

The Brownian plane

Nicolas Curien and Jean-François Le Gall

Ecole normale supérieure and Université Paris-Sud

Abstract

We introduce and study the random non-compact metric space called the Brownian plane, which is obtained as the scaling limit of the uniform infinite planar quadrangulation. Alternatively, the Brownian plane is identified as the Gromov-Hausdorff tangent cone in distribution of the Brownian map at its root vertex, and it also arises as the scaling limit of uniformly distributed (finite) planar quadrangulations with n faces when the scaling factor tends to 0 less fast than $n^{-1/4}$. We discuss various properties of the Brownian plane. In particular, we prove that the Brownian plane is homeomorphic to the plane, and we get detailed information about geodesic rays to infinity.

1 Introduction

A planar map is a finite connected (multi)graph drawn on the two-dimensional sphere and viewed up to orientation-preserving homeomorphisms of the sphere. The faces of a planar map are the connected components of the complement of edges, and the degree of a face counts the number of edges in its boundary, with the special convention that if both sides of an edge are incident to the same face then this edge is counted twice in the degree of this face. Important special cases of planar maps are triangulations, where all faces have degree three, and quadrangulations, where all faces have degree four. It is usual to deal with rooted planar maps, meaning that there is a distinguished edge, which is also oriented and whose origin is called the root vertex.

Much recent work has been devoted to understanding the asymptotic properties of large planar maps chosen uniformly at random in a particular class, e.g. the class of all triangulations, or of all quadrangulations, with a fixed number n of faces tending to infinity. There are two basic kinds of limit theorems giving information about large random planar maps.

First, local limit theorems consider for every fixed integer $r \geq 1$ the combinatorial ball of radius r in the planar map (this is the new planar map obtained by keeping only those faces whose boundary contains at least one vertex whose graph distance from the root vertex is smaller than r), and show that this ball converges in distribution as the size of the map tends to infinity towards the corresponding ball in a random infinite planar lattice. Such local limits were studied first in the case of triangulations by Angel and Schramm [3, 4] and the limiting object is called the uniform infinite planar triangulation (UIPT). The analogous result for quadrangulations was obtained later by Krikun [12] (see also Chassaing and Durhuus [7] and Ménard [20]), leading to the uniform infinite planar quadrangulation (UIPQ).

Secondly, scaling limits consist in looking at the vertex set of a planar map with n faces as a metric space for the graph distance rescaled by the factor $n^{-1/4}$, and studying the convergence of this metric space when n tends to infinity, in the sense of the Gromov-Hausdorff distance familiar to geometers. The factor $n^{-1/4}$ is chosen so that the diameter of the rescaled planar

map remains bounded in probability: It was first noticed by Chassaing and Schaeffer [8] that the diameter of a random quadrangulation with n faces is of order $n^{1/4}$, and a similar result holds for much more general random planar maps, including triangulations. The existence of a scaling limit for (uniformly distributed) random quadrangulations was obtained recently in the papers [15, 21], leading to a limiting random compact metric space called the Brownian map. In [15], it is also proved that the Brownian map is the universal scaling limit of more general random planar maps including triangulations.

Our main goal in the present work is to provide a connection between the preceding limit theorems, by introducing a random (non-compact) metric space called the Brownian plane, which can be viewed either as the scaling limit of the UIPQ or as the Gromov-Hausdorff tangent cone in distribution of the Brownian map at its root. The Brownian plane can also be obtained as the limit of rescaled random quadrangulations with n faces if the graph distance is multiplied by a factor ε_n such that $\varepsilon_n \rightarrow 0$ and $\varepsilon_n n^{1/4} \rightarrow \infty$ as $n \rightarrow \infty$.

Let us give a precise definition of the Brownian plane before stating our main results. We consider two independent three-dimensional Bessel processes R and R' started from 0 (see e.g. [22] for basic facts about Bessel processes). We then define a process $X = (X_t)_{t \in \mathbb{R}}$ indexed by the real line, by setting

$$X_t = \begin{cases} R_t & \text{if } t \geq 0, \\ R'_{-t} & \text{if } t \leq 0. \end{cases}$$

Then, for every $s, t \in \mathbb{R}$, we set

$$\overline{st} = \begin{cases} [s \wedge t, s \vee t] & \text{if } st \geq 0, \\ (-\infty, s \wedge t] \cup [s \vee t, \infty) & \text{if } st < 0, \end{cases}$$

and

$$m_X(s, t) = \inf_{r \in \overline{st}} X_r.$$

We define a random pseudo-distance on \mathbb{R} by

$$d_X(s, t) = X_s + X_t - 2m_X(s, t)$$

and put $s \sim_X t$ if $d_X(s, t) = 0$. The quotient space $\mathcal{T}_\infty = \mathbb{R} / \sim_X$ equipped with d_X is a (non-compact) random real tree, which is sometimes called the infinite Brownian tree. This tree corresponds to Process 2 in Aldous [1]. It can be verified that \mathcal{T}_∞ is the tangent cone in distribution of Aldous' CRT at its root vertex, in a sense that will be explained below (see Theorem 11 in [1] for a closely related result). We write $p_\infty : \mathbb{R} \rightarrow \mathcal{T}_\infty$ for the canonical projection and set $\rho_\infty = p_\infty(0)$, which plays the role of the root of \mathcal{T}_∞ .

We next consider Brownian motion indexed by \mathcal{T}_∞ . Formally, we consider a real-valued process $(Z_t)_{t \in \mathbb{R}}$ such that, conditionally given the process X , Z is a centered Gaussian process with covariance

$$E[Z_s Z_t | X] = m_X(s, t)$$

so that we have $Z_0 = 0$ and $E[(Z_s - Z_t)^2 | X] = d_X(s, t)$. It is not hard to verify that the process Z has a modification with continuous paths, and we consider this modification from now on. Then a.s. we have $Z_s = Z_t$ for every $s, t \in \mathbb{R}$ such that $d_X(s, t) = 0$, and thus we may (and sometimes will) view Z as indexed by \mathcal{T}_∞ .

For every $s, t \in \mathbb{R}$, we set

$$D_\infty^\circ(s, t) = Z_s + Z_t - 2 \min_{r \in [s \wedge t, s \vee t]} Z_r.$$

We extend the definition of D_∞° to $\mathcal{T}_\infty \times \mathcal{T}_\infty$ by setting for $a, b \in \mathcal{T}_\infty$,

$$D_\infty^\circ(a, b) = \min\{D_\infty^\circ(s, t) : s, t \in \mathbb{R}, p_\infty(s) = a, p_\infty(t) = b\}.$$

Finally, we set, for every $a, b \in \mathcal{T}_\infty$,

$$D_\infty(a, b) = \inf_{a_0=a, a_1, \dots, a_p=b} \sum_{i=1}^p D_\infty^\circ(a_{i-1}, a_i)$$

where the infimum is over all choices of the integer $p \geq 1$ and of the finite sequence a_0, a_1, \dots, a_p in \mathcal{T}_∞ such that $a_0 = a$ and $a_p = b$. It is not hard to verify that D_∞ is a pseudo-distance on \mathcal{T}_∞ . We put $a \approx b$ if $D_\infty(a, b) = 0$ (as will be show in Proposition 11, this is equivalent to the property $D_\infty^\circ(a, b) = 0$). The Brownian plane is the quotient space $\mathcal{P} = \mathcal{T}_\infty / \approx$, which is equipped with the metric induced by D_∞ and with the distinguished point which is the equivalence class of ρ_∞ . We simply write ρ_∞ for this equivalence class, and use the notation $\mathcal{P} = (\mathcal{P}, D_\infty, \rho_\infty)$ for the Brownian plane viewed as a pointed metric space (recall that a metric space (E, d) is said to be pointed if there is a distinguished point $\alpha \in E$).

In order to state our first result, we introduce the notation $(\mathfrak{m}_\infty, D^*)$ for the Brownian map, as defined in the introduction of [15] or in [21] for instance. Recall that \mathfrak{m}_∞ is obtained as a quotient set of Aldous' CRT [1, 2], and that in this construction, the Brownian map comes with a distinguished point, which is the equivalence class of the root ρ of the CRT. We will abuse notation and also write ρ for the equivalence class of ρ in \mathfrak{m}_∞ . From our perspective, it will be important to view the Brownian map as a triplet $\mathfrak{m}_\infty := (\mathfrak{m}_\infty, D^*, \rho)$, which is a (random) pointed compact metric space. Note that, in a sense that can be made precise, ρ is uniformly distributed over \mathfrak{m}_∞ .

For every $\varepsilon > 0$, let $B_\varepsilon(\mathcal{P})$ be the closed ball of radius ε centered at ρ_∞ in \mathcal{P} , and let $B_\varepsilon(\mathfrak{m}_\infty)$ be the closed ball of radius ε centered at ρ in \mathfrak{m}_∞ . Each of these balls is pointed at its center and thus viewed as a pointed metric space.

