Elect. Comm. in Probab. **13** (2008), 248–257

ELECTRONIC COMMUNICATIONS in PROBABILITY

ON THE SPHERICITY OF SCALING LIMITS OF RANDOM PLANAR QUADRANGULATIONS

GRÉGORY MIERMONT[1]
Université Pierre et Marie Curie, 175 rue du Chevaleret, F-75013 Paris.
email: `gregory.miermont@math.u-psud.fr`

Submitted January 3, 2008, *accepted in final form* April 7, 2008

AMS 2000 Subject classification: 60C05, 60F05, 60D05
Keywords: Random planar maps, scaling limits, Gromov-Hausdorff convergence, spherical topology

Abstract
We give a new proof of a theorem by Le Gall & Paulin, showing that scaling limits of random planar quadrangulations are homeomorphic to the 2-sphere. The main geometric tool is a reinforcement of the notion of Gromov-Hausdorff convergence, called 1-regular convergence, that preserves topological properties of metric surfaces.

1 Introduction

A planar map is a combinatorial embedding of a connected graph into the 2-dimensional sphere $\mathbb{S}^2$. Formally, it is a class of proper drawings (without edge-crossings) of a connected graph into $\mathbb{S}^2$, where two drawings are considered equivalent if there exists an orientation-preserving homeomorphism of $\mathbb{S}^2$ that corresponds the two drawings. The connected components of the complement of the drawing are called the faces of the map, and their degrees are the number of edges that are incident to each of them.
Random planar maps have drawn much attention in the recent probability literature due to mathematical physics motivations [2] and a powerful encoding of planar maps in terms of labeled trees due to Schaeffer [15, 5]. In turn, scaling limits of labeled trees are well-understood thanks to the works of Aldous, Le Gall and others [1, 8, 9]. Using this line of reasoning, many results have been obtained on the geometric aspects of large random quadrangulations (where faces all have degree 4), and other families of maps. Le Gall [10] showed in particular that scaling limits of random quadrangulations are homeomorphic to the Brownian map introduced by Marckert & Mokkadem [13], and Le Gall & Paulin [11] showed that the topology of the latter is that of the 2-dimensional sphere, hence giving a mathematical content to the claim made by physicists that summing over large random quadrangulations amounts to integrating with respect to some measure over surfaces.
The aim of this note is to give an alternative proof of Le Gall & Paulin's result. We strongly

[1]RESEARCH SUPPORTED BY THE FONDATION DES SCIENCES MATHÉMATIQUES DE PARIS

rely on the results established by Le Gall [10], but use very different methods from [11], where the reasoning uses geodesic laminations and a theorem due to Moore on the topology of quotients of the sphere. We feel that our approach is somewhat more economic, as it only needs certain estimates from [10] and not the technical statements [11, Lemmas 3.1, 3.2] that are necessary to apply Moore's theorem. On the other hand, this is at the cost of checking that quadrangulations are close to being path metric spaces, which is quite intuitive but needs justification (see definitions below). Our main geometric tool is a reinforcement of Hausdorff convergence, called 1-regular convergence and introduced by Whyburn, and which has the property of conserving the topology of surfaces. We will see that random planar quadrangulations converge 1-regularly, therefore entailing that their limits are of the same topological nature. In the case, considered in this paper, of surfaces with the topology of the sphere, the 1-regularity property is equivalent to [11, Corollary 1], stating that there are no small loops separating large random quadrangulations into two large parts. We prove this by a direct argument rather than obtaining it as a consequence of the theorem.

The basic notations are the following. We let $\mathbf{Q}_n$ be the set of rooted[2] quadrangulations of the sphere with n faces, which is a finite set of cardinality $2 \cdot 3^n (2n)!/(n!(n+2)!)$, see [5]. We let $\mathbf{q}_n$ be a random variable picked uniformly in $\mathbf{Q}_n$, and endow the set $V(\mathbf{q}_n)$ of its vertices with the usual graph distance d_n^{gr}, i.e. $d_n^{\mathrm{gr}}(x,y)$ is the length of a minimal (geodesic) chain of edges going from x to y.

We briefly give the crucial definitions on the Gromov-Hausdorff topology, referring the interested reader to [4] for more details. The isometry class $[X,d]$ of the metric space (X,d) is the collection of all metric spaces isometric to (X,d). We let $\mathbb{M}$ be the set of isometry-equivalence classes of *compact* metric spaces. The latter is endowed with the Gromov-Hausdorff distance $\mathtt{d}_{\mathrm{GH}}$, where $\mathtt{d}_{\mathrm{GH}}(\mathcal{X},\mathcal{X}')$ is defined as the least $r>0$ such that there exist a metric space (Z,δ) and subsets $X, X' \subset Z$ such that $[X,\delta]=\mathcal{X}, [X',\delta]=\mathcal{X}'$, and such that the Hausdorff distance between X and X' in (Z,δ) is less than or equal to r. This turns $\mathbb{M}$ into a complete separable metric space, see [6] (this article focuses on compact $\mathbb{R}$-trees, which form a closed subspace of $\mathbb{M}$, but the proofs apply without change to $\mathbb{M}$). Le Gall & Paulin's result states as follows.

