We define the equivalence relation associated with it by saying that $s \sim t$ if $d_\infty^\sigma(s, t) = 0$, and we set $q_\infty^\sigma := [0, 1]_{/\sim}$. We also denote by $q_\infty^\sigma(s)$ the equivalence class of s in q_∞^σ . We can then show that this yields the first assertion of Theorem 1.

From now on, we fix a subsequence $(n_k)_{k \geq 0}$ along which the previous convergence holds and we let $(q_\infty^\sigma, d_\infty^\sigma)$ be the corresponding limiting space. Recall also that, by Skorokhod’s representation theorem, we assumed that the convergence of Proposition 15 held almost surely.

4.3.2. *Seeing the limit as a quotient of a simpler “map”*

This is the longest and most difficult step. In some sense, we can view it as the continuous version of the mapping from Section 3.1. The idea is to construct a continuous counterpart to m_n and then to perform a continuous version of the mapping: because of the scaling, the counterpart of the arcs added in the discrete mapping is, in the continuous case, identifications of points. The main ingredients come from [23] but we need some extra arguments because of the more intricate combinatorial structure of the quadrangulations we consider. There are mainly two difficulties of different nature. The first difficulty arises from the more complex structure of the map m_n ; we overcome this difficulty with the arguments of [6,7]. The second technicality is caused by the boundary and is harder to grasp at first. The fact that the Brownian motions have different diffusion factors on the boundary edges and on the internal edges of the scheme (the factors κ^e from Proposition 15) induces a technical complication; we use the arguments of [8] to treat this problem.

We now properly define the continuous analog \mathcal{M} of m_n , which was briefly introduced during Section 2.2. Roughly speaking, we glue together Brownian forests coded by the C_∞^e ’s according to the scheme structure. Precisely, let $e^{[1]}, \dots, e^{[|\vec{F}(\mathfrak{s}_\infty)|]}$ be the half-edges of $\vec{F}(\mathfrak{s}_\infty)$ arranged according to the contour order around f^\bullet , starting from the root of \mathfrak{s}_∞ . For every $s \in [0, 1)$, there exists a unique $1 \leq k \leq |\vec{F}(\mathfrak{s}_\infty)|$ such that

$$\sum_{i=1}^{k-1} m_\infty^{e^{[i]}} \leq s < \sum_{i=1}^k m_\infty^{e^{[i]}}.$$

We let $e_s := e^{[k]}$ and $\langle s \rangle := s - \sum_{i=1}^{k-1} m_\infty^{e^{[i]}} \in [0, m_\infty^{e_s})$. By convention, we set $e_1 = e^{[1]}$ and $\langle 1 \rangle = 0$. We define the relation \simeq on $[0, 1]$ as the coarsest equivalence relation for which $s \simeq t$ if one of the following occurs:

$$e_s = e_t \quad \text{and} \quad C_\infty^{e_s} \langle s \rangle = C_\infty^{e_t} \langle t \rangle = \inf_{\{ \langle s \rangle \wedge \langle t \rangle, \langle s \rangle \vee \langle t \rangle \}} C_\infty^{e_s}; \tag{2a}$$

$$e_s = \overline{e_t}, \quad C_\infty^{e_s} \langle s \rangle = \underline{C_\infty^{e_s} \langle s \rangle}, \quad C_\infty^{e_t} \langle t \rangle = \underline{C_\infty^{e_t} \langle t \rangle} \quad \text{and} \quad C_\infty^{e_s} \langle s \rangle = \xi_\infty^{e_t} - C_\infty^{e_t} \langle t \rangle; \tag{2b}$$

$$\langle s \rangle = \langle t \rangle = 0 \quad \text{and} \quad e_s^- = e_t^-, \tag{2c}$$




Fig. 13. On this picture, we can see the floor fl , an example of tree τ_a , and the set $[[a, b]]$.

where we wrote $C_\infty^{e_s}\langle s \rangle$ instead of $C_\infty^{e_s}(\langle s \rangle)$ for short. Equation (2a) identifies numbers coding the same point in one of the forests, Equation (2b) identifies the floors of forests “facing each other”: the numbers s and t should code floor points (second and third equalities) of forests facing each other (first equality) and correspond to the same point (fourth equality). Finally, Equation (2c) identifies the nodes. We set $\mathcal{M} := [0, 1]_{\simeq}$ and, for $s, t \in [0, 1]$, we let $\mathcal{M}(s)$ be the equivalence class of s in the quotient and $\mathcal{M}([s, t]) := \{\mathcal{M}(r) : r \in [s, t]\}$.

The following notions will be used in what follows. They are the continuous counterpart of previous definitions. We call *floor* of \mathcal{M} the set

$$fl := \mathcal{M}(\{s : C_\infty^{e_s}\langle s \rangle = \underline{C}_\infty^{e_s}\langle s \rangle\}).$$

For $a = \mathcal{M}(s) \in \mathcal{M} \setminus fl$, let $l := \inf\{t \leq s : e_t = e_s, \underline{C}_\infty^{e_t}\langle t \rangle = \underline{C}_\infty^{e_s}\langle s \rangle\}$ and $r := \sup\{t \geq s : e_t = e_s, \underline{C}_\infty^{e_t}\langle t \rangle = \underline{C}_\infty^{e_s}\langle s \rangle\}$. We set $\tau_a := \mathcal{M}([l, r])$ and we call *tree* of \mathcal{M} a set of the form τ_a for any $a \in \mathcal{M} \setminus fl$. We say that $[l, r]$ is the *interval coding* the tree τ_a . For a, b in a tree of \mathcal{M} , we denote by $[[a, b]]$ the range of the unique injective path linking a to b . See Figure 13.

Definition 6. *The order of a point $a \in \mathcal{M}$ is the integer*

$$|\{s \in [0, 1] : \mathcal{M}(s) = a\}| + \mathbb{1}_{\{a \in fl, \exists s : a = \mathcal{M}(s), e_s \in \bar{B}(\mathfrak{s}_\infty)\}}.$$

A point of order 1 is called a leaf.

Remark. *The indicator function in the previous definition accounts for the fact that the boundaries of the holes are not visited “from the inside” of the holes. This definition differs from the one we gave in Section 2.2; it is a simple exercise to verify that they are equivalent.*

Because the functions C_∞^e are essentially Brownian motions, it is not hard to see that, a.s., all the points in \mathcal{M} are of order less than 3 and that, if U is uniformly distributed over $[0, 1]$, then $\mathcal{M}(U)$ is a.s. a leaf. Moreover, the floor contains no leaves.

The trivial fact that $m_n((2n + |\sigma_n|)s) = m_n((2n + |\sigma_n|)t) \Rightarrow d_{(n)}(s, t) = 0$ easily passes to the limit in the sense that $s \simeq t \Rightarrow s \sim t$ (see [7], Lemma 6). As a result, we may define a pseudo-metric and an equivalence relation on \mathcal{M} , which we still denote by d_∞^σ and \sim , by setting $d_\infty^\sigma(\mathcal{M}(s), \mathcal{M}(t)) := d_\infty^\sigma(s, t)$ and declaring $\mathcal{M}(s) \sim \mathcal{M}(t)$ if $s \sim t$. It allows us to see q_∞^σ as a quotient of \mathcal{M} : more precisely, the metric space $(q_\infty^\sigma, d_\infty^\sigma)$ is isometric to $(\mathcal{M}_{/\sim}, d_\infty^\sigma)$.

We now describe how the points of \mathcal{M} are identified in the quotient $\mathcal{M}_{/\sim}$. We aim to assign labels to the points of \mathcal{M} . We first define a labeling function on $[0, 1]$ by setting

$$\mathfrak{L}(s) := l_\infty^{e_s^-} + L_\infty^{e_s}\langle s \rangle + B_\infty^{e_s}(\xi_\infty^{e_s} - \underline{C}_\infty^{e_s}\langle s \rangle), \quad 0 \leq s \leq 1.$$

The reason for this definition is the observation we made at the end of Section 3.4.3. Note that this function is continuous.

Lemma 16. *The set of points where \mathcal{L} reaches its minimum is a.s. a singleton.*

We denote by s^\bullet the element of this singleton and we call *base point* the corresponding point $\rho^\bullet := q_\infty^\sigma(s^\bullet)$ in q_∞^σ . It is the continuous analog to the point v_n^\bullet . (Recall that this point is linked to the vertices with minimum label in m_n .) As the labels in m_n represent the distances to v_n^\bullet (up to an additive constant), a simple argument consisting in constructing a discrete approximation in q_n of $q_\infty^\sigma(s)$ shows that

$$d_\infty^\sigma(s, s^\bullet) = \mathcal{L}(s) - \mathcal{L}(s^\bullet). \tag{3}$$

The triangle inequality then entails that $s \sim t$ implies $\mathcal{L}(s) = \mathcal{L}(t)$. A fortiori, $s \simeq t$ implies $\mathcal{L}(s) = \mathcal{L}(t)$ so that we can view \mathcal{L} as a function on \mathcal{M} by setting $\mathcal{L}(\mathcal{M}(s)) := \mathcal{L}(s)$. The following proposition tells which points of \mathcal{M} are identified in q_∞^σ . It is the continuous counterpart of the arc construction of Section 3.3, roughly stating that points are identified if and only if they have the same label and all the points located inbetween when following the contour have greater labels.

Proposition 17. *Let $a = \mathcal{M}(s)$ and $b = \mathcal{M}(t)$ be two distinct points in \mathcal{M} . Then $a \sim b$ if and only if a and b are leaves and*

$$\mathcal{L}(a) = \mathcal{L}(b) \geq \mathcal{L}(c)$$

either for all $c \in \mathcal{M}([s \vee t, s \wedge t])$ or for all $c \in \mathcal{M}([s \wedge t, 1] \cup [0, s \vee t])$.

The two different sets appearing at the end come from the fact that there are two possible ways to follow the contour between a and b : in some sense, one corresponds to linking a to b and the other one to linking b to a .

The proofs of these lemma and proposition are rather technical and use fine estimates on the Brownian snake. Happily, as the building blocks of our space are Brownian forests (the pieces coded by the sets $\overline{\{s < 1: e_s = e\}}$), we may use the results of [8] to quickly conclude.

Proof of Lemma 16. Every half-edge $e \in \vec{F}(\mathfrak{s}_\infty)$ corresponds to a Brownian forest. They are not independent as the forests corresponding to e and \bar{e} when $e \in \vec{I}(\mathfrak{s}_\infty)$ share the same labels on their floors. However, we may consider these two forests as a single one having a density of trees multiplied by two. This merely changes the intensity of the Poisson process of the trees. With this convention, the edges of \mathfrak{s}_∞ correspond to independent forests so that they admit distinct minimums. The proof of [8], Lemma 11, then allows to conclude that the minimum is reached only once a.s. □

Proof of Proposition 17. The proof of this proposition is very similar to the proofs in the previously cited references. We will only give the main idea and refer the reader to these references for more details. The starting point is the following lemma. An *increase point* for a function is a point t such that the function is greater than its value at t on a small interval of the form $[t - \varepsilon, t]$ or $[t, t + \varepsilon]$ for some $\varepsilon > 0$.

Lemma 18. *The coding functions $s \in [0, 1] \mapsto C_\infty^{\varepsilon_s}(s)$ and \mathcal{L} do not share any increase points.*

This lemma is shown thanks to [8], Lemma 12, which states the result for a forest. Beware that, as explained in the latter reference, the result only holds for values of $\kappa^e \leq \sqrt{3}$, which is the case here. As the pair of functions considered here is a concatenation of pairs of functions for which the results holds, the result is straightforward. Some care is needed at the extremities of the intervals; we can show that such points are never increase points for the label process.