Theorem 1. *For every $\delta > 0$, we can find $\varepsilon > 0$ and construct on the same probability space copies of the Brownian plane \mathcal{P} and of the Brownian map \mathfrak{m}_∞ , in such a way that the balls $B_\varepsilon(\mathcal{P})$ and $B_\varepsilon(\mathfrak{m}_\infty)$ are isometric with probability at least $1 - \delta$. Furthermore, we have*

$$(\mathfrak{m}_\infty, \lambda D^*, \rho) \xrightarrow[\lambda \rightarrow \infty]{(d)} (\mathcal{P}, D_\infty, \rho_\infty) \quad (1)$$

in distribution for the local Gromov-Hausdorff topology.

Let us briefly discuss the local Gromov-Hausdorff topology (for more details see Chapter 8.1 in [6] and Section 2 below). First recall that a metric space (E, d) is called a length space if, for every $a, b \in E$, the distance $d(a, b)$ coincides with the infimum of the lengths of continuous curves connecting a to b . We also say that (E, d) is boundedly compact if all closed balls are compact. Then a sequence (E_n, d_n, α_n) of pointed boundedly compact length spaces is said to converge to (E, d, α) in the local Gromov-Hausdorff topology if, for every $r > 0$, the closed ball of radius r centered at α_n in E_n converges to the closed ball of radius r centered at α in E , for the usual Gromov-Hausdorff distance between pointed compact metric spaces. The space of all pointed boundedly compact length spaces (modulo isometries) can be equipped with a metric which is compatible with the preceding notion of convergence, and is then separable and complete for this metric. The convergence in Theorem 1 is just the usual convergence in distribution for random variables with values in a Polish space.

If (E, d) is a boundedly compact length space and $\alpha \in E$, the Gromov-Hausdorff tangent cone of (E, d) at α , if it exists, is the limit in the local Gromov-Hausdorff topology of the rescaled spaces $(E, \lambda d, \alpha)$ when λ tends to infinity (see Section 8.2 in [6]). The convergence (1) can thus be interpreted by saying that the Brownian plane is the Gromov-Hausdorff tangent cone in distribution of the Brownian map at its root.

Remark. It follows from the convergence (1) that the Brownian plane is scale invariant, meaning that, for every $\lambda > 0$, $(\mathcal{P}, \lambda D_\infty, \rho_\infty)$ has the same distribution as $(\mathcal{P}, D_\infty, \rho_\infty)$. This can also be verified directly from the definition, using the similar property for the infinite Brownian tree.

In order to state our second theorem, let us write Q_∞ for the uniform infinite planar quadrangulation and $V(Q_\infty)$ for the vertex set of Q_∞ . The root vertex of Q_∞ is denoted by $\rho_{(\infty)}$. Similarly, for every integer $n \geq 1$, Q_n stands for a uniformly distributed rooted quadrangulation with n faces, $V(Q_n)$ is the vertex set of Q_n and $\rho_{(n)}$ is the root vertex of Q_n . Finally d_{gr} denotes the graph distance on either $V(Q_n)$ or $V(Q_\infty)$.

Theorem 2. *We have*

$$(V(Q_\infty), \lambda d_{\text{gr}}, \rho_{(\infty)}) \xrightarrow[\lambda \rightarrow 0]{(d)} (\mathcal{P}, D_\infty, \rho_\infty), \quad (2)$$

in distribution for the local Gromov-Hausdorff topology. Furthermore, let $(k_n)_{n \geq 0}$ be a sequence of non-negative real numbers such that $k_n \rightarrow \infty$ and $k_n = o(n^{1/4})$ as n tends to infinity. Then,

$$(V(Q_n), k_n^{-1} d_{\text{gr}}, \rho_{(n)}) \xrightarrow[n \rightarrow \infty]{(d)} (\mathcal{P}, D_\infty, \rho_\infty) \quad (3)$$

in distribution for the local Gromov-Hausdorff topology.

The reader may have noticed that neither $(V(Q_\infty), d_{\text{gr}})$ nor $(V(Q_n), d_{\text{gr}})$ is a length space, so that the discussion following Theorem 1 does not seem to apply to the convergences (2) and (3). However it is very easy to approximate $(V(Q_\infty), d_{\text{gr}})$ (resp. $(V(Q_n), d_{\text{gr}})$) by a pointed boundedly compact length space, in such a way that the balls centered at the distinguished point in $(V(Q_\infty), d_{\text{gr}})$ and in this approximating space are within Gromov-Hausdorff distance 1 (and a similar result holds for the balls in $(V(Q_n), d_{\text{gr}})$). The convergences (2) and (3) make sense, and will indeed be proved for these approximating length spaces.




Figure 1: Illustration of Theorems 1 and 2.

The convergences in Theorems 1 and 2, as well as the convergence of rescaled finite quadrangulations to the Brownian map and the local convergence to the UIPQ, are summarized in the diagram of Figure 1.

The proofs of both Theorems 1 and 2 are based on coupling methods. The coupling argument already appears in the first statement of Theorem 1, which is in fact much stronger than the local Gromov-Hausdorff convergence (1). In the discrete setting, we prove in Proposition 9 that we can couple Q_∞ with Q_n so that the balls of radius $o(n^{1/4})$ are the same with high probability. This allows us to partially answer a question of Krikun on separating cycles in the UIPQ (see Corollary 10).

The coupling result given in Theorem 1 makes it possible to derive several important properties of the Brownian plane from known results about the Brownian map. In Section 5, we prove that the Brownian plane has dimension 4 and is homeomorphic to the plane. We also study geodesic rays in the Brownian plane (a geodesic ray is a semi-infinite path converging to infinity whose restriction to any finite interval is a geodesic). We prove in particular that any two geodesic rays coalesce in finite time, a property that is reminiscent of the results derived in [14] for the Brownian map and in [9] for the UIPQ. As a corollary of our study of geodesics, we also show that the “labels” Z_a can be interpreted as measuring relative distances in \mathcal{P} from the point at infinity, in the sense that, a.s. for every $a, b \in \mathcal{T}_\infty$,

$$Z_a - Z_b = \lim_{x \rightarrow \infty} (D_\infty(a, x) - D_\infty(b, x))$$

where the limit holds when x tends to the point at infinity in the Alexandroff compactification of \mathcal{P} . This should be compared to Theorem 2 in [9]. At this point it is worth mentioning that, even if the Brownian map and the Brownian plane share many important properties, the Brownian plane enjoys the additional scale invariance property, which may play a significant role in future investigations of these (still mysterious) random objects.

Finally, we conjecture that the convergence (2) in Theorem 2 still holds if the UIPQ Q_∞ is replaced by the UIPT, and more generally that this convergence can be extended to many infinite random lattices that are obtained as local limits of random planar maps.

The paper is organized as follows. Section 2 gathers some preliminaries about the local Gromov-Hausdorff topology and the convergence of rescaled finite quadrangulations towards the Brownian map. Section 3 is devoted to the proof of Theorem 1. Theorem 2 is proved in Section 4. Finally Section 5 studies properties of the Brownian plane.

Acknowledgment. We are indebted to Grégoire Miermont for a number of very useful discussions.

2 Preliminaries

2.1 Gromov-Hausdorff convergence

In this subsection, we recall some basic notions from metric geometry. For more details we refer to [6].

A pointed metric space $\mathbf{E} = (E, d, \alpha)$ is a metric space given with a distinguished point $\alpha \in E$. We will generally use bold letters $\mathbf{E}, \mathbf{Q}, \dots$ for pointed spaces. For every $r \geq 0$, we denote the closed ball of radius r centered at α by $B_r(\mathbf{E})$. We view $B_r(\mathbf{E})$ as a pointed metric space (where obviously α is the distinguished point).

Recall that if X and Y are two compact subsets of a metric space (E, d) , the Hausdorff distance between X and Y is

$$d_{\mathbb{H}}^E(X, Y) := \inf\{\varepsilon > 0 : X \subset Y^\varepsilon \text{ and } Y \subset X^\varepsilon\},$$

where $X^\varepsilon := \{x \in E : d(x, X) \leq \varepsilon\}$ denotes the ε -neighborhood of X . If $\mathbf{E} = (E, d, \alpha)$ and $\mathbf{E}' = (E', d', \alpha')$ are two pointed compact metric spaces, the Gromov-Hausdorff distance between \mathbf{E} and \mathbf{E}' is

$$d_{\text{GH}}(\mathbf{E}, \mathbf{E}') := \inf\{d_{\mathbb{H}}^F(\phi(E), \phi'(E')) \vee \delta(\phi(\alpha), \phi'(\alpha'))\},$$

where the infimum is taken over all choices of the metric space (F, δ) and of the isometric embeddings $\phi : E \rightarrow F$ and $\phi' : E' \rightarrow F$ of E and E' into F . The Gromov-Hausdorff distance is indeed a metric on the space \mathbb{K} of all isometry classes of pointed compact metric spaces (see [6, Section 7.4] for more details), and the space $(\mathbb{K}, d_{\text{GH}})$ is a Polish space, that is, a separable and complete metric space.