Theorem 1 ([11])**.** *A limit in distribution of $[V(\mathbf{q}_n), n^{-1/4} d_n^{\mathrm{gr}}]$ for the Gromov-Hausdorff topology, where $n \to \infty$ along some subsequence, is homeomorphic to the 2-sphere.*

Remarks. • One of the main open questions in the topic of scaling limits of random quadrangulations is to uniquely characterize the limit, i.e. to get rid of the somewhat annoying "along some subsequence" in the previous statement.

• To be perfectly accurate, Le Gall & Paulin showed the same result for uniform $2k$-angulations (maps with degree-$2k$ faces) with n faces. Our methods also apply in this setting (and possibly to more general families of maps), but we will restrict ourselves to the case of quadrangulations for the sake of brevity.

• We plan to provide a generalization of this result to higher genera (i.e. maps on orientable surfaces other than the sphere).

As we are quite strongly relying on Le Gall's results in [10], we will mainly focus on the new aspects of our approach. As a consequence, this paper contains two statements whose proofs will not be detailed (Proposition 2 and Lemma 2), because they are implicit in [10] and follow directly from the arguments therein, and also because their accurate proof would

[2] Which means that one oriented edge of the quadrangulation is distinguished as the root

need a space-consuming introduction to continuum tree and snake formalisms. Taking these statements for granted, the proofs should in a large part be accessible to readers with no particular acquaintance with continuum trees.

2 Gromov-Hausdorff convergence and regularity

We say that a metric space (X, d) is a path metric space if any two points $x, y \in X$ can be joined by a path isometric to a real segment, necessarily with length $d(x, y)$. We let $\mathrm{P}\mathbb{M}$ be the set of isometry classes of compact path metric spaces, and the latter is a closed subspace of $(\mathbb{M}, \mathtt{d}_{\mathrm{GH}})$, see [4, Theorem 7.5.1]. One of the main tools needed in this article is a notion that reinforces the convergence in the metric space $(\mathrm{P}\mathbb{M}, \mathtt{d}_{\mathrm{GH}})$, which was introduced by Whyburn in 1935 and was extensively studied in the years 1940's. Our main source is Begle [3].

Definition 1. *Let $(\mathcal{X}_n, n \geq 1)$ be a sequence of spaces in $\mathrm{P}\mathbb{M}$ converging to a limit $\mathcal{X}$. We say that $\mathcal{X}_n$ converges 1-regularly to $\mathcal{X}$ if for every $\varepsilon > 0$, one can find $\delta, N > 0$ such that for all $n \geq N$, every loop in X_n with diameter $\leq \delta$ is homotopic to 0 in its ε-neighborhood.*

There are a couple of slight differences between this definition and that in [3]. In the latter reference, the setting is that X_n are compact subsets of a common compact space, converging in the Hausdorff sense to a limiting set X. This is not restrictive as Gromov-Hausdorff convergence entails Hausdorff convergence of representative spaces in a common compact space, see for instance [7, Lemma A.1]. It is also assumed in the definition of 1-regular convergence that for every $\varepsilon > 0$, there exists $\delta, N > 0$ such that any two points that lie at distance $\leq \delta$ are in a connected subset of $\mathcal{X}_n$ of diameter $\leq \varepsilon$, but this condition is tautologically satisfied for path metric spaces. Last, the definition in [3] is stated in terms of homology, so our definition in terms of homotopy is in fact stronger.
The following theorem is due to Whyburn, see [3, Theorem 6] and comments before.

Theorem 2. *Let $(\mathcal{X}_n, n \geq 1)$ be a sequence of elements of $\mathrm{P}\mathbb{M}$ that are all homeomorphic to $\mathbb{S}^2$. Assume that $\mathcal{X}_n$ converges to $\mathcal{X}$ for the Gromov-Hausdorff distance, where $\mathcal{X}$ is not reduced to a point, and that the convergence is 1-regular. Then $\mathcal{X}$ is homeomorphic to $\mathbb{S}^2$ as well.*

3 Quadrangulations

Rooted quadrangulations are rooted maps whose faces all have degree 4, and their set is denoted by $\mathbf{Q} := \bigcup_{n \geq 1} \mathbf{Q}_n$ with the notations of the Introduction. For $\mathbf{q} \in \mathbf{Q}$ we let $V(\mathbf{q}), E(\mathbf{q}), F(\mathbf{q})$ be the set of vertices, edges and faces of $\mathbf{q}$, and denote by $d^{\mathrm{gr}}_{\mathbf{q}}$ the graph distance on $V(\mathbf{q})$.