The second property we need is a property roughly stating the following. If a and b have the same label, if there exists a subtree τ rooted at ρ such that $\inf_\tau \mathcal{L} < \mathcal{L}(a) < \mathcal{L}(\rho)$ and if, for infinitely many n 's, there exists a geodesic from a_n to b_n that completely passes through τ_n , then $a \approx b$. Here, a_n, b_n and τ_n designate discrete approximations in m_n of a, b and τ . See [8], Lemma 15, for a rigorous statement. As explained in [8], in order to prove this property

in our case, we only need an analog of [8], Lemma 31, which is easily obtained by decomposing \mathfrak{L} according to the half-edges of \mathfrak{s}_∞ .

These two properties together allow us to conclude as in [8], Section 4.4. The idea, coming from [23], can be sketched as follows. The absence of common increase points entails with a little work that a point a of order at least 2 will be “surrounded” by subtrees arbitrarily close with a root having a label strictly greater than $\mathfrak{L}(a)$ and a minimum label strictly smaller than $\mathfrak{L}(a)$. As a result, it cannot be identified with any other point by the previous property. We then argue by contradiction and suppose that $a = \mathcal{M}(s)$ and $b = \mathcal{M}(t)$ are identified leaves but such that the property of the statement does not hold. After some work, it follows that on both $\mathcal{M}([s \vee t, s \wedge t])$ and $\mathcal{M}([s \wedge t, 1] \cup [0, s \vee t])$ there exists a subtree with properties similar as above and we conclude that $a \approx b$, which is a contradiction. \square

4.3.3. Topology of the limit

We follow the approach of [29] to identify the topology of the limit. The first observation is that the metric space associated with q_n is not far from being homeomorphic to $\Sigma_{g,p}^\partial$. Informally, we fill the internal faces with small cells and the holes with annuli. The rigorous way to proceed is to consider, for every internal face a copy of the space $[0, 1]^3 \setminus ((0, 1)^2 \times [0, 1])$ and for every hole h_i a space $(\mathcal{G}_i \times [0, 1]) \setminus (\mathring{\mathcal{G}}_i \times [0, 1])$ where \mathcal{G}_i is a regular $2\sigma_n^i$ -sided polygon with unit length edges embedded in \mathbb{R}^2 . We endow these spaces with the intrinsic metric inherited from the Euclidean metric, rescaled by $(9/8n)^{1/4}$. We can then define a surface (with a boundary) by identifying, according to the map structure, the boundaries of these spaces and endow it with the quotient metric. The metric space we obtain is homeomorphic to $\Sigma_{g,p}^\partial$ and the choice of cells and annuli we made ensures that, on the set of points corresponding to the vertices of the map, the metric coincides with the graph metric of the map (in other words, we did not create any shortcuts as it would have been the case if we simply had glued Euclidean polygons, for instance). As a result, this surface is at Gromov–Hausdorff distance $o(1)$ from $(V(q_n), (9/8n)^{1/4}d_{q_n})$. See [8], Section 5.2, for more details.

The second step is a criterion ensuring that the Gromov–Hausdorff limit of a sequence of metric spaces all homeomorphic to some surface is also homeomorphic to this surface. This criterion is called regularity and was introduced by Whyburn [36] and later studied by Begle [4]. We will not go into much detail as the arguments are exactly the same as in [7,8]. The only difference is that, a priori, the regularity criterion only works for $\Sigma_{g,0}^\partial$ ([7], Proposition 19) and for $\Sigma_{0,1}^\partial$ ([8], Proposition 16). We may adapt it to work for $\Sigma_{g,p}^\partial$ by noticing that $\Sigma_{g,p}^\partial$ can be constructed by removing from $\Sigma_{g,0}^\partial$ the interior of p disjoint subspaces homeomorphic to $\Sigma_{0,1}^\partial$. We first “fill” the holes (formally by changing $(\mathcal{G}_i \times [0, 1]) \setminus (\mathring{\mathcal{G}}_i \times [0, 1])$ into $(\mathcal{G}_i \times [0, 1]) \setminus (\mathring{\mathcal{G}}_i \times [0, 1])$) and show a convergence toward a surface homeomorphic to $\Sigma_{g,0}^\partial$ thanks to the first theorem. We then show that the p closed holes (corresponding in the filled surface to $\mathcal{G}_i \times \{1\}$) converge to p disjoint disks thanks to the second theorem. In a few words, this comes from the fact that the holes of q_n correspond to the holes of m_n and, if they were not converging to disjoint disks, in the limit, we would either have holes “sharing” a vertex in the scheme or we would have points of the floor of \mathcal{M} identified together. Neither of these situations can happen: as observed after Definition 5, there are no vertices incident to two different holes in a dominant scheme, and Proposition 17 prohibits the latter situation.

4.3.4. Hausdorff dimension of the limit

To compute the Hausdorff dimension of the limit, we use the same method as in [26]. The idea is roughly the following. To prove that the Hausdorff dimension is less than 4, we use the fact that the processes L_∞^e of Proposition 15 are almost surely α -Hölder for all $\alpha \in (0, 1/4)$, yielding that the canonical projection from $([0, 1], |\cdot|)$ to $(q_\infty^\sigma, d_\infty^\sigma)$ is also α -Hölder for the same values of α . To prove that it is greater than 4, we show that the size of the balls of diameter δ is of order δ^4 . To see this, we first bound from below the distances in terms of label variation along the branches of \mathcal{M} , and then use twice the law of the iterated logarithm: this tells us that, for a fixed $s \in [0, 1]$, the points outside of the set $[s - \delta^4, s + \delta^4]$ code points that are at distance at least δ^2 from $q_\infty^\sigma(s)$ in \mathcal{M} , so that their distance from $q_\infty^\sigma(s)$ is at least δ in the map. See [6], Section 6.4, for a complete proof.

We now turn to the boundary. Let $\pi: \mathcal{M} \rightarrow q_\infty^\sigma$ be the canonical projection. A straightforward adaptation of [8], Proposition 21, shows that the boundary of q_∞^σ is

$$\partial q_\infty^\sigma = \pi(\mathcal{M}(\{s: e_s \in \vec{B}(\mathfrak{s}_\infty)\}) \cap fl). \quad (4)$$

For every connected component of the boundary, we may use the same method as in [8] to conclude that its dimension is a.s. equal to 2.

5. Typical geodesics from the base point

From now on, we work on the set of full probability where \mathcal{L} reaches its minimum only once. Recall that π denotes the canonical projection from \mathcal{M} to q_∞^σ , that s^\bullet is the unique point where \mathcal{L} reaches its minimum and that we set $\rho^\bullet := q_\infty^\sigma(s^\bullet)$. We also include back the case $(g, p) = (0, 0)$ from now on. In this case, $(\mathfrak{m}_n, [l_n]) \in \mathfrak{M}_{n, \emptyset}^1$ is a uniform labeled tree and its rescaled contour pair converges toward the Brownian snake (\mathfrak{e}, Z) driven by a normalized Brownian excursion. Thus \mathcal{M} is Aldous's CRT coded by the excursion \mathfrak{e} and $\mathcal{L} := Z$. Lemma 16 and Proposition 17 also hold (see [23]).

5.1. Uniqueness

Proposition 19. *Let S be uniformly distributed over $[0, 1]$ and independent of $(q_\infty^\sigma, d_\infty^\sigma)$. Then, almost surely, there is only one geodesic from the base point ρ^\bullet to $X := q_\infty^\sigma(S)$ in q_∞^σ .*

Proof. This was shown in [30] in the slightly different context of Boltzmann maps. We reformulate the ideas to our setting.

It will be sufficient to show that, for any rational numbers $0 < \alpha < \beta$, the event $\mathcal{E}_{\alpha, \beta}$ that there exist, for every $u \in [\alpha, \beta]$, at least two different points y and $y' \in q_\infty^\sigma$ such that $d_\infty^\sigma(\rho^\bullet, y) = d_\infty^\sigma(\rho^\bullet, y') = u$ and $d_\infty^\sigma(\rho^\bullet, y) + d_\infty^\sigma(y, X) = d_\infty^\sigma(\rho^\bullet, y') + d_\infty^\sigma(y', X) = d_\infty^\sigma(\rho^\bullet, X)$ has probability 0. For the remaining of the proof, we fix rational numbers $0 < \alpha < \beta$ and we consider a random variable U uniformly distributed over $[\alpha, \beta]$ and independent of all other variables.

The idea is that the previous points y and y' correspond to global minimums of the labels on the interface between the two internal faces of the map obtained by a continuous analog of the two-point mapping of Section 3.2. As the labels are essentially Brownian, this almost surely never happens. Let us proceed rigorously.

Step 1. We go back to the discrete picture and start with the construction of a suitable approximation of X in q_n . We want to find a sequence (S_n) of integers such that $q_n(S_n)$ is roughly uniformly distributed over $V(q_n)$ and $S_n/2n \rightarrow S$, at least along a subsequence. When \mathfrak{m}_n is a tree, this is quite easy. In the general case, it is a bit trickier; we proceed as follows. Let w be an arbitrarily chosen point at d_{q_n} -distance 1 from v_n^\bullet . We claim that we can find a sequence (S_n) of integers such that $X_n := q_n(S_n)$ is uniformly distributed over $V(q_n) \setminus \{v_n^\bullet, w\}$ and $S_n/2n \rightarrow S$ along some subsequence. To do so, we will associate with every of these $n + |\sigma_n| - p - 2g$ vertices exactly two unit length left-open subintervals of $[0, 2(n + |\sigma_n| - p - 2g)]$ as follows. It is possible to remove $2g + p$ edges from \mathfrak{m}_n without disconnecting the map; we obtain a plane tree. In the resulting tree, we bijectively associate with every vertex different from w an edge in such a way that the edge is incident to the corresponding vertex. Then, in the contour of f^\bullet , an edge of \mathfrak{m}_n corresponds to one or two unit length subintervals of $[0, 2n + |\sigma_n|]$ (one if the edge is incident to a hole, two otherwise). We suppose that $|\sigma_n| \geq p + 2g$, which happens for large n , and we associate with every vertex of \mathfrak{m}_n the corresponding subintervals. We associate with the vertices corresponding to only one subinterval a second subinterval arbitrarily chosen among the remaining ones. Let k be such that $2(n + |\sigma_n| - p - 2g)S \in (k, k + 1]$. If the vertex corresponding to $(k, k + 1]$ is either $\mathfrak{m}_n(k)$ or $\mathfrak{m}_n(k + 1)$, we set S_n to k or $k + 1$ accordingly. Otherwise, we arbitrarily set S_n such that $\mathfrak{m}_n(S_n)$ corresponds to $(k, k + 1]$. With this construction, X_n is clearly uniformly distributed over $V(q_n) \setminus \{v_n^\bullet, w\}$ as claimed. Moreover, the probability that we cannot extract a subsequence along which $S_n/2n \rightarrow S$ is bounded by

$$\mathbb{P}\left(\liminf_{n \rightarrow \infty} \{|S_n - 2(n + |\sigma_n| - p - 2g)S| > 1\}\right) \leq \liminf_{n \rightarrow \infty} \frac{|\sigma_n|}{2(n + |\sigma_n| - p - 2g)} = 0.$$

So, up to discarding a zero-probability event, we may suppose that $S_n/2n \rightarrow S$ along a subsequence of $(n_k)_{k \geq 0}$.

Step 2. We look at the scaling limit of the map obtained by the two-point mapping. We set $\lambda_n := \lfloor (8n/9)^{1/4} U \rfloor$. If $1 \leq \lambda_n \leq d_{q_n}(v_n^\bullet, X_n) - 1$, it is possible to define $(\mathfrak{m}_n^{\bullet\bullet}, [l_n^{\bullet\bullet}]) \in \mathfrak{M}_{n, \sigma_n}^2$ as the labeled map corresponding to $(q_n, (v_n^\bullet, X_n), \lambda_n)$ via the two-point mapping of Section 3.2.