In order to extend the Gromov-Hausdorff convergence to the non-compact case, we will restrict our attention to boundedly compact length spaces. This restriction is not really necessary (see [6, Section 8.1]) but it will avoid certain technicalities, which are not relevant to our purposes.

If (E, d) is a metric space and $\gamma : [0, T] \rightarrow E$ is a continuous path in E , the length of γ is defined by

$$L(\gamma) := \sup_{0=t_0 < \dots < t_k=T} \sum_{i=0}^{k-1} d(\gamma(t_i), \gamma(t_{i+1})),$$

where the supremum is over all choices of the subdivision $0 = t_0 < t_1 < \dots < t_k = T$ of $[0, T]$. The space (E, d) is called a length space if, for every $x, y \in E$, the distance $d(x, y)$ coincides with the infimum of the lengths of continuous paths connecting x to y . We say that (E, d) is boundedly compact if the closed balls in (E, d) are compact. For a length space, this is equivalent to saying that (E, d) is complete and locally compact [6, Proposition 2.5.22]. In a boundedly compact length space, a path with minimal length (or geodesic) exists between any pair of points.

Let $(\mathbf{E}_n)_{n \geq 0}$ and \mathbf{E} be pointed boundedly compact length spaces. We say that \mathbf{E}_n converges towards \mathbf{E} in the local Gromov-Hausdorff sense if, for every $r \geq 0$, we have

$$d_{\text{GH}}(B_r(\mathbf{E}_n), B_r(\mathbf{E})) \xrightarrow{n \rightarrow \infty} 0.$$

This notion of convergence is compatible with the distance

$$d_{\text{LGH}}(\mathbf{E}_1, \mathbf{E}_2) = \sum_{k=1}^{\infty} 2^{-k} (d_{\text{GH}}(B_k(\mathbf{E}_1), B_k(\mathbf{E}_2)) \wedge 1).$$

As observed in [10, Proposition 3.3], the set \mathbb{K}_{bcl} of all isometry classes of pointed boundedly compact length spaces endowed with d_{LGH} is a Polish space.

We will use the following easy consequence of the preceding considerations. If $(\mathbf{E}_n)_{n \geq 0}$ is now a sequence of random variables with values in \mathbb{K}_{bcl} , this sequence converges in distribution to $\mathbf{E} \in \mathbb{K}_{\text{bcl}}$ if and only if, for every $r \geq 0$, $B_r(\mathbf{E}_n)$ converges in distribution to $B_r(\mathbf{E})$ in \mathbb{K} . Furthermore, we may restrict our attention to integer values of r .

Notation. If $\mathbf{E} = (E, d, \alpha)$ is a pointed metric space and λ is a positive real number, the notation $\lambda \cdot \mathbf{E}$ stands for the rescaled metric space $(E, \lambda \cdot d, \alpha)$. In particular for any $\lambda, \eta > 0$ we have $\lambda \cdot B_\eta(\mathbf{E}) = B_{\lambda\eta}(\lambda \cdot \mathbf{E})$.

2.2 Convergence to the Brownian map

In this subsection, we recall the definition of the Brownian map, and briefly discuss the convergence of rescaled finite quadrangulations to this random metric space. The construction of the Brownian map is very similar to that of the Brownian plane given in the introduction, but the role of the process X is now played by a Brownian excursion.

We fix $T > 0$. For reasons that will become clear later, it is convenient to index the processes that we will define by the parameter $\lambda = T^{1/4}$. Let $(e_t^\lambda)_{0 \leq t \leq T}$ be a Brownian excursion with duration $T = \lambda^4$. For the purposes of this subsection, it would be sufficient to take $T = \lambda = 1$, but later we will deal with scaled versions of the Brownian map for which arbitrary values of T will be useful.

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

$$d_{e^\lambda}(s, t) = e_s^\lambda + e_t^\lambda - 2 \min_{s \wedge t \leq r \leq s \vee t} e_r^\lambda$$

and we set $s \sim_{e^\lambda} t$ if and only if $d_{e^\lambda}(s, t) = 0$. The tree coded by e^λ is the quotient space $\mathcal{T}_{e^\lambda} := [0, T] / \sim_{e^\lambda}$, which is equipped with the distance induced by d_{e^λ} . Note that \mathcal{T}_{e^λ} is a scaled version of the CRT: We refer to Section 3 of [16] for basic facts about the CRT and the coding of real trees by functions. We write $p_{e^\lambda} : [0, T] \rightarrow \mathcal{T}_{e^\lambda}$ for the canonical projection, and we set $\rho_\lambda = p_{e^\lambda}(0) = p_{e^\lambda}(T)$. To simplify notation, we write $\rho = \rho_1$.

Conditionally given e^λ , let $Z^\lambda = (Z_s^\lambda)_{0 \leq s \leq T}$ be the centered Gaussian process with covariance

$$E[Z_s^\lambda Z_t^\lambda | e^\lambda] = \min_{r \in [s \wedge t, s \vee t]} e_r^\lambda.$$

The process $(Z_s^\lambda)_{0 \leq s \leq T}$ has a continuous modification, which we consider from now on. Then, a.s. for every $s \in [0, T]$, Z_s^λ only depends on $p_{e^\lambda}(s)$ and so, for every $a \in \mathcal{T}_{e^\lambda}$, we may set $Z_a^\lambda = Z_s^\lambda$, where s is an arbitrary element of $[0, T]$ such that $p_{e^\lambda}(s) = a$.

For every $s, t \in [0, T]$ such that $s \leq t$, we set

$$D_\lambda^\circ(s, t) = D_\lambda^\circ(t, s) = Z_s^\lambda + Z_t^\lambda - 2 \max \left(\min_{r \in [s, t]} Z_r^\lambda, \min_{r \in [t, T] \cup [0, s]} Z_r^\lambda \right).$$

Furthermore, we set for $a, b \in \mathcal{T}_{e^\lambda}$,

$$D_\lambda^\circ(a, b) = \min \{ D_\lambda^\circ(s, t) : s, t \in [0, T], p_{e^\lambda}(s) = a, p_{e^\lambda}(t) = b \},$$

and

$$D_\lambda^*(a, b) = \inf_{a_0=a, a_1, \dots, a_p=b} \sum_{i=1}^p D_\lambda^\circ(a_{i-1}, a_i),$$

where the infimum is over all choices of the integer $p \geq 1$ and of the finite sequence a_0, a_1, \dots, a_p in \mathcal{T}_{e^λ} such that $a_0 = a$ and $a_p = b$.

Since clearly $D_\lambda^\circ(a, b) \geq |Z_a^\lambda - Z_b^\lambda|$, it is immediate that we have also

$$D_\lambda^*(a, b) \geq |Z_a^\lambda - Z_b^\lambda|$$

for every $a, b \in \mathcal{T}_{e^\lambda}$. More precisely, the pseudo-distance D_λ^* satisfies the so-called ‘‘cactus bound’’

$$D_\lambda^*(a, b) \geq Z_a^\lambda + Z_b^\lambda - 2 \min_{c \in [a, b]} Z_c^\lambda \quad (4)$$

where $\llbracket a, b \rrbracket$ denotes the geodesic segment between a and b in $\mathcal{T}_{\mathbf{e}^\lambda}$. To see this, first note that the cactus bound holds for $D_\lambda^\circ(a, b)$ instead of $D_\lambda^*(a, b)$: This is an immediate consequence of the definition of D_λ° and the fact that, if $p_{\mathbf{e}^\lambda}(s) = a$ and $p_{\mathbf{e}^\lambda}(t) = b$, the set $p_{\mathbf{e}^\lambda}([s \wedge t, s \vee t])$ must contain $\llbracket a, b \rrbracket$. Then, given a finite sequence $a_0 = a, a_1, \dots, a_p = b$ we observe that $\llbracket a, b \rrbracket$ is contained in the union of the segments $\llbracket a_{i-1}, a_i \rrbracket$ for $1 \leq i \leq p$. Hence, there must exist an index $j \in \{1, \dots, p\}$ such that

$$\min_{c \in \llbracket a, b \rrbracket} Z_c^\lambda \geq \min_{c \in \llbracket a_{j-1}, a_j \rrbracket} Z_c^\lambda$$

and it follows that

$$\begin{aligned} \sum_{i=1}^p D_\lambda^\circ(a_{i-1}, a_i) &\geq \sum_{i=1}^{j-1} |Z_{a_i}^\lambda - Z_{a_{i-1}}^\lambda| + (Z_{a_{j-1}}^\lambda + Z_{a_j}^\lambda - 2 \min_{c \in \llbracket a_{j-1}, a_j \rrbracket} Z_c^\lambda) + \sum_{i=j+1}^p |Z_{a_i}^\lambda - Z_{a_{i-1}}^\lambda| \\ &\geq Z_a^\lambda + Z_b^\lambda - 2 \min_{c \in \llbracket a, b \rrbracket} Z_c^\lambda. \end{aligned}$$