3.1 A metric surface representation

One of the issues that must be addressed in order to apply Theorem 2 is that the metric space $[V(\mathbf{q}), d^{\mathrm{gr}}_{\mathbf{q}}]$ is not a surface, rather, it is a finite metric space. We take care of this by constructing a particular graphical representative of $\mathbf{q}$ which is a path metric space whose restriction to the vertices of the graph is isometric to $(V(\mathbf{q}), d^{\mathrm{gr}}_{\mathbf{q}})$.
Let $(X_f, d_f), f \in F(\mathbf{q})$ be copies of the emptied unit cube "with bottom removed"

$$X_f = [0,1]^3 \setminus \big((0,1)^2 \times [0,1)\big) \ ,$$

endowed with the intrinsic metric d_f inherited from the Euclidean metric (the distance between two points of X_f is the minimal Euclidean length **of a path in** X_f). Obviously each (X_f, d_f) is a path metric space homeomorphic to a closed disk of $\mathbb{R}^2$. The boundary of each face $f \in F(\mathbf{q})$, when explored turning counterclockwise, is made of four oriented edges e_1, e_2, e_3, e_4, where the labeling is arbitrary among the 4 possible labelings preserving the cyclic order. Then define

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

In these notations, we keep the subscript f to differentiate points of different spaces X_f. In this way, for every oriented edge e of the map $\mathbf{q}$, we have defined a path c_e which goes along one of the four edges of the square $\partial X_f = ([0,1]^2 \setminus (0,1)^2) \times \{0\}$, where f is the face located to the left of e.

We define a relation $\sim$ on the disjoint union $\amalg_{f\in F(\mathbf{q})} X_f$, as the coarsest equivalence relation such that for every oriented edge e of $\mathbf{q}$, and every $t \in [0,1]$, we have $c_e(t) \sim c_{\overline{e}}(1-t)$, where $\overline{e}$ is e with reversed orientation. By identifying points of the same class, we glue the oriented sides of the squares ∂X_f pairwise, in a way that is consistent with the map structure. More precisely, the topological quotient $\mathcal{S}_\mathbf{q} := \amalg_{f\in F(\mathbf{q})} X_f/\sim$ is a surface which has a 2-dimensional cell complex structure, whose 1-skeleton $\mathcal{E}_\mathbf{q} := \amalg_{f\in F(\mathbf{q})} \partial X_f/\sim$ is a graph drawing of the map $\mathbf{q}$, with faces (2-cells) $X_f \setminus \partial X_f$. In particular, $\mathcal{S}_\mathbf{q}$ is homeomorphic to $\mathbb{S}^2$ by [14, Lemma 3.1.4]. With an oriented edge e of $\mathbf{q}$, one associates an edge of the graph drawing $\mathcal{E}_\mathbf{q}$ in $\mathcal{S}_\mathbf{q}$, more simply called an edge of $\mathcal{S}_\mathbf{q}$, made of the equivalence classes of points in $c_e([0,1])$ (or $c_{\overline{e}}([0,1])$). We also let $\mathcal{V}_\mathbf{q}$ be the 0-skeleton of this complex, i.e. the vertices of the graph — these are the equivalent classes of the corners of the squares ∂X_f. We call them the vertices of $\mathcal{S}_\mathbf{q}$ for simplicity.

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

Proposition 1. *The space $(\mathcal{S}_\mathbf{q}, d_\mathbf{q})$ is a path metric space homeomorphic to $\mathbb{S}^2$. Moreover, the restriction of $\mathcal{S}_\mathbf{q}$ to the set $\mathcal{V}_\mathbf{q}$ is isometric to $(V(\mathbf{q}), d^{\mathrm{gr}}_\mathbf{q})$, and any geodesic path in $\mathcal{S}_\mathbf{q}$ between two elements of $\mathcal{V}_\mathbf{q}$ is a concatenation of edges of $\mathcal{S}_\mathbf{q}$. Last,*

$$\mathtt{d}_{\mathrm{GH}}([V(\mathbf{q}), d^{\mathrm{gr}}_\mathbf{q}], [\mathcal{S}_\mathbf{q}, d_\mathbf{q}]) \le 3\,.$$

Proof. What we first have to check is that $d_\mathbf{q}$ is a true metric on $\mathcal{S}_\mathbf{q}$, i.e. that it separates points. To see this, we use the fact [4, Theorem 3.1.27] that $D_\mathbf{q}$ admits the constructive expression:

$$D_\mathbf{q}(a,b) = \inf\left\{ \sum_{i=0}^{n} d(x_i, y_i) : n \ge 0, x_0 = a, y_n = b, y_i \sim x_{i+1} \right\},$$