We decompose a labeled map $(\mathfrak{m}, [l]) \in \mathfrak{M}_{n, \sigma_n}^2$ into a scheme and a collection of integers and continuous functions as in Section 4. We denote by $\mathcal{V}_{(n)}(\mathfrak{m}, [l])$ the corresponding rescaled vector (the analog of the one from Proposition 15). By convention, we set $\mathcal{V}_{(n)}(\mathfrak{m}_n^{\bullet\bullet}, [l_n^{\bullet\bullet}]) := 0$ whenever $(\mathfrak{m}_n^{\bullet\bullet}, [l_n^{\bullet\bullet}])$ is not defined (that is, whenever $\lambda_n \notin \{1, \dots, d_{q_n}(v_n^\bullet, X_n) - 1\}$). We compare the distribution of $\mathcal{V}_{(n)}(\mathfrak{m}_n^{\bullet\bullet}, [l_n^{\bullet\bullet}])$ with the distribution of $\mathcal{V}_{(n)}(M_n, [L_n])$,

where $(M_n, [L_n])$ denotes a random variable uniformly distributed over $\mathfrak{M}_{n, \sigma_n}^2$: For any bounded measurable function F such that $F(0) = 0$, a straightforward computation yields that

$$\frac{\mathbb{E}[F(\mathcal{V}_{(n)}(\mathfrak{m}_n^{\bullet\bullet}, [\mathfrak{l}_n^{\bullet\bullet}]))]}{\mathbb{P}(\lambda_n < d_{q_n}(v_n^\bullet, X_n))} = \frac{\mathbb{E}[F(\mathcal{V}_{(n)}(M_n, [L_n])) \Delta_n(\mathcal{V}_{(n)}(M_n, [L_n]))]}{\mathbb{E}[\Delta_n(\mathcal{V}_{(n)}(M_n, [L_n]))]},$$

where Δ_n is defined as follows. For $(\mathfrak{m}, [\mathfrak{l}]) \in \mathfrak{M}_{n, \sigma_n}^2$, the number $\Delta_n(\mathcal{V}_{(n)}(\mathfrak{m}, [\mathfrak{l}]))$ is $(8n/9)^{1/4}$ times the probability that λ_n equals the λ corresponding to $(\mathfrak{m}, [\mathfrak{l}])$ via the two-point mapping. It is a simple piecewise linear function of $(9/8n)^{1/4}\lambda$, equal to 0 on $(-\infty, \alpha - (9/8n)^{1/4}] \cup [\beta, \infty)$ and equal to $1/(\beta - \alpha)$ on $[\alpha, \beta - (9/8n)^{1/4}]$. Moreover, by Lemma 12(iii), $(9/8n)^{1/4}(\lambda - 1)$ may be written as a deterministic continuous function Λ of $\mathcal{V}_{(n)}(\mathfrak{m}, [\mathfrak{l}])$. The explicit expression of Λ is quite intricate and will not be needed: loosely speaking, Λ assigns to a vector the minimum of the labels on the interface between the two faces minus the minimum of the labels inside f^\bullet .

Using the same machinery as in Section 4, we obtain that the random vector $\mathcal{V}_{(n)}(M_n, [L_n])$ converges to a random vector \mathcal{V}_∞ whose precise distribution will not be needed. We easily deduce that $\mathcal{V}_{(n)}(\mathfrak{m}_n^{\bullet\bullet}, [\mathfrak{l}_n^{\bullet\bullet}])$ converges weakly toward a random vector $\mathcal{V}_\infty^{\bullet\bullet}$ with distribution

$$\mathbb{P}(U > d_\infty^\sigma(\rho^\bullet, X))\delta_0 + \mathbb{P}(U < d_\infty^\sigma(\rho^\bullet, X))\mathcal{L},$$

where \mathcal{L} is the probability distribution given by

$$\mathcal{L}(G) = \mathbb{E}\left[G(\mathcal{V}_\infty) \frac{\mathbb{1}_{\{\alpha < \Lambda(\mathcal{V}_\infty) < \beta\}}}{\mathbb{P}(\alpha < \Lambda(\mathcal{V}_\infty) < \beta)}\right] \tag{5}$$

for any measurable function G . In particular, the distribution \mathcal{L} is absolutely continuous with respect to the distribution of \mathcal{V}_∞ so that every property holding a.s. for \mathcal{V}_∞ also holds a.s. under \mathcal{L} .

Step 3. We now work on the event $\mathcal{E}_{\alpha, \beta}$. We denote by y and y' two different points satisfying the relations appearing in the definition of $\mathcal{E}_{\alpha, \beta}$ for $u = U$ and we let t and t' be such that $y = q_\infty(t)$ and $y' = q_\infty(t')$. We construct discrete approximations of y and y' as follows.

As $0 < U < d_\infty^\sigma(\rho^\bullet, X)$, we see that $(\mathfrak{m}_n^{\bullet\bullet}, [\mathfrak{l}_n^{\bullet\bullet}])$ is always defined for n large enough. We denote by f_n^\bullet and $f_n^{\bullet\bullet}$ its two internal faces corresponding respectively to v_n^\bullet and X_n . We want to construct an approximation y_n on the interface $f_n^\bullet \cap f_n^{\bullet\bullet}$. Let $z_n := q_n(\lfloor 2nt \rfloor)$. If $z_n \in f_n^\bullet \cap f_n^{\bullet\bullet}$, then we set $y_n := z_n$. Otherwise, let suppose for now that $z_n \in f_n^{\bullet\bullet}$. We consider a geodesic from z_n to v_n^\bullet and we let y_n be a point on this geodesic that belongs both to f_n^\bullet and $f_n^{\bullet\bullet}$. By construction and by Lemma 12(ii), we have $\lambda_n \leq d_{q_n}(v_n^\bullet, y_n) \leq d_{q_n}(v_n^\bullet, z_n)$. The triangle inequality yields that $d_{q_n}(y_n, z_n) = o(n^{1/4})$ and another application of Lemma 12 gives

$$0 \leq \mathfrak{l}_n^{\bullet\bullet}(y_n) - \min_{f_n^\bullet \cap f_n^{\bullet\bullet}} \mathfrak{l}_n^{\bullet\bullet} = o(n^{1/4}). \tag{6}$$

Finally, if $z_n \in f_n^\bullet$, we construct y_n in a symmetrical way and obtain the same equations by similar manipulations.

We also construct the points z'_n and y'_n by applying the same method to t' instead of t . As $(9/8n)^{1/4}d_{q_n}(z_n, z'_n) \rightarrow d_\infty^\sigma(y, y')$, we also obtain that $(9/8n)^{1/4}d_{q_n}(y_n, y'_n) \rightarrow d_\infty^\sigma(y, y') > 0$.

The interface $f_n^\bullet \cap f_n^{\bullet\bullet}$ is a ‘‘cycle’’ in the map $\mathfrak{m}_n^{\bullet\bullet}$. We parametrize this cycle by $[0, 1]$, the point 0 being for example a node chosen in a deterministic way from the scheme. We denote by $r_{(n)} \in [0, 1)$ and $r'_{(n)} \in [0, 1)$ the real numbers corresponding to y_n and y'_n . Up to further extraction, we may suppose that $r_{(n)} \rightarrow r$ and $r'_{(n)} \rightarrow r'$. As $(9/8n)^{1/4}d_{q_n}(y_n, y'_n)$ converges toward a positive number, we have $r \neq r'$: This is obtained by the argument we used in Section 4.3.2. In short, we may construct an equivalence relation from the limiting vector $\mathcal{V}_\infty^{\bullet\bullet}$ by a construction similar to the one of Section 4.3.2 and show that two points identified by this relation are also identified in q_∞^σ . In other words, because two corners identified in $\mathfrak{m}_n^{\bullet\bullet}$ correspond to the same point in q_n , the Brownian surface q_∞^σ may be seen as a quotient of a continuous analog of $\mathfrak{m}_n^{\bullet\bullet}$. This is very similar to what we did in Section 4.3.2 so that we leave the details to the reader.

Now, Step 2 implies that the labels on the interface $f_n^\bullet \cap f_n^{\bullet\bullet}$ converge in distribution, once properly rescaled, toward a concatenation of some functions of $\mathcal{V}_\infty^{\bullet\bullet}$ (up to an irrelevant time scaling because of the parametrization by $[0, 1]$). Let us denote by $B_\infty^{\bullet\bullet} : [0, 1] \rightarrow \mathbb{R}$ this limiting function. Equation (6) yields that $B_\infty^{\bullet\bullet}(r) = B_\infty^{\bullet\bullet}(r') = \min B_\infty^{\bullet\bullet}$. As, on the event $\{\mathcal{V}_\infty^{\bullet\bullet} \neq 0\}$, the process $B_\infty^{\bullet\bullet}$ is, up to the bias (5), a Brownian bridge, this means that we are on an event of probability zero, as wanted. \square

5.2. Simple geodesics

We use in this section the terminology and ideas of [24], developed in the spherical case and readily adaptable to our more general setting. We define simple geodesics and give some elementary properties. Recall that $d_\infty^\sigma(s^\bullet, s) = \mathfrak{L}(s) - \mathfrak{L}(s^\bullet)$. For $s, t \in [0, 1]$, we set

$$[s \rightarrow t] := \begin{cases} [s, t] & \text{if } s \leq t, \\ [s, 1] \cup [0, t] & \text{if } t < s. \end{cases}$$

Definition 7. *The simple geodesic of index $s \in [0, 1]$ is the path Φ_s defined by*

$$\Phi_s(w) := q_\infty^\sigma\left(\sup\left\{r: \inf_{[r \rightarrow s]} \mathfrak{L} = \mathfrak{L}(s^\bullet) + w\right\}\right), \quad 0 \leq w \leq d_\infty^\sigma(s^\bullet, s).$$

It is not very hard to see that Φ_s is a geodesic from ρ^\bullet to $q_\infty^\sigma(s)$ (see [24], Section 4.1, for more details). If we set $t := \sup\{r: \inf_{[r \rightarrow s]} \mathfrak{L} = \mathfrak{L}(s^\bullet) + w\}$ then we must have $\mathfrak{L}(t) = \inf_{[t \rightarrow s]} \mathfrak{L}$ by continuity. If $w < d_\infty^\sigma(s^\bullet, s)$, then $t \neq s$ so that t is an increase point of \mathfrak{L} and, by Lemma 18, $\mathcal{M}(t)$ is a leaf. This entails that Φ_s only visit images of leaves, except possibly at its end $q_\infty^\sigma(s)$. Moreover, if $t' := \inf\{r: \inf_{[s \rightarrow r]} \mathfrak{L} = \mathfrak{L}(s^\bullet) + w\}$ then Proposition 17 implies that $t \sim t'$. This provides a dual definition of Φ_s :

$$\Phi_s(w) = q_\infty^\sigma\left(\inf\left\{r: \inf_{[s \rightarrow r]} \mathfrak{L} = \mathfrak{L}(s^\bullet) + w\right\}\right), \quad 0 \leq w \leq d_\infty^\sigma(s^\bullet, s).$$

Observe that, if t'' is such that $\mathfrak{L}(t'') = \inf_{[t'' \rightarrow s]} \mathfrak{L}$, then $t'' \sim \sup\{r: \inf_{[r \rightarrow s]} \mathfrak{L} = \mathfrak{L}(s^\bullet) + (\mathfrak{L}(t'') - \mathfrak{L}(s^\bullet))\}$. As a result, the range of Φ_s is given by

$$\Phi_s([0, d_\infty^\sigma(s^\bullet, s)]) = q_\infty^\sigma\left(\left\{t: \mathfrak{L}(t) = \inf_{[t \rightarrow s]} \mathfrak{L}\right\}\right) = q_\infty^\sigma\left(\left\{t: \mathfrak{L}(t) = \inf_{[s \rightarrow t]} \mathfrak{L}\right\}\right).$$

Finally, let $s \sim s'$ be two distinct numbers. If $s \not\sim s'$ then, by Proposition 17, we have either $\mathfrak{L}(s) = \mathfrak{L}(s') = \inf_{[s \rightarrow s']} \mathfrak{L}$ or $\mathfrak{L}(s) = \mathfrak{L}(s') = \inf_{[s' \rightarrow s]} \mathfrak{L}$; thus $\Phi_s = \Phi_{s'}$. If $s \simeq s'$ then s and s' cannot be increase points of \mathfrak{L} , so that both $\inf_{[s \rightarrow s']} \mathfrak{L} < \mathfrak{L}(s)$ and $\inf_{[s' \rightarrow s]} \mathfrak{L} < \mathfrak{L}(s)$. It is easy to see that $\Phi_s(w) = \Phi_{s'}(w)$ if and only if $w \in [0, \inf_{[s \rightarrow s']} \mathfrak{L} \vee \inf_{[s' \rightarrow s]} \mathfrak{L} - \mathfrak{L}(s^\bullet)] \cup \{\mathfrak{L}(s) - \mathfrak{L}(s^\bullet)\}$. In particular, $\Phi_s \neq \Phi_{s'}$.