Finally, we put $a \approx_\lambda b$ if and only if $D_\lambda^*(a, b) = 0$. Specializing now to the case $T = 1$, the Brownian map \mathbf{m}_∞ is defined as the quotient space $\mathcal{T}_{\mathbf{e}^1} / \approx_1$, which is equipped with the distance induced by D_1^* . We view \mathbf{m}_∞ as a pointed metric space with distinguished point ρ (here and later we abuse notation and identify ρ with its equivalence class in $\mathcal{T}_{\mathbf{e}^1} / \approx_1$), and we write $\mathbf{m}_\infty = (\mathbf{m}_\infty, D_1^*, \rho)$ for this pointed metric space. With the notation introduced at the end of the preceding subsection, we have, by a simple scaling argument,

$$\lambda \cdot \mathbf{m}_\infty \stackrel{(d)}{=} (\mathcal{T}_{\mathbf{e}^\lambda} / \approx_\lambda, D_\lambda^*, \rho_\lambda), \quad (5)$$

with the same abuse of notation for ρ_λ . This identity in distribution explains why we considered an arbitrary value of $T > 0$ (and not only the case $T = 1$) in the preceding discussion.

Let us recall the convergence of rescaled random quadrangulations towards the Brownian map [15, 21]. For every integer $n \geq 1$, let Q_n and $V(Q_n)$ be as in the introduction. We view $V(Q_n)$ as a metric space for the graph distance d_{gr} , which is pointed at the root vertex $\rho_{(n)}$ of Q_n . Then, we have

$$(V(Q_n), n^{-1/4} d_{\text{gr}}, \rho_{(n)}) \xrightarrow[n \rightarrow \infty]{(d)} \left(\mathbf{m}_\infty, \left(\frac{8}{9}\right)^{1/4} D_1^*, \rho \right) \quad (6)$$

in distribution in \mathbb{K} .

3 The Brownian plane as a tangent cone of the Brownian map

In this section, we prove Theorem 1. To do so, we establish a coupling result (Proposition 4) from which Theorem 1 easily follows. This lemma shows that we can couple the realizations of the Brownian plane and of the Brownian map in such a way that the balls of small radius centered at the distinguished point coincide with high probability.

3.1 Absolute continuity properties of excursion measures

The branching structures of the (scaled) CRT $\mathcal{T}_{\mathbf{e}^\lambda}$ and of the infinite Brownian tree \mathcal{T}_∞ are encoded respectively in a Brownian excursion \mathbf{e}^λ of duration $T = \lambda^4$ and in a pair (R, R') of independent three-dimensional Bessel processes. In the proof of the forthcoming coupling result, it will be important to have information about the absolute continuity properties of the law of a

pair of processes corresponding respectively to the initial and the final part of \mathbf{e}^λ , with respect to the law of (R, R') .

We fix $T > 0$ and set $\lambda = T^{1/4}$ as previously. We also consider two reals $\alpha, \beta > 0$ such that $\alpha + \beta < T$. We will then consider the pair $((\mathbf{e}_t^\lambda)_{0 \leq t \leq \alpha}, (\mathbf{e}_{T-t}^\lambda)_{0 \leq t \leq \beta})$, which is viewed as a random element of the space $C([0, \alpha], \mathbb{R}_+) \times C([0, \beta], \mathbb{R}_+)$. The generic element of the latter space will be denoted by (ω, ω') .

For $t > 0$ and $y \in \mathbb{R}$, we let

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

be the usual Brownian transition density. The transition density of Brownian motion killed upon hitting 0 is

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

for $x, y > 0$. We also set

$$q_t(x) = \frac{x}{t} p_t(x)$$

for $t > 0$ and $x > 0$. We recall that the transition density of the three-dimensional Bessel process is given by

$$\begin{cases} p_t^{(3)}(x, y) = \frac{y}{x} p_t^*(x, y) & \text{if } t > 0, x > 0, y \geq 0, \\ p_t^{(3)}(0, y) = 2y q_t(y) & \text{if } t > 0, y \geq 0. \end{cases}$$

Proposition 3. *The law of the triplet*

$$\left((\mathbf{e}_t^\lambda)_{0 \leq t \leq \alpha}, (\mathbf{e}_{T-t}^\lambda)_{0 \leq t \leq \beta}, \min_{\alpha \leq t \leq T-\beta} \mathbf{e}_t^\lambda \right)$$

is absolutely continuous with respect to the law of

$$\left((R_t)_{0 \leq t \leq \alpha}, (R'_t)_{0 \leq t \leq \beta}, \inf_{t \geq \alpha} R_t \wedge \inf_{t \geq \beta} R'_t \right),$$

with density given by the function $\Delta_{T, \alpha, \beta}(\omega(\alpha), \omega'(\beta), z)$, where, for every $x, y, z > 0$,

$$\Delta_{T, \alpha, \beta}(x, y, z) = \varphi_{T, \alpha, \beta}(x, y) \psi_{T-(\alpha+\beta)}(x, y, z)$$

with

$$\varphi_{T, \alpha, \beta}(x, y) = \frac{\sqrt{2\pi T^3}}{2xy} p_{T-(\alpha+\beta)}^*(x, y)$$

and, for every $s > 0$,

$$\psi_s(x, y, z) = \frac{(2(x+y-2z)/s) \exp(-((x+y-2z)^2 - (y-x)^2)/2s)}{(1 - \exp(-2xy/s)) \left(\frac{1}{x} + \frac{1}{y} - \frac{2z}{xy}\right)} \mathbf{1}_{\{z < x \wedge y\}}.$$

Moreover, for every $x, y, z > 0$ such that $z < x \wedge y$, we have

$$\lim_{T \rightarrow \infty} \Delta_{T, \alpha, \beta}(x, y, z) = 1.$$

Proof. We first recall a well-known fact about the Brownian excursion with fixed duration. If $0 < t_1 < t_2 < \dots < t_p < T$, the density of the law of the p -tuple $(\mathbf{e}_{t_1}^\lambda, \dots, \mathbf{e}_{t_p}^\lambda)$ is

$$2\sqrt{2\pi T^3} q_{t_1}(x_1) p_{t_2-t_1}^*(x_1, x_2) \dots p_{t_p-t_{p-1}}^*(x_{p-1}, x_p) q_{T-t_p}(x_p).$$

See e.g. Chapter XII of [22] for a proof. To simplify notation, we set, for $0 < t_1 < \dots < t_p$ and $x_1, \dots, x_p > 0$,

$$F_{t_1, \dots, t_p}(x_1, \dots, x_p) = q_{t_1}(x_1) p_{t_2-t_1}^*(x_1, x_2) \dots p_{t_p-t_{p-1}}^*(x_{p-1}, x_p).$$

Then fix $0 < t_1 < \dots < t_p = \alpha$ and $0 < t'_1 < \dots < t'_q = \beta$. We see that the density of $(\mathbf{e}_{t_1}^\lambda, \dots, \mathbf{e}_{t_p}^\lambda, \mathbf{e}_{T-t'_q}^\lambda, \dots, \mathbf{e}_{T-t'_1}^\lambda)$ is the function

$$\begin{aligned} g_{t_1, \dots, t_p, t'_1, \dots, t'_q}^T(x_1, \dots, x_{p+q}) \\ = 2\sqrt{2\pi T^3} F_{t_1, \dots, t_p}(x_1, \dots, x_p) F_{t'_1, \dots, t'_q}(x_{p+q}, \dots, x_{p+q}) p_{T-(\alpha+\beta)}^*(x_p, x_{p+1}). \end{aligned}$$

On the other hand, from the explicit formulas for the transition density of the three-dimensional Bessel process, we have

$$F_{t_1, \dots, t_p}(x_1, \dots, x_p) = \frac{1}{2x_p} p_{t_1}^{(3)}(0, x_1) p_{t_2-t_1}^{(3)}(x_1, x_2) \dots p_{t_p-t_{p-1}}^{(3)}(x_{p-1}, x_p) = \frac{1}{2x_p} G_{t_1, \dots, t_p}^{(3)}(x_1, \dots, x_p),$$

where $G_{t_1, \dots, t_p}^{(3)}$ is the density of $(R_{t_1}, \dots, R_{t_p})$. Consequently, we have also

$$\begin{aligned} g_{t_1, \dots, t_p, t'_1, \dots, t'_q}^T(x_1, \dots, x_{p+q}) \\ = 2\sqrt{2\pi T^3} G_{t_1, \dots, t_p}^{(3)}(x_1, \dots, x_p) G_{t'_1, \dots, t'_q}^{(3)}(x_{p+q}, \dots, x_{p+q}) \times \frac{p_{T-(\alpha+\beta)}^*(x_p, x_{p+1})}{4x_p x_{p+1}}. \end{aligned}$$

It follows from this calculation that the density of the law of the pair

$$\left((\mathbf{e}_t^\lambda)_{0 \leq t \leq \alpha}, (\mathbf{e}_{T-t}^\lambda)_{0 \leq t \leq \beta} \right)$$

with respect to the law of

$$\left((R_t)_{0 \leq t \leq \alpha}, (R'_t)_{0 \leq t \leq \beta} \right),$$

is given by the function $\varphi_{T, \alpha, \beta}(\omega(\alpha), \omega'(\beta))$.