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

It remains to treat the case $a \in \partial X_f, b \in \partial X_{f'}$ for some f, f'. The crucial observation is that a shortest path in X_f between two points of ∂X_f is entirely contained in ∂X_f. It is then a simple exercise to check that for a, b in distinct classes, the distance $D_{\mathbf{q}}(a, b)$ will be larger than the length of some fixed non-trivial path with values in $\mathcal{E}_{\mathbf{q}}$. More precisely, if (the classes of) a, b belong to the same edge of $\mathcal{S}_{\mathbf{q}}$, then we can find representatives a', b' in the same X_f and we will have $D_{\mathbf{q}}(a, b) \geq d_f(a', b')$. If the class of a is not a vertex of $\mathcal{S}_{\mathbf{q}}$ but that of b is, then $D_{\mathbf{q}}(a, b)$ is at least equal to the distance of $a \in X_f$ to the closest corner of the square ∂X_f. Finally, if the (distinct) equivalence classes of a, b are both vertices, then $D_{\mathbf{q}}(a, b) \geq 1$. One deduces that $d_{\mathbf{q}}$ is a true distance on $\mathcal{S}_{\mathbf{q}}$, which makes it a path metric space by [4, Corollary 3.1.24]. Since $\mathcal{S}_{\mathbf{q}}$ is a compact topological space, the metric space $(\mathcal{S}_{\mathbf{q}}, d_{\mathbf{q}})$ is homeomorphic to $\mathbb{S}^2$ by [4, Exercise 3.1.14].
From the last paragraph's observations, a shortest path between vertices of $\mathcal{S}_{\mathbf{q}}$ takes all its values in $\mathcal{E}_{\mathbf{q}}$. Since an edge of $\mathcal{S}_{\mathbf{q}}$ is easily checked to have length 1 for the distance $d_{\mathbf{q}}$, such a shortest path will have same length as a geodesic path for the (combinatorial) graph distance between the two vertices. Hence $(\mathcal{V}_{\mathbf{q}}, d_{\mathbf{q}})$ is indeed isometric to $(V(\mathbf{q}), d^{\mathrm{gr}}_{\mathbf{q}})$. The last statement follows immediately from this and the fact that $\operatorname{diam}(X_f, d_f) \leq 3$, entailing that $\mathcal{V}_{\mathbf{q}}$ is 3-dense in $(\mathcal{S}_{\mathbf{q}}, d_{\mathbf{q}})$, i.e. its 3-neighborhood in $(\mathcal{S}_{\mathbf{q}}, d_{\mathbf{q}})$ equals $\mathcal{S}_{\mathbf{q}}$. □

3.2 Tree encoding of quadrangulations

We briefly introduce the second main ingredient, the Schaeffer bijection (see e.g. [10] for details). Let $\mathbf{T}_n$ be the set of pairs $(\mathbf{t}, \mathbf{l})$ where $\mathbf{t}$ is a rooted plane tree with n edges, and $\mathbf{l}$ is a function from the set of vertices of $\mathbf{t}$ to $\{1, 2, \ldots\}$, such that $|\mathbf{l}(x) - \mathbf{l}(y)| \leq 1$ if x and y are neighbors, and $\mathbf{l}(x^0) = 1$, where x^0 is the root vertex of $\mathbf{t}$. Then the set $\mathbf{Q}_n$ is in one-to-one correspondence with $\mathbf{T}_n$. More precisely, this correspondence is such that given a graph representation of $\mathbf{q} \in \mathbf{Q}_n$ on a surface, the corresponding $(\mathbf{t}, \mathbf{l}) \in \mathbf{T}_n$ can be realized as a graph drawn on the same surface, whose vertices are $V(\mathbf{t}) = V(\mathbf{q}) \setminus \{x^*\}$, where x^* is the origin of the root edge of $\mathbf{q}$, and $\mathbf{l}$ is the restriction to $V(\mathbf{t})$ of the function $\mathbf{l}(x) = d^{\mathrm{gr}}_{\mathbf{q}}(x, x^*), x \in V(\mathbf{q})$. Moreover, the edges of $\mathbf{t}$ and $\mathbf{q}$ only intersect at vertices. The root vertex x^0 of $\mathbf{t}$ is the tip of the root edge of $\mathbf{q}$, so it lies at $d^{\mathrm{gr}}_{\mathbf{q}}$-distance 1 from x^*.
Let $\xi_{\mathbf{t}}(0) = x^0$. Recursively, given $\{\xi_{\mathbf{t}}(0), \ldots, \xi_{\mathbf{t}}(i)\}$, let $\xi_{\mathbf{t}}(i+1)$ be the first[3] child of $\xi_{\mathbf{t}}(i)$ not in $\{\xi_{\mathbf{t}}(0), \ldots, \xi_{\mathbf{t}}(i)\}$ if there is any, or the parent of $\xi_{\mathbf{t}}(i)$ otherwise. This procedure stops at $i = 2n$, when we are back to the root and have explored all vertices of the tree. We let $C_i = d^{\mathrm{gr}}_{\mathbf{t}}(\xi_{\mathbf{t}}(i), \xi_{\mathbf{t}}(0))$, and $L_i = \mathbf{l}(\xi_{\mathbf{t}}(i))$. Both C and L are extended by linear interpolation between integer times into continuous functions, still called C, L, with duration $2n$. The process C, called the contour process, is the usual Dyck path encoding of the rooted tree $\mathbf{t}$, and the pair (C, L) determines $(\mathbf{t}, \mathbf{l})$ completely.