Let $(s_i)_{i \geq 0}$ be a sequence of i.i.d. random variables uniformly distributed over $[0, 1]$. We denote by $a_i := \mathcal{M}(s_i)$ and $x_i := q_\infty^\sigma(s_i)$ the corresponding points in \mathcal{M} and q_∞^σ . By Proposition 19, it is immediate that almost surely, for all i , Φ_{s_i} is the only geodesic from ρ^\bullet to x_i . Note also that, almost surely, the set $\{s_i: i \geq 0\}$ is a dense subset of $[0, 1]$. Up to discarding yet another event of zero probability, we will suppose that both previous properties hold.

In what follows, we will show that it is actually sufficient to know that the geodesics to the x_i 's are simple geodesics in order to conclude that all the geodesics from ρ^\bullet are simple geodesics.

6. General geodesics from the base point

6.1. Geodesics from the base point to images of leaves

In this section, we treat the case of the image under π of the set of leaves in \mathcal{M} . Let $s \in [0, 1]$ be such that $a := \mathcal{M}(s)$ is a leaf. We will show that Φ_s is the only geodesic from ρ^\bullet to $x := q_\infty^\sigma(s)$ in q_∞^σ . There is nothing to prove in the case $s = s^\bullet$, so that we suppose $s \neq s^\bullet$. The first step is to find points a_i and a_j before and after a (in the contour of \mathcal{M}), arbitrarily close to a and such that no geodesics from ρ^\bullet to x intersect the set $\pi([a_i, a_j])$. Note that, as a is a leaf, the tree τ_a is well defined.

Lemma 20. *Let $\varepsilon > 0$ be such that $[s - \varepsilon, s + \varepsilon]$ is included in the interval coding τ_a . There exist i and j satisfying $s - \varepsilon < s_i < s < s_j < s + \varepsilon$ and, for all $b \in [[a_i, a_j]]$,*

$$d_\infty^\sigma(\rho^\bullet, a) < d_\infty^\sigma(\rho^\bullet, b) + d_\infty^\sigma(b, a).$$




Fig. 14. Notation used in the proof of Lemma 20.

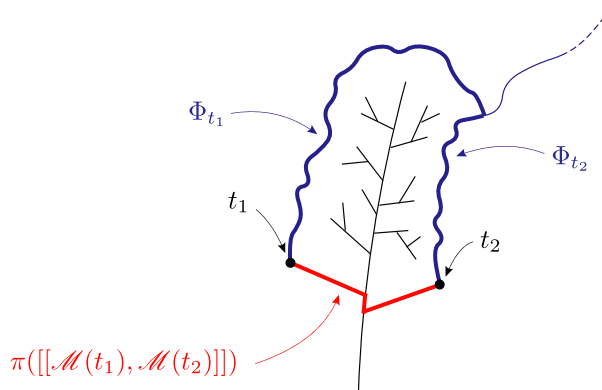


Fig. 15. Topological boundary of $q_\infty^\sigma([t_1, t_2])$.

Proof. We argue by contradiction and suppose that there exists an $\varepsilon > 0$ such that $[s - \varepsilon, s + \varepsilon]$ is included in the interval coding τ_a , and that for all i and j satisfying $s - \varepsilon < s_i < s < s_j < s + \varepsilon$, we can find $b \in [[a_i, a_j]]$ for which $d_\infty^\sigma(\rho^\bullet, a) = d_\infty^\sigma(\rho^\bullet, b) + d_\infty^\sigma(b, a)$. For $0 \leq \xi \leq C_\infty^{e_s}(s) - \underline{C}_\infty^{e_s}(s)$, we define

$$l_\xi := \sup\{t \leq s : C_\infty^{e_t}(t) = C_\infty^{e_s}(s) - \xi\} \quad \text{and} \quad r_\xi := \inf\{t \geq s : C_\infty^{e_t}(t) = C_\infty^{e_s}(s) - \xi\}$$

and we set $c_\xi := \mathcal{M}(l_\xi) = \mathcal{M}(r_\xi)$. See Figure 14. As a is a leaf, there exists $\xi_0 > 0$ such that $s - \varepsilon < l_\xi < s < r_\xi < s + \varepsilon$ as soon as $\xi \in (0, \xi_0]$. Until further notice, we fix such a ξ .

Let $0 < \eta < (s - l_\xi) \wedge (r_\xi - s)$ and i, j be such that $l_\xi \leq s_i < l_\xi + \eta$ and $r_\xi - \eta < s_j \leq r_\xi$. By hypothesis, we can find $b_\eta \in [[a_i, a_j]]$ such that $d_\infty^\sigma(\rho^\bullet, a) = d_\infty^\sigma(\rho^\bullet, b_\eta) + d_\infty^\sigma(b_\eta, a)$. We can furthermore find $t_\eta \in [l_\xi, s_i] \cup [s_j, r_\xi]$ such that $b_\eta := \mathcal{M}(t_\eta)$. As $\eta \rightarrow 0$, t_η converges, up to extraction, either to l_ξ or to r_ξ . This implies that $d_\infty^\sigma(c_\xi, a) = d_\infty^\sigma(\rho^\bullet, a) - d_\infty^\sigma(\rho^\bullet, c_\xi) = \mathfrak{L}(a) - \mathfrak{L}(c_\xi)$, by (3). Moreover, the classical ‘‘cactus bound’’ (see [19] or [6], Lemma 20, for a similar statement) states that $d_\infty^\sigma(c_\xi, a) \geq \mathfrak{L}(c_\xi) + \mathfrak{L}(a) - 2 \min_{[[c_\xi, a]]} \mathfrak{L}$, so that $\mathfrak{L}(c_\xi) = \min_{[[c_\xi, a]]} \mathfrak{L}$. As a result, the process $\xi \in [0, \xi_0] \mapsto \mathfrak{L}(c_\xi)$ is nonincreasing. Basic properties of Brownian snakes show that this is a contradiction. \square

We will need to consider sets of the form $\mathcal{M}([s_i, s_j])$. The following lemma identifies the topological boundary of the image in q_∞^σ of such a set.

Lemma 21. *Let τ be a tree of \mathcal{M} coded by $[l, r]$ and $l \leq t_1 \leq t_2 \leq r$. We set $w' := \inf\{w : \Phi_{t_1}(w) \neq \Phi_{t_2}(w)\}$. Then the topological boundary in q_∞^σ of the subset $q_\infty^\sigma([t_1, t_2])$ is*

$$\partial(q_\infty^\sigma([t_1, t_2])) = \pi([\mathcal{M}(t_1), \mathcal{M}(t_2)]) \cup \Phi_{t_1}([w', d_\infty^\sigma(s^\bullet, t_1)]) \cup \Phi_{t_2}([w', d_\infty^\sigma(s^\bullet, t_2)]).$$

In other words, the boundary of $q_\infty^\sigma([t_1, t_2])$ is composed of three parts: the image of the set $[[\mathcal{M}(t_1), \mathcal{M}(t_2)]]$ and the ranges of Φ_{t_1} and Φ_{t_2} after they split. See Figure 15.

Proof of Lemma 21. The properties stated during Section 5.2 show that

$$\Phi_{t_1}([\mathcal{M}(w'), d_\infty^\sigma(s^\bullet, t_1)]) = q_\infty^\sigma\left(\left\{t \in [t_1, t_2]: \mathfrak{L}(t) = \inf_{[t_1, t]} \mathfrak{L}\right\}\right).$$

Let $t \in [t_1, t_2]$ be such that $\mathfrak{L}(t) = \inf_{[t_1, t]} \mathfrak{L}$, and let us define $t' := \sup\{r: \inf_{[r \rightarrow t_1]} \mathfrak{L} = \mathfrak{L}(t)\}$. Then $t' \sim t$, so that $q_\infty^\sigma(t') = q_\infty^\sigma(t)$. (Possibly, $t' = t$ if $t \in \{t_1, t_2\}$.) For all integer k , we can find a point $t_k \in [t' - 1/k, t']$ such that $\mathcal{M}(t_k)$ is a leaf not identified with any other points by \sim (it suffices to select a point that is an increase point for neither functions of Lemma 18). As a result, $q_\infty^\sigma(t_k) \notin q_\infty^\sigma([t_1, t_2])$ and $d_\infty^\sigma(q_\infty^\sigma(t_k), q_\infty^\sigma(t)) \rightarrow 0$ as $k \rightarrow \infty$, so that $q_\infty^\sigma(t) \in \partial(q_\infty^\sigma([t_1, t_2]))$. The same argument shows that $\Phi_{t_2}([\mathcal{M}(w'), d_\infty^\sigma(s^\bullet, t_2)]) \subseteq \partial(q_\infty^\sigma([t_1, t_2]))$ as well. To see that $\pi([\mathcal{M}(t_1), \mathcal{M}(t_2)]) \subseteq \partial(q_\infty^\sigma([t_1, t_2]))$, we notice that

$$\pi([\mathcal{M}(t_1), \mathcal{M}(t_2)]) = q_\infty^\sigma\left(\left\{t \in [t_1, t_2]: C_\infty^{e_t}(t) = \inf_{r \in [t_1, t]} C_\infty^{e_r}(r) \text{ or } C_\infty^{e_t}(t) = \inf_{r \in [t, t_2]} C_\infty^{e_r}(r)\right\}\right)$$

and apply a similar argument.

Conversely, let $y \in \partial(q_\infty^\sigma([t_1, t_2]))$ and $t \in [t_1, t_2]$ be such that $y = q_\infty^\sigma(t)$ (note that $q_\infty^\sigma([t_1, t_2])$ is a closed subset of q_∞^σ as the image of a closed subset of $[0, 1]$ under the projection from $[0, 1]$ into q_∞^σ). We further suppose for the moment that y is not the root of τ as this point needs a special argument. There exists a sequence $(r_k)_k$ of numbers in $[0, 1] \setminus [t_1, t_2]$ such that $q_\infty^\sigma(r_k) \rightarrow q_\infty^\sigma(t)$. Up to extraction, we may suppose that $r_k \rightarrow t' \in [0, 1] \setminus (t_1, t_2)$. Then $q_\infty^\sigma(t') = q_\infty^\sigma(t)$, so that either $t \simeq t'$, in which case $C_\infty^{e_t}(t) = \inf_{r \in [t_1, t]} C_\infty^{e_r}(r)$ or $C_\infty^{e_t}(t) = \inf_{r \in [t, t_2]} C_\infty^{e_r}(r)$ (we are necessarily in the case (2a) of the definition of \simeq from Section 4.3.2, as we supposed that y is not the root of τ), or $t \not\sim t'$ but $t \sim t'$, in which case $\mathfrak{L}(t) = \inf_{[t_1, t]} \mathfrak{L}$ or $\mathfrak{L}(t) = \inf_{[t, t_2]} \mathfrak{L}$, by Proposition 17.