We next consider the conditional density of the third component of each triplet in the proposition, given its first two components. We note that, for every $z > 0$,

$$\begin{aligned} P \left[\inf_{t \geq \alpha} R_t \wedge \inf_{t \geq \beta} R'_t > z \mid (R_t)_{0 \leq t \leq \alpha}, (R'_t)_{0 \leq t \leq \beta} \right] &= P \left[\inf_{t \geq \alpha} R_t > z \mid R_\alpha \right] P \left[\inf_{t \geq \beta} R'_t > z \mid R'_\beta \right] \\ &= \left(1 - \frac{z \wedge R_\alpha}{R_\alpha} \right) \left(1 - \frac{z \wedge R'_\beta}{R'_\beta} \right). \end{aligned}$$

Hence the conditional density of

$$\inf_{t \geq \alpha} R_t \wedge \inf_{t \geq \beta} R'_t$$

given $(R_t)_{0 \leq t \leq \alpha}$ and $(R'_t)_{0 \leq t \leq \beta}$ is

$$h_{R_\alpha, R'_\beta}(z) = \mathbf{1}_{\{0 < z < R_\alpha \wedge R'_\beta\}} \left(\frac{1}{R_\alpha} + \frac{1}{R'_\beta} - \frac{2z}{R_\alpha R'_\beta} \right).$$

Similarly, to get the conditional distribution of

$$\min_{\alpha \leq t \leq T-\beta} \mathbf{e}_t^\lambda,$$

we observe that, conditionally on $(\mathbf{e}_t^\lambda)_{0 \leq t \leq \alpha}$ and $(\mathbf{e}_{T-t}^\lambda)_{0 \leq t \leq \beta}$, the process $(\mathbf{e}_{\alpha+t}^\lambda)_{0 \leq t \leq T-(\alpha+\beta)}$ is a Brownian bridge of duration $T - (\alpha + \beta)$, starting from $\mathbf{e}_\alpha^\lambda$ and ending at $\mathbf{e}_{T-\beta}^\lambda$, and conditioned not to hit 0.

Fix $s > 0$ and $x > 0$, and recall that for a standard Brownian motion B starting from x the density of the pair $(B_s, \min_{0 \leq r \leq s} B_r)$ is the function $g_s(y, z) = 2q_s(x + y - 2z) \mathbf{1}_{\{z < x \wedge y\}}$. Hence, if we also fix $y > 0$, the conditional density of $\min_{0 \leq r \leq s} B_r$ knowing that $B_s = y$ is

$$f_{s,x,y}(z) = \frac{2q_s(x + y - 2z)}{p_s(y - x)} \mathbf{1}_{\{z < x \wedge y\}}.$$

In particular,

$$P\left(\min_{0 \leq r \leq s} B_r > 0 \mid B_s = y\right) = \int_0^{x \wedge y} dz f_{s,x,y}(z) = 1 - \exp\left(-\frac{2xy}{s}\right).$$

We conclude from these calculations that the density of $\min_{\alpha \leq t \leq T-\beta} \mathbf{e}_t^\lambda$ knowing $(\mathbf{e}_t^\lambda)_{0 \leq t \leq \alpha}$ and $(\mathbf{e}_{T-t}^\lambda)_{0 \leq t \leq \beta}$ is the function $f_{T-(\alpha+\beta), \mathbf{e}_\alpha^\lambda, \mathbf{e}_{T-\beta}^\lambda}^*(z)$, where, for $s > 0$ and $x, y, z > 0$,

$$f_{s,x,y}^*(z) = \frac{f_{t,x,y}(z)}{1 - \exp\left(-\frac{2xy}{s}\right)} \mathbf{1}_{\{0 < z < x \wedge y\}}.$$

Set

$$m_{\alpha,\beta} = \inf_{t \geq \alpha} R_t \wedge \inf_{t \geq \beta} R'_t$$

to simplify notation. It follows that, for every nonnegative measurable function Γ on $C([0, \alpha], \mathbb{R}_+) \times C([0, \beta], \mathbb{R}_+) \times \mathbb{R}$,

$$\begin{aligned} & E\left[\Gamma\left((\mathbf{e}_t^\lambda)_{0 \leq t \leq \alpha}, (\mathbf{e}_{T-t}^\lambda)_{0 \leq t \leq \beta}, \min_{\alpha \leq t \leq T-\beta} \mathbf{e}_t^\lambda\right)\right] \\ &= E\left[\int_0^\infty dz f_{T-(\alpha+\beta), \mathbf{e}_\alpha^\lambda, \mathbf{e}_{T-\beta}^\lambda}^*(z) \Gamma\left((\mathbf{e}_t^\lambda)_{0 \leq t \leq \alpha}, (\mathbf{e}_{T-t}^\lambda)_{0 \leq t \leq \beta}, z\right)\right] \\ &= E\left[\varphi_{T,\alpha,\beta}(R_\alpha, R'_\beta) \int_0^\infty dz f_{T-(\alpha+\beta), R_\alpha, R'_\beta}^*(z) \Gamma\left((R_t)_{0 \leq t \leq \alpha}, (R'_t)_{0 \leq t \leq \beta}, z\right)\right] \\ &= E\left[\int_0^\infty dz h_{R_\alpha, R'_\beta}(z) \Delta_{T,\alpha,\beta}(R_\alpha, R'_\beta, z) \Gamma\left((R_t)_{0 \leq t \leq \alpha}, (R'_t)_{0 \leq t \leq \beta}, z\right)\right] \\ &= E\left[\Delta_{T,\alpha,\beta}(R_\alpha, R'_\beta, m_{\alpha,\beta}) \Gamma\left((R_t)_{0 \leq t \leq \alpha}, (R'_t)_{0 \leq t \leq \beta}, m_{\alpha,\beta}\right)\right]. \end{aligned}$$

In the second equality, we applied the first part of the proof, and in the third one, we used the identity

$$\varphi_{T,\alpha,\beta}(R_\alpha, R'_\beta) f_{T-(\alpha+\beta), R_\alpha, R'_\beta}^*(z) = h_{R_\alpha, R'_\beta}(z) \Delta_{T,\alpha,\beta}(R_\alpha, R'_\beta, z),$$

which follows from our definitions. The proof of the first assertion of the proposition is now complete. The last assertion follows from the explicit expressions for $\varphi_{T,\alpha,\beta}(x, y)$ and $\psi_{T-(\alpha+\beta)}(x, y, z)$. \square

3.2 Proof of Theorem 1

In this section, we prove Theorem 1. Recall the notation $\mathcal{P} = (\mathcal{P}, D_\infty, \rho_\infty)$ for the Brownian plane viewed as a pointed metric space.

Proposition 4. *Let $\delta > 0$ and $r \geq 0$. There exists $\lambda_0 > 0$ such that, for every $\lambda \geq \lambda_0$, we can construct copies of $\lambda \cdot \mathbf{m}_\infty$ and \mathcal{P} on the same probability space, in such a way that the equality*

$$B_r(\lambda \cdot \mathbf{m}_\infty) = B_r(\mathcal{P}) \quad (7)$$

holds with probability at least $1 - \delta$.

Proposition 4 immediately implies that, for every $r \geq 0$,

$$B_r(\lambda \cdot \mathbf{m}_\infty) \xrightarrow[\lambda \rightarrow \infty]{(d)} B_r(\mathcal{P}) \quad (8)$$

in distribution in \mathbb{K} . From the observations of the end of subsection 2.1, this suffices to get the convergence (1) in Theorem 1. The first assertion in Theorem 1 also readily follows from Proposition 4 and the scale invariance of the Brownian plane.

Proof. We will rely on the identity in distribution (5). To simplify notation, we write $Y^\lambda = \mathcal{T}_{e^\lambda} / \approx_\lambda$. We also write \mathbf{Y}^λ for the pointed space $(Y^\lambda, D_\lambda^*, \rho_\lambda)$, so that $\lambda \cdot \mathbf{m}_\infty$ has the same distribution as \mathbf{Y}^λ .