3.3 Estimates on the lengths of geodesics

Our last ingredient is a slight rewriting of the estimates of Le Gall [10] on distances in quadrangulations in terms of encoding processes. Precisely, let C^n, L^n be the contour and label process of a uniform random element $\mathbf{t}_n$ of $\mathbf{T}_n$, and let $\mathbf{q}_n$ be the quadrangulation that is the image of this element by Schaeffer's bijection. In particular, $\mathbf{q}_n$ is a uniform random element of $\mathbf{Q}_n$. Also, recall that a graphical representation $\mathcal{T}_n$ of $\mathbf{t}_n$ can be drawn on the space $\mathcal{S}_{\mathbf{q}_n}$ of Sect. 3.1, in such a way that the vertices of $\mathcal{T}_n$ are $\mathcal{V}_{\mathbf{q}_n} \setminus \{x_n^*\}$, where x_n^* is the root vertex of $\mathbf{q}_n$, considered as an element of $\mathcal{V}_{\mathbf{q}_n}$, and $\mathcal{T}_n$ intersects $\mathcal{E}_{\mathbf{q}_n}$ only at vertices. For simplicity we

[3] For the natural order inherited from the planar structure of $\mathbf{t}$

let $V_n = V(\mathbf{q}_n)$, $d_n^{\mathrm{gr}} = d_{\mathbf{q}_n}^{\mathrm{gr}}$, $\mathcal{S}_n = \mathcal{S}_{\mathbf{q}_n}$, $d_n = d_{\mathbf{q}_n}$ and $\xi_n(\cdot) = \xi_{\mathbf{t}_n}(\cdot)$. We will also assimilate the vertex $\xi_n(i)$ of $\mathbf{t}_n$ with a point of $\mathcal{T}_n \subset \mathcal{S}_n$, which is a vertex of $\mathcal{S}_n$.
The main result of [9] says that the convergence in distribution in $\mathcal{C}([0,1],\mathbb{R})^2$ holds:

$$\left(\left(\frac{1}{\sqrt{2n}}C^n_{2nt}\right)_{0\le t\le 1}, \left(\left(\frac{9}{8n}\right)^{1/4} L^n_{2nt}\right)_{0\le t\le 1}\right) \xrightarrow[n\to\infty]{(d)} (\overline{\mathbb{e}}, \overline{Z}), \tag{1}$$

where $(\overline{\mathbb{e}}, \overline{Z})$ is the Brownian snake conditioned to be positive introduced by Le Gall & Weill [12]. Moreover, it is shown in [10] that the laws of $[V_n, n^{-1/4}d_n^{\mathrm{gr}}]$ form a relatively compact family in the set of probability measures on $\mathbb{M}$ endowed with the weak topology. Since V_n is 3-dense in $\mathcal{S}_n$, the same holds for $[\mathcal{S}_n, n^{-1/4}d_n]$. We argue as in [10], and assume by Skorokhod's theorem that the trees $\mathbf{t}_n$ (hence also the quadrangulations $\mathbf{q}_n$) are defined on the same probability space, on which we have, almost-surely

- $[\mathcal{S}_n, n^{-1/4}d_n] \to [S, d]$, some random limiting space in $\mathbb{PM}$, along some (fixed, nonrandom) subsequence $n_k \to \infty$, and
- the convergence (1) holds a.s. along this subsequence.

From this point on, we will always assume that n is taken along this subsequence. In particular, we have that $\operatorname{diam} S = \lim_n n^{-1/4}\operatorname{diam}\mathcal{S}_n \ge \lim_n \sup n^{-1/4}L^n = (8/9)^{1/4}\sup \overline{Z} > 0$ a.s., so S is not reduced to a point and Theorem 2 may be applied if we check that the convergence is 1-regular. We are going to rely on proposition 4.2 of [10], which can be rephrased as follows.

Proposition 2. *The following property is true with probability* 1. *Let i_n, j_n be integers such that $i_n/2n \to s, j_n/2n \to t$ in $[0,1]$, where $s < t$ satisfy*

$$\overline{\mathbb{e}}_s = \inf_{s\le u\le t} \overline{\mathbb{e}}_u < \overline{\mathbb{e}}_t\,.$$

Then it holds that

$$\liminf_{n\to\infty} n^{-1/4} d_n(\xi_n(i_n), \xi_n(j_n)) > 0\,.$$

In [10], this proposition was a first step in the proof of the fact that a limit in distribution of (V_n, d_n^{gr}) can be expressed as a quotient of the continuum tree with contour function $\overline{\mathbb{e}}$. Proposition 2 says that two points of the latter such that one is an ancestor of the other are not identified. Le Gall completed this study by exactly characterizing which are the points that are identified.