Finally, if the root of τ belongs to $\partial(q_\infty^\sigma([t_1, t_2]))$, it also belongs to the set we claim to be $\partial(q_\infty^\sigma([t_1, t_2]))$ by a simple continuity argument. □

We may now conclude that there is a unique geodesic from ρ^\bullet to x .

Proposition 22. *The path Φ_s is the only geodesic from ρ^\bullet to x .*

Proof. We consider a geodesic $\wp: [0, d_\infty^\sigma(\rho^\bullet, x)] \rightarrow q_\infty^\sigma$ from ρ^\bullet to x . Let $\varepsilon > 0$ be such that $[s - \varepsilon, s + \varepsilon]$ is included in the interval coding τ_a . We take s_i and s_j satisfying the hypotheses of Lemma 20 with this value of ε and we set $r := \sup\{w: \wp(w) \notin q_\infty^\sigma([s_i, s_j])\}$, so that $\wp(r) \in \partial(q_\infty^\sigma([s_i, s_j]))$. Our choice of s_i and s_j ensures that $\wp(r) \notin \pi([\mathcal{M}(a_i), \mathcal{M}(a_j)])$, hence, by Lemma 21,

$$\wp(r) \in \Phi_{s_i}([\mathcal{M}(w'), d_\infty^\sigma(s^\bullet, s_i)]) \cup \Phi_{s_j}([\mathcal{M}(w'), d_\infty^\sigma(s^\bullet, s_j)]) \quad \text{where } w' := \inf\{w: \Phi_{s_i}(w) \neq \Phi_{s_j}(w)\}.$$

As Φ_{s_i} and Φ_{s_j} are the only geodesics from ρ^\bullet to x_i and x_j , this yields that \wp restricted to $[0, r]$ is equal to either Φ_{s_i} or Φ_{s_j} restricted to the same interval. In particular, $\wp(w) = \Phi_{s_i}(w) = \Phi_{s_j}(w)$ for $w \in [0, w']$. See Figure 16. Conducting this reasoning with Φ_s , we obtain that $\wp(w) = \Phi_s(w)$ for $w \in [0, w']$.

As

$$\inf_{[s-\varepsilon, s+\varepsilon]} \mathfrak{L} - \mathfrak{L}(s^\bullet) \leq w' = \inf_{[s_i \wedge s_j, s_i \vee s_j]} \mathfrak{L} - \mathfrak{L}(s^\bullet) \leq \mathfrak{L}(s) - \mathfrak{L}(s^\bullet),$$

we conclude by continuity of \mathfrak{L} that $\wp = \Phi_s$. □

6.2. Geodesics from the base point to all points

Proposition 23. *Let $a \in \mathcal{M}$. The only geodesics from ρ^\bullet to the point $\pi(a)$ are the paths (all distinct) Φ_s , for all s such that $\mathcal{M}(s) = a$.*

Proof. We have already treated the case of leaves in Section 6.1. The other cases are a little more intricate, but use exactly the same arguments, so that we only give a sketch of proof and leave the details to the reader.

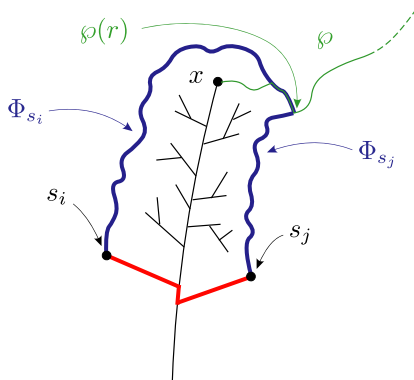


Fig. 16. The geodesic \wp is “trapped” between Φ_{s_i} and Φ_{s_j} .

If $a \in \mathcal{M} \setminus fl$ is a point of order 2, let $s < s'$ be such that $a = \mathcal{M}(s) = \mathcal{M}(s')$. For ε small enough, the set $[s - \varepsilon, s' + \varepsilon]$ is included in the interval coding τ_a . For such an ε , following the ideas of Lemma 20, we can find $s - \varepsilon < s_i < s < s_j < s + \varepsilon$ and $s' - \varepsilon < s_{i'} < s' < s_{j'} < s' + \varepsilon$ such that for all $b \in [[a_i, a_{j'}]] \cup [[a_{i'}, a_j]]$, we have $d_\infty^\sigma(\rho^\bullet, a) < d_\infty^\sigma(\rho^\bullet, b) + d_\infty^\sigma(b, a)$.

We then follow the proof of Proposition 22, replacing $q_\infty^\sigma([s_i, s_j])$ with $q_\infty^\sigma([s_i, s_j] \cup [s_{i'}, s_{j'}])$, which is also a closed subset of q_∞^σ , and whose boundary is the reunion of $\pi([a_i, a_{j'}] \cup [a_{i'}, a_j])$, the ranges of Φ_{s_i} and Φ_{s_j} from the point where they separate up to x_i and x_j , as well as the ranges of $\Phi_{s_{i'}}$ and $\Phi_{s_{j'}}$ from the point where they separate up to $x_{i'}$ and $x_{j'}$. We eventually conclude that the only geodesic from ρ^\bullet to $\pi(a)$ are Φ_s and $\Phi_{s'}$ by the same technique.

All remaining cases are treated in a very similar fashion. □

We immediately deduce the following result by the same method as Le Gall.

Corollary 24 ([24], Corollary 7.7). *Almost surely, for every $\varepsilon > 0$, there exists $\eta \in (0, \varepsilon)$ such that the following holds. Let \wp and \wp' be geodesics from ρ^\bullet to points x and x' such that $d_\infty^\sigma(\rho^\bullet, x) \geq \varepsilon$ and $d_\infty^\sigma(\rho^\bullet, x') \geq \varepsilon$. Then $\wp(w) = \wp'(w)$ for every $w \in [0, \eta]$.*

7. Remaining proofs

7.1. Proof of Theorems 3 and 2

Proof of Theorem 3. Note that, by (4), the indicator in Definition 6 is equal to $\mathbb{1}_{\{\pi(a) \in \partial q_\infty^\sigma\}}$. Assertion (i) is then a direct consequence of Proposition 23. We now turn to the three remaining assertions. There is nothing to show in the case $(g, p) = (0, 0)$ as all the sets in question are empty. If $(g, p) = (0, 1)$, by the Jordan curve Theorem, all the cycles are homotopic to 0, so that we only need to see that $\pi(\mathcal{B}) = \partial q_\infty^\sigma$; this easily follows from (4). We thus suppose from now on that $(g, p) \notin \{(0, 0), (0, 1)\}$.

In this section, it will be convenient to “remove” the root edge dangling from s_∞ . Recall that, as s_∞ is dominant, one extremity of the root e_* has degree 1 and the other one has degree 3. Let $\epsilon_1, \epsilon_2, \bar{\epsilon}_2, \epsilon_3$ be the sequence of successive half-edges in the contour order of f^\bullet such that $\{\epsilon_2, \bar{\epsilon}_2\}$ is the root edge of s_∞ . Note that this implies that $\bar{\epsilon}_1$ is the half-edge following $\bar{\epsilon}_3$ in the contour of the incident face. We define the (nondominant, unrooted) scheme \tilde{s}_∞ as the unrooted map obtained from s_∞ by removing the root edge $\{\epsilon_2, \bar{\epsilon}_2\}$ and merging $\{\epsilon_1, \bar{\epsilon}_1\}$ with $\{\epsilon_3, \bar{\epsilon}_3\}$. We associate with every half-edge $e \in \bar{E}(s_\infty)$ a half-edge $\tilde{e} \in \bar{E}(\tilde{s}_\infty)$ as follows: if $e \in \{\epsilon_1, \epsilon_2, \bar{\epsilon}_2, \epsilon_3\}$ then \tilde{e} is the merged edge, oriented from ϵ_1^- to ϵ_3^+ ; if $e \in \{\bar{\epsilon}_1, \bar{\epsilon}_3\}$ then \tilde{e} is the reverse of the previous half-edge and otherwise $\tilde{e} := e$.

Recall from Section 4.3.2 that e_s is the half-edge of s_∞ “visited” at time s . We also denote by $\bar{\Phi}_s$ the reverse of Φ_s and by $\wp \bullet \wp'$ the concatenation of \wp with \wp' .

Lemma 25. *Let $s \simeq s'$. The path $\Phi_s \bullet \bar{\Phi}_{s'}$ is homotopic to 0 if and only if $\tilde{e}_s = \tilde{e}_{s'}$.*

We postpone the proof of this lemma to the end of this section and conclude. Assertion (ii) is shown as follows. First note that

$$\mathbf{B} = \mathcal{M}(\{s: C_\infty^{e_s}(s) = \underline{C}_\infty^{e_s}(s) \text{ and } e_s \notin \{e_*, \bar{e}_*\}\}).$$

This set splits into $\mathbf{I} \cup \mathbf{B}$ where

$$\begin{aligned} \mathbf{I} &:= \mathcal{M}(\{s: C_\infty^{e_s}(s) = \underline{C}_\infty^{e_s}(s) \text{ and } e_s \in \vec{I}(\mathfrak{s}_\infty) \setminus \{e_*, \bar{e}_*\}\}), \\ \mathbf{B} &:= \mathcal{M}(\{s: C_\infty^{e_s}(s) = \underline{C}_\infty^{e_s}(s) \text{ and } e_s \in \vec{B}(\mathfrak{s}_\infty)\}). \end{aligned} \tag{7}$$

We already know by (4) that $\pi(\mathbf{B}) = \partial q_\infty^\sigma$ so that it will be sufficient to show that $\pi(\mathbf{I}) = \mathcal{N}(\rho^\bullet, q_\infty^\sigma)$. If s lies in the set defining \mathbf{I} , then there exists s' such that $s \simeq s'$ and $\tilde{e}_{s'} \neq \tilde{e}_s$ and as a consequence $q_\infty^\sigma(s) \in \mathcal{N}(\rho^\bullet, q_\infty^\sigma)$, by Lemma 25. Conversely, if $q_\infty^\sigma(s) \in \mathcal{N}(\rho^\bullet, q_\infty^\sigma)$, then Lemma 25 implies the existence of s' such that $s \simeq s'$ and $\tilde{e}_s \neq \tilde{e}_{s'}$; as a result, the pair $\{s, s'\}$ satisfies (2b) or (2c) and thus $\mathcal{M}(s) \in \mathbf{I}$.

Assertion (iii) is also an easy consequence of Lemma 25: a point belongs to the set in question if and only if it is of the form $q_\infty^\sigma(s) = q_\infty^\sigma(s') = q_\infty^\sigma(s'')$ with $\tilde{e}_s, \tilde{e}_{s'}, \tilde{e}_{s''}$ pairwise distinct. These points are in one-to-one correspondence with the vertices of \mathfrak{s}_∞ that are not incident to any holes. Let $N(\mathfrak{s}_\infty)$ denotes their number. If $p = 0$, these are all the vertices of \mathfrak{s}_∞ : an easy application of Euler characteristic formula shows that $N(\mathfrak{s}_\infty) = 4g - 2$. If $p \geq 1$, the situation is a bit more complicated as $N(\mathfrak{s}_\infty)$ actually depends on \mathfrak{s}_∞ . It is a simple functional of \mathfrak{s}_∞ so that its distribution may be computed, as the distribution of \mathfrak{s}_∞ is given by Proposition 15. Moreover, the distribution of \mathfrak{s}_∞ shows that \mathfrak{s}_∞ is equal to every dominant scheme of type $(p, 1)$ with positive probability. This yields that the support of $N(\mathfrak{s}_\infty)$ is merely the range of this functional. Let us now compute this range.