We first choose $A > 1$ sufficiently large so that, if $\beta = (\beta_t)_{t \geq 0}$ is a standard linear Brownian motion, we have

$$P\left[\left\{\min_{0 \leq t \leq A} \beta_t < -10r\right\} \cap \left\{\min_{A \leq t \leq A^2} \beta_t < -10r\right\} \cap \left\{\min_{A^2 \leq t \leq A^4} \beta_t < -10r\right\}\right] \geq 1 - \delta/8.$$

Next we choose $\alpha > 0$ large enough so that the property

$$\inf_{t \geq \alpha} R_t \wedge \inf_{t \geq \alpha} R'_t > A^4$$

holds with probability at least $1 - \delta/8$.

From Proposition 3 (especially the last assertion of this proposition) and standard coupling arguments, we can find $\lambda_0 > 0$, such that $\lambda_0^4 > 2\alpha$, and such that, for every $\lambda \geq \lambda_0$, we can construct both processes \mathbf{e}^λ and (R, R') on the same probability space, in such a way that the probability of the event

$$\mathcal{E}_\lambda := \{\mathbf{e}_t^\lambda = R_t \text{ and } \mathbf{e}_{\lambda^4 - t}^\lambda = R'_t, \forall t \leq \alpha\} \cap \left\{\min_{\alpha \leq t \leq \lambda^4 - \alpha} \mathbf{e}_t^\lambda = \inf_{t \geq \alpha} R_t \wedge \inf_{t \geq \alpha} R'_t\right\}$$

is bounded below by $1 - \delta/2$.

We fix $\lambda \geq \lambda_0$ and we write $T = \lambda^4$ as in subsection 3.1. We suppose that \mathbf{e}^λ and the pair (R, R') have been constructed so that $P(\mathcal{E}_\lambda) > 1 - \delta/2$. We then observe that, conditionally given \mathbf{e}^λ , the process $(W_t^\lambda)_{t \in [-\alpha, \alpha]}$ given by

$$W_t^\lambda = \begin{cases} Z_t^\lambda & \text{if } t \in [0, \alpha] \\ Z_{T+t}^\lambda & \text{if } t \in [-\alpha, 0] \end{cases}$$

is a Gaussian process whose covariance is a function of the triplet

$$\left((\mathbf{e}_t^\lambda)_{0 \leq t \leq \alpha}, (\mathbf{e}_{T-t}^\lambda)_{0 \leq t \leq \alpha}, \min_{\alpha \leq t \leq T-\alpha} \mathbf{e}_t^\lambda\right).$$

Similarly, conditionally given (R, R') , the process $(Z_t)_{t \in [-\alpha, \alpha]}$ is also Gaussian with covariance given by the *same* function of the triplet

$$\left((R_t)_{0 \leq t \leq \alpha}, (R'_t)_{0 \leq t \leq \alpha}, \inf_{t \geq \alpha} R_t \wedge \inf_{t \geq \alpha} R'_t \right).$$

Note that \mathcal{E}_λ is precisely the event where the latter two triplets of processes coincide. It follows that we can construct the processes Z^λ and Z in such a way that, on the event \mathcal{E}_λ , we have also the identities

$$Z_t^\lambda = Z_t \quad , \quad Z_{T-t}^\lambda = Z_{-t} \quad , \quad \forall t \in [0, \alpha].$$

From now on, we assume that these identities hold on \mathcal{E}_λ .

To simplify notation, we write $\mathcal{T}_{(\lambda)} = \mathcal{T}_{\mathbf{e}^\lambda}$, and $p_{(\lambda)}$ for the canonical projection from $[0, T]$ onto $\mathcal{T}_{(\lambda)}$. For every $x \in [0, \mathbf{e}_{T/2}^\lambda]$, we set

$$\gamma_\lambda(x) = \sup\{t \leq T/2 : \mathbf{e}_t^\lambda = x\} \quad , \quad \eta_\lambda(x) = \inf\{t \geq T/2 : \mathbf{e}_t^\lambda = x\}.$$

Then $p_{(\lambda)}(\gamma_\lambda(x)) = p_{(\lambda)}(\eta_\lambda(x))$ is the vertex at distance x from the root in the ancestral line of the vertex $p_{(\lambda)}(T/2)$ in the tree $\mathcal{T}_{(\lambda)}$. Similarly, we define, for every $x \geq 0$,

$$\gamma_\infty(x) = \sup\{t \geq 0 : R_t = x\} \quad , \quad \eta_\infty(x) = \sup\{t \geq 0 : R'_t = x\}.$$

With the notation of the introduction, we have $p_\infty(\gamma_\infty(x)) = p_\infty(-\eta_\infty(x))$ for every $x \geq 0$. Furthermore, the process $(Z_{\gamma_\infty(x)})_{x \geq 0}$ is distributed as a standard linear Brownian motion.

Write \mathcal{F}_λ for the intersection of \mathcal{E}_λ with the event where we have both

$$\inf_{t \geq \alpha} R_t \wedge \inf_{t \geq \alpha} R'_t > A^4$$

and

$$\min_{0 \leq x \leq A} Z_{\gamma_\infty(x)} < -10r \quad , \quad \min_{A \leq x \leq A^2} Z_{\gamma_\infty(x)} < -10r \quad , \quad \min_{A^2 \leq x \leq A^4} Z_{\gamma_\infty(x)} < -10r. \quad (9)$$

By our choice of the constants A and α , we have $P(\mathcal{F}_\lambda) \geq 1 - \delta$.

On the event \mathcal{F}_λ , we have

$$\min_{\alpha \leq t \leq T-\alpha} \mathbf{e}_t^\lambda = \inf_{t \geq \alpha} R_t \wedge \inf_{t \geq \alpha} R'_t > A^4$$

and therefore, for every $x \in [0, A^4]$, $\gamma_\lambda(x) = \gamma_\infty(x) < \alpha$ and $T - \eta_\lambda(x) = \eta_\infty(x) < \alpha$. It follows that, for every $x \in [0, A^4]$,

$$Z_{\gamma_\lambda(x)}^\lambda = Z_{\gamma_\infty(x)} = Z_{-\eta_\infty(x)} = Z_{\eta_\lambda(x)}^\lambda.$$

Before stating another lemma, we introduce one more piece of notation. Let $s, t \in [0, T]$. If s and t both belong to $[0, T/2]$ or if they both belong to $[T/2, T]$, we set

$$\tilde{D}_\lambda^\circ(s, t) = Z_s^\lambda + Z_t^\lambda - 2 \min_{u \in [s \wedge t, s \vee t]} Z_u^\lambda.$$

Otherwise we set

$$\tilde{D}_\lambda^\circ(s, t) = Z_s^\lambda + Z_t^\lambda - 2 \min_{u \in [0, s \wedge t] \cup [s \vee t, T]} Z_u^\lambda.$$

Lemma 5. *Suppose that \mathcal{F}_λ holds.*

(i) *For every $t \in [\gamma_\lambda(A), \eta_\lambda(A)]$,*

$$D_\lambda^*(\rho_\lambda, p_\lambda(t)) > r.$$

For every $s, t \in [0, \gamma_\lambda(A)] \cup [\eta_\lambda(A), T]$ such that $D_\lambda^(\rho_\lambda, p_\lambda(s)) \leq r$ and $D_\lambda^*(\rho_\lambda, p_\lambda(t)) \leq r$, we have*

$$D_\lambda^*(p_\lambda(s), p_\lambda(t)) = \inf_{s_0, t_0, s_1, t_1, \dots, s_p, t_p} \sum_{i=1}^p \tilde{D}_\lambda^\circ(t_{i-1}, s_i) \quad (10)$$

where the infimum is over all choices of the integer $p \geq 0$ and of the reals $s_0, s_1, \dots, s_p, t_0, t_1, \dots, t_p$ belonging to $[0, \gamma_\lambda(A^2)] \cup [\eta_\lambda(A^2), T]$ such that $s_0 = s$, $t_p = t$, and $p_\lambda(s_i) = p_\lambda(t_i)$ for every $i \in \{0, 1, \dots, p\}$.

(ii) *For every $t' \in (-\infty, -\eta_\infty(A)] \cup [\gamma_\infty(A), \infty)$,*

$$D_\infty(\rho_\infty, p_\infty(t')) > r.$$

For every $s', t' \in [-\eta_\infty(A), \gamma_\infty(A)]$ such that $D_\infty(\rho_\infty, p_\infty(s')) \leq r$ and $D_\infty(\rho_\infty, p_\infty(t')) \leq r$, we have

$$D_\infty(p_\infty(s'), p_\infty(t')) = \inf_{s'_0, t'_0, s'_1, t'_1, \dots, s'_p, t'_p} \sum_{i=1}^p D_\infty^\circ(t'_{i-1}, s'_i) \quad (11)$$

where the infimum is over all choices of the integer $p \geq 0$ and of the reals $s'_0, s'_1, \dots, s'_p, t'_0, t'_1, \dots, t'_p$ belonging to $[-\eta_\infty(A^2), \gamma_\infty(A^2)]$ such that $s'_0 = s'$, $t'_p = t'$, and $p_\infty(s'_i) = p_\infty(t'_i)$ for every $i \in \{0, 1, \dots, p\}$.