4 Proof of Theorem 1

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

Proof. Assume that with positive probability, along some (random) subsequence, there exist simple loops γ_n made of edges of $\mathcal{S}_n$, with diameters $o(n^{1/4})$ as $n \to \infty$, such that the two Jordan domains bounded by γ_n are of diameters $\ge n^{1/4}\varepsilon$, where $\varepsilon > 0$ is some fixed constant. Reasoning on this event, let l_n be the minimal label on γ_n, i.e. $l_n = d_n(x_n^*, \gamma_n)$. Then the

labels of vertices that are in a connected component D_n of $\mathcal{S}_n \setminus \gamma_n$ not containing x_n^* are all larger than l_n, since a geodesic from x_n^* to any such vertex must pass through γ_n.
The intuitive idea of the proof is the following. Starting from the root of the tree $\mathcal{T}_n$, follow a simple path in $\mathcal{T}_n$ that enters in D_n at some stage. If all such paths remained in D_n after entering, then all the descendents of the entrance vertices would have labels larger than l_n, which is close in the scale $n^{1/4}$ to the label of the entrance vertex. The property of admitting a subtree with labels all larger than that of its root is of of zero probability under the limiting Brownian snake measure, see [10, Lemma 2.2] and Lemma 2 below. Thus, some of these paths must go out of D_n after entering, but they can do it only by passing through γ_n again, which entails that strict ancestors in $\mathcal{T}_n$ will be at distance $o(n^{1/4})$. This is prohibited by Proposition 2. This reasoning is summed up in Figure 1, which gathers some of the notations to come.

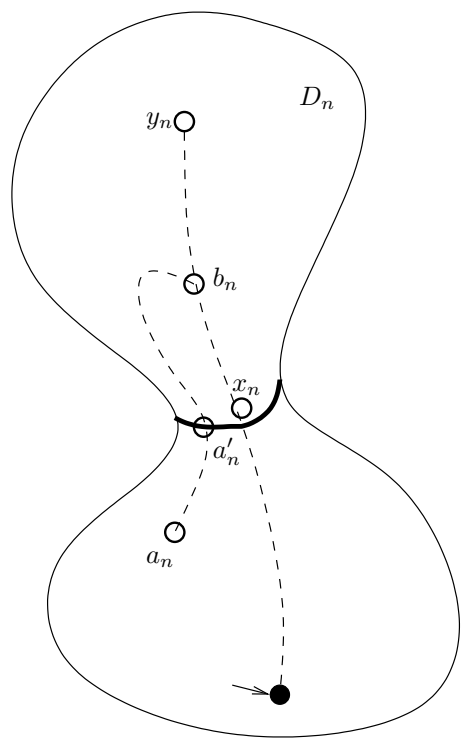

Figure 1: Illustration of the proof. The surface $\mathcal{S}_n$ is depicted as a sphere with a bottleneck circled by γ_n (thick line). The root edge of the quadrangulation is drawn at the bottom, and the tree $\mathcal{T}_n$ originates from its tip. In dashed lines are represented the two branches of $\mathcal{T}_n$ that are useful in the proof: one enters the component D_n, and the other goes out after entering, identifying strict ancestors in the limit

We proceed to the rigorous proof. Let y_n be a vertex in D_n at maximal distance from γ_n. Since every point in $\mathcal{S}_n$ is at distance at most 3 from some vertex, this shows $\operatorname{diam}(D_n) \leq \operatorname{diam}(\gamma_n) + 2d_n(y_n, \gamma_n) + 6$. Thus $d_n(y_n, \gamma_n) \geq (\varepsilon/2 + o(1))n^{1/4}$. Since every path in $\mathcal{S}_n$ from y_n to x_n^* has to cross γ_n, we obtain that $\mathbf{l}(y_n) = d_n(y_n, x_n^*) \geq d_n(y_n, \gamma_n) + d_n(\gamma_n, x_n^*)$, so that

$$\mathbf{l}(y_n) \geq l_n + n^{1/4}\varepsilon/4 \qquad \text{for large enough } n\,. \tag{2}$$

We let x_n be the highest ancestor of y_n in $\mathcal{T}_n$ lying at d_n-distance ≤ 1 from γ_n, so that $\mathbf{l}(x_n) = l_n + o(n^{1/4})$ and all vertices in the ancestral line from x_n to y_n are in D_n. Then, we can find integers $i_n < j_n$ such that $x_n = \xi_n(i_n), y_n = \xi_n(j_n)$, so that $\mathbf{l}(x_n) = L^n_{i_n}, \mathbf{l}(y_n) = L^n_{j_n}$. Up to further extraction, we may and will assume that

$$(9/8n)^{1/4} l_n \to l\,, \qquad i_n/2n \to s\,, \qquad j_n/2n \to t\,, \qquad s \leq t\,. \tag{3}$$

Since $x_n \prec y_n$, we have $C^n_{i_n} = \inf_{i_n \leq r \leq j_n} C^n_r$, as a basic property of Dyck path encodings of trees. Using (1), we have $\overline{\mathbb{e}}_s = \inf_{u\in[s,t]} \overline{\mathbb{e}}_u$. More precisely, by (1), (2) and (3), we have $\overline{Z}_s = l$ and $\overline{Z}_t \geq l + (9/8)^{1/4}\varepsilon/4$, which implies $s < t$, and $\overline{\mathbb{e}}_s < \overline{\mathbb{e}}_t$ (outside of a set of probability zero). In terms of the continuum tree admitting $\overline{\mathbb{e}}$ as contour process, this amounts to the fact that s, t encode two vertices such that the first is an ancestor of the second, and that are not the same because the snake $\overline{Z}$ takes distinct values at these points. Now, we need the following statement:

Lemma 2. *Assume that $s > 0$. Outside of a set of probability 0, there exist $\eta > 0$ and vertices a_n, b_n of $\mathcal{T}_n$ such that $x_n \prec b_n \prec y_n$ and $b_n \prec a_n$, where $\prec$ denotes "is an ancestor of", with labels satisfying*

$$\mathbf{l}(b_n) > \mathbf{l}(x_n) + \eta n^{1/4}\,,$$

and

$$\mathbf{l}(a_n) \leq \mathbf{l}(x_n) - \eta n^{1/4}\,.$$

In words, there exist subtrees of $\mathcal{T}_n$ branching on a vertex b_n of the ancestral line from x_n to y_n that attain labels that are significantly smaller (in the scale $n^{1/4}$) than $\mathbf{l}(y_n)$, but such that $\mathbf{l}(b_n)$ is significantly larger than $\mathbf{l}(x_n)$.

Proof (sketch). Standard properties of Dyck paths encodings imply that x_n is an ancestor in $\mathcal{T}_n$ of $\xi_n(i)$ for every integer $i \in [i_n, j_n]$. Let $s' = \sup\{u \in [s,t] : \overline{\mathbb{e}}_{s'} = \overline{\mathbb{e}}_s\} \leq t$. Then $\inf_{u\in[s,s']} \overline{\mathbb{e}}_u = \overline{\mathbb{e}}_s = \overline{\mathbb{e}}'_s$, which yields $\overline{Z}_s = \overline{Z}_{s'}$, by a basic property of the label process $\overline{Z}$ (see [12] for instance). Thus $s' < t$ and

$$\overline{\mathbb{e}}_{s'} < \overline{\mathbb{e}}_u \qquad \text{for every} \quad u \in (s', t]\,. \tag{4}$$

By [10, Lemma 2.2], this implies that for some $\alpha > 0$, and with full probability,

$$\inf_{u\in[s',s'+\alpha]} \overline{Z}_u < \overline{Z}_{s'} = l\,. \tag{5}$$

We take integers $i'_n \in [i_n, j_n]$ so that $i'_n/2n \to s'$ and $C^n_i \geq C^n_{i'_n}$ for all $i \in [i'_n, j_n]$. Then $x'_n := \xi_n(i'_n)$ satisfies $x_n \prec x'_n \prec y_n$.
Finally, if $m_n \in [i'_n, j_n]$ is such that $x'_n \prec \xi_n(m_n) \prec y_n$, then $\mathbf{l}(\xi_n(m_n)) = L^n_{m_n} \geq l_n$, because $\xi_n(m_n)$ lies in D_n. If moreover $m_n/2n \to u \in [s', t]$, we obtain that $(9/8n)^{1/4} L^n_{m_n} \to \overline{Z}_u > \overline{Z}_{s'} = l$. Thus $\overline{Z}_u \geq \overline{Z}_{s'} = l$ for every $u \in [s', t]$ such that $\overline{\mathbb{e}}_u = \inf_{v\in[u,t]} \overline{\mathbb{e}}_v$, a continuous counterpart for the fact that all labels on the ancestral line from x'_n to y_n are larger than $\mathbf{l}(x'_n) + o(n^{1/4}) = (8n/9)^{1/4}(l + o(1))$.
At this point, the conclusion is obtained from this last fact and (4), (5) by reasoning along the exact same lines as in the proof of [10, Proposition 4.2 pp. 649-650]. The inequalities (15) and (17) therein give the existence of $\eta > 0, a_n, b_n$ with $x'_n \prec b_n \prec y_n$ and $b_n \prec a_n$, satisfying the stated inequalities. This implies the result since $x_n \prec x'_n$. □

Now back to the proof of Lemma 1. Take k_n, r_n with $i_n < k_n < r_n < j_n$, such that $\xi_n(k_n) = a_n, \xi_n(r_n) = b_n$. Because of the property of the label of a_n, it does not lie in D_n, however, its ancestor b_n does because it is on the ancestral path from x_n to y_n. Hence some ancestor of a_n must belong to γ_n. Let a'_n be the highest in the tree, and take $k'_n \in (k_n, r_n)$ with $\xi_n(k'_n) = a'_n$. Since x_n is at distance at most 1 from γ_n, we obtain that $d_n(a'_n, x_n) = o(n^{1/4})$. However, if $k'_n/2n \to v, r_n/2n \to u$, taking again an extraction if necessary, then we have $\overline{\mathbb{e}}_s < \overline{\mathbb{e}}_u \leq \overline{\mathbb{e}}_v$ because of the ancestral relations $C^n_{i_n} \leq C^n_{r_n} \leq C^n_{k'_n}$, and the fact

$$\overline{Z}_s = \lim_{n\to\infty} (9/8n)^{1/4} L^n_{i_n} = l < l + (9/8)^{1/4}\eta \leq \lim_{n\to\infty} (9/8n)^{1/4} L^n_{r_n} = \overline{Z}_u \,.$$

Now the statements $d_n(a'_n, x_n) = o(n^{1/4})$ and $\overline{\mathbb{e}}_s < \overline{\mathbb{e}}_v$ together hold with zero probability by Proposition 2, a contradiction.