By Euler characteristic formula, $|V(\mathfrak{s}_\infty)| = 4g + 2p - 2$. Because every hole of \mathfrak{s}_∞ is incident to at least one vertex, and because no vertices can be incident to more than one hole, we see that $N(\mathfrak{s}_\infty) \leq 4g + p - 2$. Now, for any given $k \in \{0, 1, \dots, 4g + p - 2\}$, it is not hard to construct a dominant scheme \mathfrak{s} of type $(p, 1)$ such that $N(\mathfrak{s}) = k$ by using the three following operations (see Figure 17). Operation (a) consists in appending an edge with an extremity incident to a degree 1-hole: this increases the number N by one. Operation (b) consists in adding a degree 2-hole on some edge and operation (c) consists in replacing a degree κ -vertex with a hole of degree κ .

If $g \geq 1$ and $4g - 2 \leq k \leq 4g + p - 2$, choose any dominant scheme of type $(0, 1)$. At this stage, it has no holes and its number N is $4g - 2$. Perform $k - (4g - 2)$ times operation (a) and $p - (k - (4g - 2))$ times operation (b) to obtain the desired result. If $g = 0$, start with a dominant scheme of type $(1, 1)$ and perform k times operation (a) and $p - 1 - k$ times operation (b). Finally, if $0 \leq k < 4g - 2$, take a (nondominant) scheme of type $(0, 1)$ having exactly $k + 1$ degree 3-vertices, one degree 1-vertex and one degree $4g - k$ -vertex. Perform operation (c) on the degree $4g - k$ -vertex and perform $p - 1$ times operation (b).

Assertion (iv) directly follows from the observation that $\partial q_\infty^\sigma \cap \mathcal{N}(\rho^\bullet, q_\infty^\sigma)$ is in one-to-one correspondence with the set of vertices of \mathfrak{s}_∞ that are incident to a hole, so that its cardinality is equal to $|V(\mathfrak{s}_\infty)| - N(\mathfrak{s}_\infty)$. \square

Proof of Theorem 2. Here again, the Jordan curve Theorem allows us to directly conclude in the cases $(g, p) \in \{(0, 0), (0, 1)\}$, so that we exclude them. By the previous discussion, $\mathcal{N}(\rho^\bullet, q_\infty^\sigma) = \pi(\mathbf{I})$, where \mathbf{I} was defined by (7).




Fig. 17. Adding a face in order to construct a dominant scheme \mathfrak{s} for which $N(\mathfrak{s}) = k$. (a). Adding one hole while increasing N by 1. (b) and (c). Adding one hole without changing N .

In order to conclude that its Hausdorff dimension is a.s. equal to 2, we apply the method used in [8] to compute the dimension of the boundary. In [8], the method was used to prove that the Hausdorff dimension of the boundary, that is, the image of the boundary edge of the scheme, was 2. In fact, the proof straightforwardly extends to any edge of the scheme and, as \mathbf{I} corresponds to a nonempty collection of scheme edges, the result follows. \square

Proof of Lemma 25. Let $s \simeq s'$ and $a := \mathcal{M}(s) = \mathcal{M}(s')$. If $s = s'$, the result is trivial so that we suppose $s \neq s'$. The condition $\tilde{e}_s = \tilde{e}_{s'}$ is equivalent to saying that one of the sets $\mathcal{M}([s \rightarrow s'])$ or $\mathcal{M}([s' \rightarrow s])$ is a real tree. Indeed, if $\tilde{e}_s = \tilde{e}_{s'}$, then s and s' are instants corresponding to the same forest, so that $\mathcal{M}([s \rightarrow s'])$ or $\mathcal{M}([s' \rightarrow s])$ is a real tree. Conversely, if $\tilde{e}_s \neq \tilde{e}_{s'}$ then a belongs to the backbone $\mathcal{B} \subseteq \mathcal{M}$ and both $\mathcal{M}([s \rightarrow s'])$ and $\mathcal{M}([s' \rightarrow s])$ contain a connected component of $\mathcal{B} \setminus \{a\}$. As a result, neither is a real tree.

Let us first suppose that $\tilde{e}_s = \tilde{e}_{s'}$ and, without loss of generality, let us suppose that $\mathcal{M}([s \rightarrow s'])$ is a real tree. Let $w_0 := \inf\{w: \Phi_s(w) \neq \Phi_{s'}(w)\}$ and

$$\mathcal{R} := \{\Phi_s(w), \Phi_{s'}(w): w_0 \leq w \leq d_\infty^\sigma(s^\bullet, s)\}. \tag{8}$$

Plainly, $\Phi_s \bullet \bar{\Phi}_{s'}$ is homotopic to \mathcal{R} so that it will be sufficient to show that \mathcal{R} is homotopic to 0. We claim that \mathcal{R} is equal to the boundary of $(\pi(\mathcal{M}([s \rightarrow s'])), d_\infty^\sigma)$: indeed, it is easy to check that if $\mathcal{X} \subseteq \mathcal{X}'$ are two surfaces with a boundary, then the boundary of \mathcal{X} is the union of its topological boundary with $(\partial\mathcal{X}') \cap \mathcal{X}$ and we obtain the claim thanks to Lemma 21 (applied to $t_1 := s$ and $t_2 := s'$).

We will show that $(\pi(\mathcal{M}([s \rightarrow s'])), d_\infty^\sigma)$ is homeomorphic to a disk and the result will immediately follow. We can easily construct sequences (i_n) and (i'_n) such that $m_n(i_n) = m_n(i'_n)$, $i_n/2n \rightarrow s$, $i'_n/2n \rightarrow s'$ and the submap of m_n composed of the elements visited between time i_n and time i'_n in the contour order is a tree t_n . The simple discrete geodesic in q_n from v_n^\bullet to i_n is the geodesic made of the arcs between the successive successors of i_n in the construction of Section 3.3. The two simple geodesics toward i_n and i'_n may have a common part. Removing this common part, we are left with a simple loop, which we denote by \mathbf{l}_n . The technique of Section 4.3.1 implies that it is possible to isometrically embed the spaces of the pairs

$$\left(\left(V(q_{n_k}), \left(\frac{9}{8n_k} \right)^{1/4} d_{q_{n_k}} \right), \left(V(t_{n_k}), \left(\frac{9}{8n_k} \right)^{1/4} d_{q_{n_k}} \right) \right)$$

as well as $((q_\infty^\sigma, d_\infty^\sigma), (\pi(\mathcal{M}([s \rightarrow s'])), d_\infty^\sigma))$ into the same metric space, in such a way that the first pair jointly converges toward the latter. It should be understood here that the metric is restricted to the considered sets. We cut the surface corresponding to q_n introduced during Section 4.3.3 along the loop \mathbf{l}_n . Clearly, the component corresponding to t_n is homeomorphic to a disk. Using [8], Proposition 16, we obtain that $(\pi(\mathcal{M}([s \rightarrow s'])), d_\infty^\sigma)$ is homeomorphic to a disk as well. (We need to check that the sequence $(\mathbf{l}_n)_n$ is 0-regular in the sense of [8], Definition 9; this is clear as \mathcal{R} is a simple curve not reduced to a single point.)

Conversely, let us suppose that neither $\mathcal{M}([s \rightarrow s'])$ nor $\mathcal{M}([s' \rightarrow s])$ is a real tree. In particular, a belongs to the backbone $\mathcal{B} \subseteq \mathcal{M}$. Two cases are possible: either $\mathcal{M}([s \rightarrow s']) \cap \mathcal{M}([s' \rightarrow s]) = \{a\}$ or $\mathcal{M}([s \rightarrow s']) \cap \mathcal{M}([s' \rightarrow s]) \neq \{a\}$. Let us first suppose that we are in the latter case. Let us define \mathcal{R} as previously by (8). Recall from Section 5.2 that the set $\pi^{-1}(\mathcal{R})$ is composed only of the point a and of leaves. We claim that this entails that the set $q_\infty^\sigma \setminus \mathcal{R}$ is connected. Indeed, as $a \in \mathcal{B}$, $\mathcal{M} \setminus \pi^{-1}(\mathcal{R})$ may only have one or two connected components. If it has one component, the claim is immediate as $q_\infty^\sigma \setminus \mathcal{R}$ is the image under π of $\mathcal{M} \setminus \pi^{-1}(\mathcal{R})$. If it has two components, then a has to be of order 3 (a corresponds to a node of \mathfrak{s}_∞ because $\mathcal{M}([s \rightarrow s']) \cap \mathcal{M}([s' \rightarrow s]) \neq \{a\}$). Let $s'' \notin \{s, s'\}$ be the third point such that $\mathcal{M}(s'') = a$. Then the two images under π of the two connected components of $\mathcal{M} \setminus \pi^{-1}(\mathcal{R})$ are connected and, for w large enough, $\Phi_{s''}(w)$ belongs to both images; the claim follows. Because \mathcal{R} is not reduced to a single point, The Jordan curve Theorem entails that \mathcal{R} cannot be homotopic to 0.

We now suppose that $\mathcal{M}([s \rightarrow s']) \cap \mathcal{M}([s' \rightarrow s]) = \{a\}$. The same method as above shows that both $(\pi(\mathcal{M}([s \rightarrow s'])), d_\infty^\sigma)$ and $(\pi(\mathcal{M}([s' \rightarrow s])), d_\infty^\sigma)$ are surfaces with a boundary different from the disk (they contain at least a boundary component or a handle of q_∞^σ), whose boundaries contain \mathcal{R} as a component. As a result, \mathcal{R} cannot be homotopic to 0. \square

7.2. Asymptotics for large quadrangulations

The proof of Proposition 5 is a straightforward adaptation of [24], Section 9, so that we leave it to the reader and concentrate on the other proofs, which also use the same ideas but with a little extra work.

Proof of Proposition 6. We argue by contradiction and suppose that there exists $\eta > 0$ such that, for infinitely many n 's, $\mathbb{P}(\exists v \in V(q_n): \text{HMult}_n^{\varepsilon_n}(v_n^\bullet, v) \geq 4) \geq \eta$. From the corresponding increasing sequence of integers, we extract a subsequence $(n_k)_{k \geq 0}$ along which the convergence of Theorem 1 holds. We consider the simple loops made of edges of q_{n_k} that are not homotopic to 0. Let δ_{n_k} denote the smallest diameter of these loops. An easy adaptation of [7], Lemma 21, shows that $\delta_{n_k}/n_k^{1/4}$ is uniformly bounded from below by an a.s. positive random variable δ_0 .

We fix an integer n belonging to $(n_k)_{k \geq 0}$. We claim that for any v and v' , $\text{HMult}_n^{\varepsilon_n}(v, v') \leq \text{Mult}_{n,\delta}^{\varepsilon_n}(v, v')$ as soon as $\delta < \delta_n/(2n^{1/4})$. To see this, it is sufficient to see that every pair $\{\varphi_n, \varphi'_n\}$ of paths such that $\varphi_n \bullet \varphi'_n$ is not homotopic to 0 satisfies $d_{\text{path}}(\varphi_n, \varphi'_n) \geq \delta n^{1/4}$. We argue by contradiction and suppose that there exists a pair $\{\varphi_n, \varphi'_n\}$ satisfying the previous hypothesis but not the conclusion. It defines at least a simple loop \mathbf{I}_n that is not homotopic to 0. Let v_1, \dots, v_j (resp. v'_1, \dots, v'_j) be the vertices of \mathbf{I}_n successively visited by φ_n (resp. φ'_n). For every $i \in \{1, \dots, j\}$, we consider a geodesic γ_i linking v_i to v'_i . These geodesics have length smaller than $\delta n^{1/4}$ by hypothesis. For $i \in \{1, \dots, j-1\}$, the simple loop resulting of the concatenation of γ_i , the edge visited by φ'_n between time i and $i+1$, $\bar{\gamma}_{i+1}$ and the edge visited by φ_n between time $i+1$ and i has diameter smaller than $2\delta n^{1/4}$ so that it is homotopic to 0. As a result, the whole loop \mathbf{I}_n is homotopic to 0, which is a contradiction.