Proof. We argue on the event \mathcal{F}_λ and prove (i). We observe that, for every $t \in [\gamma_\lambda(A), \eta_\lambda(A)]$, we have

$$D_\lambda^*(\rho_\lambda, p_\lambda(t)) \geq Z_t^\lambda - 2 \min_{c \in \llbracket \rho_\lambda, p_\lambda(t) \rrbracket} Z_c^\lambda \geq - \min_{c \in \llbracket \rho_\lambda, p_\lambda(t) \rrbracket} Z_c^\lambda \geq 10r.$$

The first inequality is the cactus bound (4), and the last one follows from the fact that if $t \in [\gamma_\lambda(A), \eta_\lambda(A)]$, the ancestral line of $p_\lambda(t)$ contains all vertices of the form $p_\lambda(\gamma_\lambda(x))$ for $x \in [0, A]$, and we use the equality $Z_{\gamma_\lambda(x)}^\lambda = Z_{\gamma_\infty(x)}$ (which holds for these values of x on \mathcal{F}_λ), together with the first bound in (9).

We next turn to the proof of the second assertion in (i). Let $s, t \in [0, \gamma_\lambda(A)] \cup [\eta_\lambda(A), T]$ such that $D^*(\rho_\lambda, p_\lambda(s)) \leq r$ and $D^*(\rho_\lambda, p_\lambda(t)) \leq r$. Note that we have then

$$|Z_s^\lambda| \leq r \quad , \quad |Z_t^\lambda| \leq r$$

by the cactus bound. Furthermore, we have by definition

$$D_\lambda^*(p_\lambda(s), p_\lambda(t)) = \inf_{s=s_0, t_0, s_1, t_1, \dots, s_p, t_p=t} \sum_{i=1}^p D_\lambda^\circ(t_{i-1}, s_i) \quad (12)$$

where the infimum is over all choices of the integer $p \geq 0$ and of the reals $s_0, s_1, \dots, s_p, t_0, t_1, \dots, t_p$ in $[0, T]$ such that $s_0 = s$, $t_p = t$ and $p_\lambda(s_i) = p_\lambda(t_i)$ for every $i \in \{0, 1, \dots, p\}$.

Since $D_\lambda^*(p_\lambda(s), p_\lambda(t)) \leq 2r$, we can obviously restrict our attention to reals s_0, s_1, \dots, t_p such that

$$\sum_{i=1}^p D_\lambda^\circ(t_{i-1}, s_i) < 5r/2. \quad (13)$$

We next claim that in the infimum in the right-hand side of (12), we can furthermore limit ourselves to choices of $s_1, \dots, s_p, t_0, t_1, \dots, t_{p-1}$ such that, for every $i \in \{0, 1, \dots, p\}$, both s_i and t_i belong to $[0, \gamma_\lambda(A^2)] \cup [\eta_\lambda(A^2), T]$.

Indeed, suppose that, for some $i \in \{0, 1, \dots, p\}$, t_i (or s_i) does not belong to $[0, \gamma_\lambda(A^2)] \cup [\eta_\lambda(A^2), T]$. From (13), we have $D_\lambda^*(p_{(\lambda)}(s), p_{(\lambda)}(t_i)) < 5r/2$. On the other hand, by the cactus bound and the property $Z_s^\lambda \geq -r$, we have

$$D_\lambda^*(p_{(\lambda)}(s), p_{(\lambda)}(t_i)) \geq -r - \min_{c \in \llbracket p_{(\lambda)}(s), p_{(\lambda)}(t_i) \rrbracket} Z_c^\lambda \geq 9r$$

because the fact that $s \in [0, \gamma_\lambda(A)] \cup [\eta_\lambda(A), T]$ and $t_i \notin [0, \gamma_\lambda(A^2)] \cup [\eta_\lambda(A^2), T]$ ensures that the geodesic segment $\llbracket p_{(\lambda)}(s), p_{(\lambda)}(t_i) \rrbracket$ contains all vertices of the form $p_{(\lambda)}(\gamma_\lambda(x))$ for $x \in [A, A^2]$, and we use the equality $Z_{\gamma_\lambda(x)}^\lambda = Z_{\gamma_\infty(x)}$ (for these values of x), together with the second bound in (9). This contradiction proves our claim.

In order to establish formula (10), we still need to justify the fact that we can replace D_λ° by \tilde{D}_λ° in (12). So let s_0, t_1, \dots, t_p be reals in $[0, \gamma_\lambda(A^2)] \cup [\eta_\lambda(A^2), T]$ such that $s_0 = s$, $t_p = t$ and (13) holds. Consider first the case where $i \in \{1, \dots, p\}$ is such that $t_{i-1} \in [0, \gamma_\lambda(A^2)]$ and $s_i \in [\eta_\lambda(A^2), T]$. Using (13) we have $D_\lambda^*(p_{(\lambda)}(s), p_{(\lambda)}(t_{i-1})) < 5r/2$, hence $Z_{t_{i-1}}^\lambda \geq Z_s^\lambda - 5r/2 \geq -7r/2$, and similarly $Z_{s_i}^\lambda \geq -7r/2$. We have also

$$- \min_{u \in [t_{i-1}, s_i]} Z_u^\lambda \geq 10r$$

because the interval $[t_{i-1}, s_i]$ must contain all $\gamma_\infty(x)$ for $x \in [A^2, A^4]$, and we use the third bound of (9). Then

$$\begin{aligned} 5r/2 > D_\lambda^\circ(t_{i-1}, s_i) &= Z_{t_{i-1}}^\lambda + Z_{s_i}^\lambda - 2 \max \left(\min_{u \in [t_{i-1}, s_i]} Z_u^\lambda, \min_{u \in [0, t_{i-1}] \cup [s_i, T]} Z_u^\lambda \right) \\ &\geq -7r - 2 \max \left(\min_{u \in [t_{i-1}, s_i]} Z_u^\lambda, \min_{u \in [0, t_{i-1}] \cup [s_i, T]} Z_u^\lambda \right) \end{aligned}$$

and the previous two displays can hold only if

$$\max \left(\min_{u \in [t_{i-1}, s_i]} Z_u^\lambda, \min_{u \in [0, t_{i-1}] \cup [s_i, T]} Z_u^\lambda \right) = \min_{u \in [0, t_{i-1}] \cup [s_i, T]} Z_u^\lambda,$$

which means that $\tilde{D}_\lambda^\circ(t_{i-1}, s_i) = D_\lambda^\circ(t_{i-1}, s_i)$. Next consider the case where both t_{i-1} and s_i belong to $[0, \gamma_\lambda(A^2)]$. Then, by the same argument as previously, we have

$$- \min_{u \in [t_{i-1} \vee s_i, T]} Z_u^\lambda \geq 10r$$

and it again follows that $\tilde{D}_\lambda^\circ(t_{i-1}, s_i) = D_\lambda^\circ(t_{i-1}, s_i)$. The other cases are treated in a similar way. This completes the proof of assertion (i).

The proof of assertion (ii) is similar. Just note that a version of the cactus bound (4) where D_λ^* is replaced by D_∞ and Z^λ by Z holds for $a, b \in \mathcal{T}_\infty$, with exactly the same proof. We omit the details. \square

The next lemma is a simple corollary of Lemma 5.

Lemma 6. *Assume that \mathcal{F}_λ holds. Let $s, t \in [0, \gamma_\lambda(A)] \cup [\eta_\lambda(A), T]$. Set $s' = s$ if $s \in [0, \gamma_\lambda(A)]$ and $s' = s - T$ if $s \in [\eta_\lambda(A), T]$, and similarly $t' = t$ if $t \in [0, \gamma_\lambda(A)]$ and $t' = t - T$ if $t \in [\eta_\lambda(A), T]$. Then we have $D_\lambda^*(\rho_\lambda, p_{(\lambda)}(s)) \leq r$ and $D_\lambda^*(\rho_\lambda, p_{(\lambda)}(t)) \leq r$ if and only if $D_\infty(\rho_\infty, p_\infty(s')) \leq r$ and $D_\infty(\rho_\infty, p_\infty(t')) \leq r$. Furthermore, if these conditions hold, we have also*

$$D_\lambda^*(p_{(\lambda)}(s), p_{(\lambda)}(t)) = D_\infty(p_\infty(s'), p_\infty(t')).$$

Proof. First notice that the condition $s, t \in [0, \gamma_\lambda(A)] \cup [\eta_\lambda(A), T]$ is equivalent to saying that $s', t' \in [-\eta_\infty(A), \gamma_\infty(A)]$ (recall that $\gamma_\infty(A) = \gamma_\lambda(A)$ and $\eta_\infty(A) = T - \eta_\lambda(A)$ on \mathcal{F}_λ).