It remains to rule out the possibility that $s = 0$, i.e. that γ_n lies at distance $o(n^{1/4})$ from x^*_n. To see that this is not possible, argue as in the beginning of the proof and take x_n, y_n respectively in the two disjoint connected components of $\mathcal{S}_n \setminus \gamma_n$, and with labels $\mathbf{l}(x_n) \wedge \mathbf{l}(y_n) \geq n^{1/4}\varepsilon/4$. By symmetry, assume that $x_n = \xi_n(i_n)$ and $y_n = \xi_n(j_n)$ with $i_n < j_n$. Now take the least integer $k_n \in [i_n, j_n]$ such that $\xi_n(k)$ belongs to γ_n. Such a k has to exist because any path from x_n to y_n in $\mathcal{S}_n$ must pass through γ_n. Then $L^n_{k_n} = \mathbf{l}(\xi_n(k_n)) = o(n^{1/4})$. Up to extraction, assume $i_n/2n \to s, k_n/2n \to u, j_n/2n \to t$. Then $\overline{Z}_u = 0 < \overline{Z}_s \wedge \overline{Z}_t$, so that $s < u < t$, and this contradicts the fact that $\overline{Z}$ is a.s. strictly positive on $(0, 1)$, which is a consequence of [12, Proposition 2.5]. □

We claim that Lemma 1 is enough to obtain 1-regularity of the convergence, and hence to conclude by Theorem 2 that the limit (S, d) is a sphere. First choose $\varepsilon < \operatorname{diam} S/3$ to avoid trivialities. Let γ_n be a loop in $\mathcal{S}_n$ with diameter $\leq n^{1/4}\delta$. Consider the union of the closures of faces of $\mathcal{S}_n$ that are visited by γ_n. The boundary of this union is made of pairwise disjoint simple loops $\mathcal{L}$ made of edges of $\mathcal{S}_n$. If x, y are elements in this family of faces, and since a face of $\mathcal{S}_n$ has diameter less than 3, there exist points x', y' of γ_n at distance at most 3 from x, y respectively, so that the diameters of the loops in $\mathcal{L}$ all are $\leq n^{1/4}\delta + 6$.

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

Acknowledgments. Thanks to Pierre Pansu for an interesting discussion on regular convergence, and to an anonymous referee, whose careful reading allowed to improve the exposition and fix some inaccuracies of an earlier version of the present paper.

References

[1] D. J. Aldous. The continuum random tree. III. *Ann. Probab.*, 21(1):248–289, 1993. MR1207226

[2] J. Ambjørn, B. Durhuus, and T. Jonsson. *Quantum geometry. A statistical field theory approach.* Cambridge Monographs on Mathematical Physics. Cambridge University Press, Cambridge, 1997. MR1465433

[3] E. G. Begle. Regular convergence. *Duke Math. J.*, 11:441–450, 1944. MR0010964

[4] D. Burago, Y. Burago, and S. Ivanov. *A course in metric geometry*, volume 33 of *Graduate Studies in Mathematics.* American Mathematical Society, Providence, RI, 2001. MR1835418

[5] P. Chassaing and G. Schaeffer. Random planar lattices and integrated superBrownian excursion. *Probab. Theory Related Fields*, 128(2):161–212, 2004. MR2031225

[6] S. N. Evans, J. Pitman, and A. Winter. Rayleigh processes, real trees, and root growth with re-grafting. *Probab. Theory Related Fields*, 134(1):81–126, 2006. MR2221786

[7] A. Greven, P. Pfaffelhuber, and A. Winter. Convergence in distribution of random metric measure spaces (Λ-coalescent measure trees). 2006.

[8] J.-F. Le Gall. *Spatial branching processes, random snakes and partial differential equations.* Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, 1999. MR1714707

[9] J.-F. Le Gall. A conditional limit theorem for tree-indexed random walk. *Stochastic Process. Appl.*, 116(4):539–567, 2006. MR2205115

[10] J.-F. Le Gall. The topological structure of scaling limits of large planar maps. *Invent. Math.*, 169(3):621–670, 2007. MR2336042

[11] J.-F. Le Gall and F. Paulin. Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere. *Geom. Funct. Anal.*, 2008. To appear.

[12] J.-F. Le Gall and M. Weill. Conditioned Brownian trees. *Ann. Inst. H. Poincaré Probab. Statist.*, 42(4):455–489, 2006. MR2242956

[13] J.-F. Marckert and A. Mokkadem. Limit of normalized random quadrangulations: the Brownian map. *Ann. Probab.*, 34(6):2144–2202, 2006. MR2294979

[14] B. Mohar and C. Thomassen. *Graphs on surfaces.* Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, 2001. MR1844449

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