This entails that, for every $\delta > 0$,

$$\mathbb{P}(\exists v \in V(q_n): \text{HMult}_n^{\varepsilon_n}(v_n^\bullet, v) \geq 4) \leq \mathbb{P}(\delta_n \leq 2\delta n^{1/4}) + \mathbb{P}(\exists v \in V(q_n): \text{Mult}_{n,\delta}^{\varepsilon_n}(v_n^\bullet, v) \geq 4).$$

Using Proposition 5, we obtain the following contradiction:

$$\eta \leq \limsup_{k \rightarrow \infty} \mathbb{P}(\exists v \in V(q_{n_k}): \text{HMult}_{n_k}^{\varepsilon_{n_k}}(v_{n_k}^\bullet, v) \geq 4) \leq \inf_{\delta > 0} \mathbb{P}(\delta_0 \leq 2\delta) = 0. \quad \square$$

Before we proceed, let us show our claim following Proposition 7 about $\mathcal{B}(q_n, v_n^\bullet) \setminus \partial V(q_n)$. Let $v \in \mathcal{B}(q_n, v_n^\bullet) \setminus \partial V(q_n)$. We can find i and i' such that $v = q_n(i) = q_n(i')$ and the i th and i' th corners of m_n do not belong to the same forest of the decomposition. It is a simple exercise to check that the simple discrete geodesics toward i and i' satisfy the required property, so that $\text{HMult}_n^0(v_n^\bullet, v) \geq 2$.

Proof of Proposition 7. We follow the lines of reasoning of [24], Section 9. We argue by contradiction and suppose that there exists $\eta > 0$ such that, for infinitely many n 's,

$$\mathbb{P}(\exists v \in V(q_n): d_{q_n}(v, \mathcal{B}(q_n, v_n^\bullet) \setminus \partial V(q_n)) \geq \delta n^{1/4}, \text{HMult}_n^{\varepsilon_n}(v_n^\bullet, v) \geq 2) \geq \eta.$$

From the corresponding increasing sequence of integers, we extract a subsequence $(n_k)_{k \geq 0}$ along which the convergence of Theorem 1 holds, and we let $(q_\infty^\sigma, d_\infty^\sigma)$ be the corresponding Brownian surface. We argue on the event

$$\limsup_{k \rightarrow \infty} \{ \exists v \in V(q_{n_k}): d_{q_{n_k}}(v, \mathcal{B}(q_{n_k}, v_{n_k}^\bullet) \setminus \partial V(q_{n_k})) \geq \delta n_k^{1/4}, \text{HMult}_{n_k}^{\varepsilon_{n_k}}(v_{n_k}^\bullet, v) \geq 2 \},$$

whose probability is greater than $\eta > 0$. For infinitely many k 's, we can find a vertex $v_{n_k} \in V(q_{n_k})$ that satisfies the above requirements; we let i_{n_k} be an integer such that $v_{n_k} = q_{n_k}(i_{n_k})$. By compactness, we can find yet another subsequence along which $i_n/2n \rightarrow s$; we set $x := q_\infty^\sigma(s)$. Until further notice, we only consider integers n belonging to this subsequence. For every $q_\infty^\sigma(s') \in \mathcal{B} \setminus \partial q_\infty^\sigma$, we can easily construct a sequence i'_n such that $q_n(i'_n) \in \mathcal{B}(q_n, v_n^\bullet) \setminus \partial V(q_n)$ and $i'_n/2n \rightarrow s'$ along our subsequence. This implies that $d_{q_n}(q_n(i_n), q_n(i'_n)) \geq \delta n^{1/4}$, so that $d_\infty^\sigma(x, q_\infty^\sigma(s')) \geq (9/8)^{1/4} \delta$ and finally $d_\infty^\sigma(x, \mathcal{B} \setminus \partial q_\infty^\sigma) > 0$. In particular, $x \notin \mathcal{N}(\rho^\bullet, q_\infty^\sigma)$.

As $\text{HMult}_n^{\varepsilon_n}(v_n^\bullet, v_n) \geq 2$, we can find two paths φ_n and φ'_n from v_n^\bullet to v_n with length less than $(1 + \varepsilon_n)d_{q_n}(v_n^\bullet, v_n)$ and such that $\varphi_n \bullet \varphi'_n$ is not homotopic to 0. There exist two integers w_n and w'_n such that the part of φ_n between

times w_n and w'_n concatenated with the part of $\bar{\wp}'_n$ between times w'_n and w_n make up a simple loop \mathbf{I}_n that is not homotopic to 0. We cut the surface of Section 4.3.3 associated with q_n along the loop \mathbf{I}_n : we obtain either one or two surfaces with a boundary. As the number of possibilities for the topology of this or these surfaces is finite, there exist infinitely many n 's for which the topology is always the same. Restricting us to these n 's, we see by the method of the previous section that this or these surfaces converge jointly with the surface associated with q_n to a surface with the same topology and the boundary component corresponding to \mathbf{I}_n converges to a loop \mathbf{I} not homotopic to 0.

Now, let $0 \leq w \leq d_\infty^\sigma(\rho^\bullet, x)$. We denote by $\wp_n(w)$ (resp. $\wp'_n(w)$) the unique vertex of \wp_n (resp. of \wp'_n) lying at distance $\lfloor (8n/9)^{1/4} w \wedge d_{q_n}(v_n^\bullet, v_n) \rfloor$ from v_n^\bullet . For each w , we fix an extraction along which both $\wp_n(w)$ and $\wp'_n(w)$ converge and we denote by $\wp(w), \wp'(w) \in q_\infty^\sigma$ the limiting points. As \wp_n is an approximate geodesic, we see that $d_\infty^\sigma(\wp(w), \wp'(w)) = |w - w'|$, so that \wp is a geodesic from ρ^\bullet to x . Similarly, \wp' is also a geodesic.

Up to further extraction, we may suppose that $w_n/(8n/9)^{1/4} \rightarrow w_\infty$ and $w'_n/(8n/9)^{1/4} \rightarrow w'_\infty$. As $w'_n - w_n$ is half the length of \mathbf{I}_n , which, after scaling, does not converge to 0, we see that $w_\infty < w'_\infty$. On the one hand, for $w_\infty < w < w'_\infty$, we have $\wp_n(w) \in \mathbf{I}_n$ for n sufficiently large, so that $\wp(w) \in \mathbf{I}$. On the other hand, by definition, the vertex of \wp_n lying at distance w_n from v_n^\bullet is the same as the vertex of \wp'_n lying at distance w_n from v_n^\bullet , so that $|\wp_n(w_\infty) - \wp'_n(w_\infty)| \rightarrow 0$ and thus $\wp(w_\infty) = \wp'(w_\infty)$. Moreover, as $\wp(w_\infty)$ is a point of a geodesic that is not an end-point, it is a leaf by the discussion of Section 5.2. As a result, \wp and \wp' restricted to $[0, w_\infty]$ are geodesics from ρ^\bullet to a leaf, so that they coincide. Summing up, we constructed two geodesics \wp and \wp' from ρ^\bullet to x such that $\wp(w) = \wp'(w)$ for $0 \leq w \leq w_\infty$ and $\wp([w_\infty, w'_\infty]) \cup \wp'([w_\infty, w'_\infty]) \subseteq \mathbf{I}$. In order to obtain a contradiction with the fact that $x \notin \mathcal{N}(\rho^\bullet, q_\infty^\sigma)$, we still need to see that $\wp \neq \wp'$. Note that this immediately entails that $w'_\infty = d_\infty^\sigma(\rho^\bullet, x)$ as otherwise, $\wp(w'_\infty)$ would be a leaf reachable from ρ^\bullet by two distinct geodesics. To do this, let us take $y \in \mathbf{I}$. As \mathbf{I}_n converges toward \mathbf{I} , y is the limit of points of the form $\wp_n(r_n)$ or $\wp'_n(r_n)$ with $r_n \geq w_n/(8n/9)^{1/4}$. We first extract a subsequence along which $r_n \rightarrow r$. Considering the extraction along which both $\wp_n(r)$ and $\wp'_n(r)$ converge, we see that $y \in \{\wp(r), \wp'(r)\}$ so that actually $\wp([w_\infty, d_\infty^\sigma(\rho^\bullet, x)]) \cup \wp'([w_\infty, d_\infty^\sigma(\rho^\bullet, x)]) = \mathbf{I}$, which concludes. \square

Proof of Proposition 8. It is sufficient to show that

$$\lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} \mathbb{P}(A_{n,\delta}^{\varepsilon_n}(j)) \leq \mathbb{P}(H \geq j) \leq \lim_{\delta \rightarrow 0} \liminf_{n \rightarrow \infty} \mathbb{P}(A_{n,\delta}^{\varepsilon_n}(j)).$$

Let us focus on the first inequality. We argue by contradiction and suppose that there exist $\eta > 0$ and $\delta > 0$ such that, for infinitely many n 's, $\mathbb{P}(A_{n,\delta}^{\varepsilon_n}(j)) \geq \mathbb{P}(H \geq j) + \eta$. We extract from these n 's a subsequence $(n_k)_{k \geq 0}$ along which the convergence of Theorem 1 holds, and we let $(q_\infty^\sigma, d_\infty^\sigma)$ be the corresponding Brownian surface. We argue on the event $\limsup_{k \rightarrow \infty} A_{n_k,\delta}^{\varepsilon_{n_k}}(j)$ whose probability is strictly greater than $\mathbb{P}(H \geq j)$. Using the same technique as in the proof of Proposition 7, we can construct j points $x_1, \dots, x_j \in q_\infty^\sigma$ lying at distance greater than $(9/8)^{1/4} \delta$ one from another and belonging to the set appearing in (iii). As a result, we obtain that $\limsup_{k \rightarrow \infty} A_{n_k,\delta}^{\varepsilon_{n_k}}(j) \subseteq \{H \geq j\}$, which is a contradiction.

To show the second inequality, we do not extract any subsequences. We argue on the event $\{H \geq j\}$, that is, the event that \tilde{s}_∞ has at least j vertices v_1, \dots, v_j that are not incident to any holes (see Section 7.1). For n large enough, $\varepsilon_n = \varepsilon_\infty$ so that these vertices correspond to vertices of $V(q_n)$. It is not hard to see that these vertices are such that $\text{HMult}_n^0(v_n^\bullet, v_i) = 3$. As a result, the event $A_{n,\delta}^{\varepsilon_n}(j)$ holds, as long as $\min_{i \neq i'} d_{q_n}(v_i, v_{i'}) \geq \delta n^{1/4}$. The ‘‘cactus bound’’ (see [19]) implies that the later condition holds for n large enough, provided that δ is smaller than some a.s. positive random variable Δ , which is an explicit function of ε_∞ and \mathcal{L} . We just showed that

$$\mathbb{P}(\{\delta \leq \Delta\} \cap \{H \geq j\}) \leq \mathbb{P}\left(\liminf_{n \rightarrow \infty} A_{n,\delta}^{\varepsilon_n}(j)\right)$$

and we conclude thanks to Fatou’s Lemma and by taking the limit $\delta \rightarrow 0$. \square

The proof of Proposition 9 is very similar to the proof of Proposition 8; we leave it to the reader.

Often used notation

$\Sigma_{g,p}^\partial$	compact orientable surface of genus g with p boundary components	613
$g \geq 0$	fixed genus in which we work	614
m, q	map, quadrangulation	614
V, E, \vec{E}	vertex set, edge set, half-edge set of a map	614
d_m	graph metric on m	614
e^-, e^+, \bar{e}	origin, end, reverse of the half-edge e	614
e_*	root half-edge	614
$v^\bullet, v^{\bullet\bullet}$	distinguished vertices in q	622
$f^\bullet, f^{\bullet\bullet}$	face of m carrying v^\bullet , face of m carrying $v^{\bullet\bullet}$	623
$(m, [l])$	labeled map	622
$p \geq 0$	number of holes	614
h_1, \dots, h_p	holes in q or in m	614
$n \geq 0$	number of internal faces of q /number of edges not incident to a hole in m	614
σ_i	half boundary length of h_i in q /boundary length of h_i in m	614
$\mathcal{Q}_{n,\sigma}$	set of genus g quadrangulations with p boundary components having n internal faces and such that h_i is of degree $2\sigma^i$, for $1 \leq i \leq p$	614
$\mathcal{M}_{n,\sigma}^1$	set of genus g maps having $n + \sigma $ edges and $p + 1$ faces denoted by $h_1, \dots, h_p, f^\bullet$ such that, for all i , h_i has a simple boundary, is of degree σ^i and is not adjacent to any other h_j , and such that the root is not incident to any hole h_j	622
$\mathfrak{M}_{n,\sigma}^1$	set of pairs $(m, [l])$ where $m \in \mathcal{M}_{n,\sigma}^1$ and l is a suitable labeling function	622
$\mathcal{M}_{n,\sigma}^2$	similar definition as $\mathcal{M}_{n,\sigma}^1$ with one additional face $f^{\bullet\bullet}$	623
$\mathfrak{M}_{n,\sigma}^2$	similar definition as $\mathfrak{M}_{n,\sigma}^1$ with one additional face $f^{\bullet\bullet}$	623
\mathfrak{s}	scheme	625
$\vec{l}(\mathfrak{s}), \vec{b}(\mathfrak{s})$	internal half-edges of \mathfrak{s} , boundary half-edges of \mathfrak{s}	625
$q \in \{1, 2\}$	number of distinguished vertices in the bijection	626
$\mathfrak{S}_{p,q}$	set of schemes of type (p, q)	625
$\mathfrak{S}_{p,q}^*$	set of dominant schemes of type (p, q)	629
(f^e, l^e)	labeled forest corresponding to the half-edge $e \in \mathfrak{s}$	627
$\xi^e \geq 1$	length of the forest f^e	626
$m^e \geq 0$	mass of the forest f^e	626
b^e	bridge corresponding to the half-edge $e \in \mathfrak{s}$	627
$C_{(n)}^e, L_{(n)}^e, B_{(n)}^e$	rescaled coding functions corresponding to the half-edge e	628
$(q_\infty^\sigma, d_\infty^\sigma)$	Brownian surface	631
∂q_∞^σ	boundary of $(q_\infty^\sigma, d_\infty^\sigma)$	634
\mathcal{M}, fl	underlying continuous map, floor of \mathcal{M}	632
$[[a, b]]$	range of the unique injective path linking a to b	632
\sim, \simeq	equivalence relation defining q_∞^σ , equivalence relation defining \mathcal{M}	631
π	canonical projection $\mathcal{M} \rightarrow q_\infty^\sigma$	634
\mathfrak{L}	continuous label function, defined on $[0, 1]$ or \mathcal{M}	633
s^\bullet, ρ^\bullet	argmin \mathfrak{L} , base point $q_\infty^\sigma(s^\bullet)$ of q_∞^σ	633
$[s \rightarrow t]$	$[s, t]$ if $s \leq t$ or $[s, 1] \cup [0, t]$ if $t < s$	637
Φ_s	simple geodesic from ρ^\bullet of index s	637

Acknowledgements

The author warmly thanks Grégory Miermont for interesting discussions leading to the realization of this work. The author also thanks Curtis McMullen for suggesting the terminology of Brownian surfaces, as well as an anonymous referee for tediously reading a previous version of this work and for making numerous improving comments.

References

- [1] C. Abraham. Rescaled bipartite planar maps converge to the Brownian map. Preprint, 2013. Available at [arXiv:1312.5959](https://arxiv.org/abs/1312.5959).
- [2] L. Addario-Berry and M. Albenque. The scaling limit of random simple triangulations and random simple quadrangulations. Preprint, 2013. Available at [arXiv:1306.5227](https://arxiv.org/abs/1306.5227).
- [3] J. Ambjørn and T. G. Budd. Trees and spatial topology change in causal dynamical triangulations. *J. Phys. A* **46** (31) (2013) 315201, 33. [MR3090757](https://arxiv.org/abs/1309.0757)
- [4] E. G. Begle. Regular convergence. *Duke Math. J.* **11** (1944) 441–450. [MR0010964](https://arxiv.org/abs/190010964)
- [5] J. Bertoin, L. Chaumont and J. Pitman. Path transformations of first passage bridges. *Electron. Commun. Probab.* **8** (2003) 155–166 (electronic). [MR2042754](https://arxiv.org/abs/12042754)
- [6] J. Bettinelli. Scaling limits for random quadrangulations of positive genus. *Electron. J. Probab.* **15** (2010) 1594–1644. [MR2735376](https://arxiv.org/abs/120735376)
- [7] J. Bettinelli. The topology of scaling limits of positive genus random quadrangulations. *Ann. Probab.* **40** (5) (2012) 1897–1944. [MR3025705](https://arxiv.org/abs/1203025705)
- [8] J. Bettinelli. Scaling limit of random planar quadrangulations with a boundary. *Ann. Inst. Henri Poincaré Probab. Stat.* **51** (2015) 432–477. [MR3335010](https://arxiv.org/abs/1503335010)
- [9] J. Bettinelli, E. Jacob and G. Miermont. The scaling limit of uniform random plane maps, via the Ambjørn–Budd bijection. *Electron. J. Probab.* **19** (2014) 1–16. [MR3256874](https://arxiv.org/abs/13023256874)
- [10] J. Bettinelli and G. Miermont. Compact Brownian surfaces I. Brownian disks. Preprint. Available at [arXiv:1507.08776](https://arxiv.org/abs/1507.08776).
- [11] J. Bettinelli and G. Miermont. Compact Brownian surfaces II. The general case. In preparation, 2016.
- [12] P. Billingsley. *Convergence of Probability Measures*. Wiley, New York, 1968. [MR0233396](https://arxiv.org/abs/120233396)
- [13] J. Bouttier, P. Di Francesco and E. Guitter. Planar maps as labeled mobiles. *Electron. J. Combin.* **11** (1) (2004), Research Paper 69, 27 pp. (electronic). [MR2097335](https://arxiv.org/abs/12097335)
- [14] D. Burago, Y. Burago and S. Ivanov. *A Course in Metric Geometry. Graduate Studies in Mathematics* **33**. American Mathematical Society, Providence, RI, 2001. [MR1835418](https://arxiv.org/abs/18035418)
- [15] G. Chapuy. The structure of unicellular maps, and a connection between maps of positive genus and planar labeled trees. *Probab. Theory Related Fields* **147** (3–4) (2010) 415–447. [MR2639711](https://arxiv.org/abs/120639711)
- [16] G. Chapuy, M. Marcus and G. Schaeffer. A bijection for rooted maps on orientable surfaces. *SIAM J. Discrete Math.* **23** (3) (2009) 1587–1611. [MR2563085](https://arxiv.org/abs/120563085)
- [17] P. Chassaing and G. Schaeffer. Random planar lattices and integrated superBrownian excursion. *Probab. Theory Related Fields* **128** (2) (2004) 161–212. [MR2031225](https://arxiv.org/abs/12031225)
- [18] R. Cori and B. Vauquelin. Planar maps are well labeled trees. *Canad. J. Math.* **33** (5) (1981) 1023–1042. [MR0638363](https://arxiv.org/abs/120638363)
- [19] N. Curien, J.-F. Le Gall and G. Miermont. The Brownian cactus I. Scaling limits of discrete cactuses. *Ann. Inst. Henri Poincaré Probab. Stat.* **49** (2) (2013) 340–373. [MR3088373](https://arxiv.org/abs/13088373)
- [20] T. Duquesne and J.-F. Le Gall. Random trees, Lévy processes and spatial branching processes. *Astérisque* **281**, vi+147 (2002). [MR1954248](https://arxiv.org/abs/19054248)
- [21] M. Gromov. *Metric Structures for Riemannian and Non-Riemannian Spaces. Progress in Mathematics* **152**. Birkhäuser, Boston, MA, 1999. Based on the 1981 French original [[MR0682063](https://arxiv.org/abs/120682063) (85e:53051)], With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates. [MR1699320](https://arxiv.org/abs/16099320)
- [22] J.-F. Le Gall. *Spatial Branching Processes, Random Snakes and Partial Differential Equations. Lectures in Mathematics ETH Zürich*. Birkhäuser Verlag, Basel, 1999. [MR1714707](https://arxiv.org/abs/1714707)
- [23] J.-F. Le Gall. The topological structure of scaling limits of large planar maps. *Invent. Math.* **169** (3) (2007) 621–670. [MR2336042](https://arxiv.org/abs/120336042)
- [24] J.-F. Le Gall. Geodesics in large planar maps and in the Brownian map. *Acta Math.* **205** (2) (2010) 287–360. [MR2746349](https://arxiv.org/abs/120746349)
- [25] J.-F. Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.* **41** (4) (2013) 2880–2960. [MR3112934](https://arxiv.org/abs/1303112934)
- [26] J.-F. Le Gall and G. Miermont. Scaling limits of random planar maps with large faces. *Ann. Probab.* **39** (1) (2011) 1–69. [MR2778796](https://arxiv.org/abs/120778796)
- [27] J.-F. Le Gall and F. Paulin. Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere. *Geom. Funct. Anal.* **18** (3) (2008) 893–918. [MR2438999](https://arxiv.org/abs/120438999)
- [28] J.-F. Marckert and A. Mokkadem. Limit of normalized quadrangulations: The Brownian map. *Ann. Probab.* **34** (6) (2006) 2144–2202. [MR22294979](https://arxiv.org/abs/1202294979)
- [29] G. Miermont. On the sphericity of scaling limits of random planar quadrangulations. *Electron. Commun. Probab.* **13** (2008) 248–257. [MR2399286](https://arxiv.org/abs/120399286)
- [30] G. Miermont. Tessellations of random maps of arbitrary genus. *Ann. Sci. Éc. Norm. Supér. (4)* **42** (5) (2009) 725–781. [MR2571957](https://arxiv.org/abs/1202571957)
- [31] G. Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.* **210** (2) (2013) 319–401. [MR3070569](https://arxiv.org/abs/13070569)
- [32] S. B. Myers. Connections between differential geometry and topology II. Closed surfaces. *Duke Math. J.* **2** (1) (1936) 95–102. [MR1545908](https://arxiv.org/abs/15045908)
- [33] D. Poulalhon and G. Schaeffer. Optimal coding and sampling of triangulations. *Algorithmica* **46** (3–4) (2006) 505–527. [MR2291967](https://arxiv.org/abs/1202291967)
- [34] D. Revuz and M. Yor. *Continuous Martingales and Brownian Motion*, 3rd edition. *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]* **293**. Springer, Berlin, 1999. [MR1725357](https://arxiv.org/abs/17025357)
- [35] G. Schaeffer. Conjugaison d’arbres et cartes combinatoires aléatoires. Ph.D. thesis, Université de Bordeaux 1, 1998.
- [36] G. T. Whyburn. On sequences and limiting sets. *Fund. Math.* **25** (1935) 408–426.