Then let $s, t \in [0, \gamma_\lambda(A)] \cup [\eta_\lambda(A), T]$ such that $D_\lambda^*(\rho_\lambda, p_{(\lambda)}(s)) \leq r$ and $D_\lambda^*(\rho_\lambda, p_{(\lambda)}(t)) \leq r$. By Lemma 5, $D_\lambda^*(p_{(\lambda)}(s), p_{(\lambda)}(t))$ is given by formula (10). We claim that the right-hand side of this formula coincides with the right-hand side of formula (11). To see this, let $s_0, t_0, s_1, t_1, \dots, s_p, t_p \in [0, T]$ such that $s_0 = s$ and $t_p = t$. For every $i \in \{0, 1, \dots, p\}$, set $s'_i = s_i$ if $s_i \leq T/2$ and $s'_i = s_i - T$ otherwise, and define t'_i analogously. Then $s_0, s_1, \dots, t_p \in [0, \gamma_\lambda(A^2)] \cup [\eta_\lambda(A^2), T]$ if and only if $s'_0, s'_1, \dots, t'_p \in [-\eta_\infty(A^2), \gamma_\infty(A^2)]$. Assume that this condition holds. Then, one immediately checks that, for every $i \in \{0, 1, \dots, p\}$, we have $p_{(\lambda)}(s_i) = p_{(\lambda)}(t_i)$ if and only if $p_\infty(s'_i) = p_\infty(t'_i)$. Finally, we have also $\tilde{D}_\lambda^\circ(t_{i-1}, s_i) = D_\infty^\circ(t'_{i-1}, s'_i)$ for every $i \in \{0, 1, \dots, p\}$ (at this point it is crucial that \tilde{D}_λ° appears instead of D_λ° in (10)). Our claim follows.

We cannot immediately infer that the right-hand side of (11) coincides with $D_\infty(p_\infty(s), p_\infty(t))$ since we do not know yet that $D_\infty(\rho_\infty, p_\infty(s)) \leq r$ and $D_\infty(\rho_\infty, p_\infty(t)) \leq r$. However, the right-hand side of (11) is clearly an upper bound for $D_\infty(p_\infty(s), p_\infty(t))$, and thus, by considering the special cases $s = 0$ or $t = 0$, we deduce from the equality between the right-hand sides of (10) and (11) that $D_\infty(\rho_\infty, p_\infty(s)) \leq r$ and $D_\infty(\rho_\infty, p_\infty(t)) \leq r$.

From a symmetric argument, we obtain that the latter conditions imply $D_\lambda^*(\rho_\lambda, p_{(\lambda)}(s)) \leq r$ and $D_\lambda^*(\rho_\lambda, p_{(\lambda)}(t)) \leq r$. Finally, the equality between the right-hand sides of (10) and (11) gives the identity $D_\lambda^*(p_{(\lambda)}(s), p_{(\lambda)}(t)) = D_\infty(p_\infty(s'), p_\infty(t'))$. \square

To complete the proof of Proposition 4, we verify that the identity (7) holds on \mathcal{F}_λ . Write $\mathbf{p}_{(\lambda)}$ for the composition of $p_{(\lambda)}$ with the canonical projection from $\mathcal{T}_{(\lambda)}$ onto Y^λ , and similarly write \mathbf{p}_∞ for the composition of p_∞ with the canonical projection from \mathcal{T}_∞ onto \mathcal{P} . Assuming that \mathcal{F}_λ holds, we construct an isometry \mathcal{I} from $B_r(\mathbf{Y}^\lambda)$ onto $B_r(\mathcal{P})$, which maps ρ_λ to ρ_∞ , in the following way. An arbitrary point of $B_r(\mathbf{Y}^\lambda)$ is of the form $\mathbf{p}_{(\lambda)}(s)$, where s has to belong to $[0, \gamma_\lambda(A)] \cup [\eta_\lambda(A), T]$ (by the first assertion in Lemma 5 (i)). We then define

$$\mathcal{I}(\mathbf{p}_{(\lambda)}(s)) = \mathbf{p}_\infty(s')$$

where $s' = s$ if $s \in [0, \gamma_\lambda(A)]$ and $s' = s - T$ if $s \in [\eta_\lambda(A), T]$, as previously. The last assertion of Lemma 6 shows that this definition does not depend on the choice of s , and then that \mathcal{I} is an isometry. Finally, using the first assertion of Lemma 5 (ii), we easily get that \mathcal{I} maps $B_r(\mathbf{Y}^\lambda)$ onto $B_r(\mathcal{P})$, and it is also immediate that $\mathcal{I}(\rho_\lambda) = \rho_\infty$. We conclude that the pointed spaces $B_r(\mathbf{Y}^\lambda)$ and $B_r(\mathcal{P})$ are isometric on the event \mathcal{F}_λ . Since $\lambda \cdot \mathbf{m}_\infty$ has the same distribution as \mathbf{Y}^λ , this completes the proof of Proposition 4 and of Theorem 1.

4 The Brownian plane as the limit of discrete quadrangulations

In this section, we prove Theorem 2, which shows that the Brownian plane also arises as a scaling limit of discrete quadrangulations. Let us briefly discuss the strategy of the proof, which

is, roughly speaking, the same as in the proof of Theorem 1. After recalling the constructions of discrete quadrangulations from labeled trees, we establish in Proposition 7 a comparison lemma for these trees, which plays the same role as Proposition 3 in the continuous setting. Using this lemma, we show in Proposition 9 that we can couple the realizations of Q_n and of Q_∞ in such a way that the balls of radius $o(n^{1/4})$ centered at the origin are the same in Q_n and Q_∞ with large probability (this is the discrete counterpart to Proposition 4). We then use the convergence towards the Brownian map (6) and Proposition 4 to complete the proof.

4.1 Discrete quadrangulations

As in the introduction, we let Q_n be uniformly distributed over the set of all rooted quadrangulations with n faces. Let us recall the definition of the uniform infinite planar quadrangulation (UIPQ). More details can be found in [9, 12, 20].

If m is a rooted planar map and $r \geq 0$, we define the “combinatorial ball” $\text{Ball}_r(m)$ as the submap of m obtained by keeping only those edges and vertices of m that are incident to (at least) one face of m having a vertex whose graph distance from the root vertex is smaller than or equal to r . Independently of the embedding chosen for m , this defines a planar map $\text{Ball}_r(m)$, which by convention has the same root as m .

Krikun [12] proved that there exists a random infinite (rooted) planar quadrangulation Q_∞ such that, for every $r \geq 0$, we have the following convergence in distribution

$$\text{Ball}_r(Q_n) \xrightarrow[n \rightarrow \infty]{(d)} \text{Ball}_r(Q_\infty). \quad (14)$$

We refer to [9] for a discussion and a precise definition of infinite planar maps. The random infinite quadrangulation Q_∞ is called the uniform infinite planar quadrangulation or UIPQ. Notice that a similar object has been defined earlier in the context of triangulations by Angel and Schramm, see [3, 4].

An alternative approach to the preceding convergence can be found in [20], where it is also proved that Q_∞ coincides with the random infinite quadrangulation that had been constructed by Chassaing and Durhuus [7] from a random infinite well-labeled tree – a different but related version of this construction, due to [9], will be presented below.

As we already mentioned after the statement of Theorem 2, it will be convenient to see discrete planar maps as length spaces: If Q is a (finite or infinite) quadrangulation, we associate with Q a pointed (boundedly compact) length space, denoted by \mathbf{Q} , which is obtained by replacing every edge of Q by a unit length Euclidean segment and pointing the resulting metric space at the root vertex of Q . More precisely, \mathbf{Q} is the union of a (finite or infinite) collection of copies of the interval $[0, 1]$ and two of these copies may intersect only at their endpoints if the associated edges of Q share one or two vertices. Obviously, the distance between two points of \mathbf{Q} is the length of a shortest path between them. If Q is finite, it is straightforward to verify that $d_{\text{GH}}(V(Q), \mathbf{Q}) \leq 1$ (here and later, we view $V(Q)$ as a pointed metric space, for the graph distance and with the root vertex of Q as distinguished point). More generally, both in the finite and the infinite case, we have $d_{\text{GH}}(B_r(V(Q)), B_r(\mathbf{Q})) \leq 1$ for every $r \geq 0$.

We write \mathbf{Q}_n , resp. \mathbf{Q}_∞ , for the pointed length space associated with Q_n , resp. with Q_∞ .

4.2 Constructing quadrangulations from discrete trees

In this section we briefly recall the construction of discrete quadrangulations from labeled trees. The finite case is presented in Section 5 of [16] and its extension to the infinite case can be found

in [9] (note that a different version of the construction in the infinite case had appeared earlier in [7]). The reader may consult the preceding references for more details.

4.2.1 Labeled trees

We start by recalling the standard formalism for plane trees. Let

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

where $\mathbb{N} = \{1, 2, \dots\}$ and $\mathbb{N}^0 = \{\emptyset\}$ by convention. An element u of \mathcal{U} is thus a (possibly empty) word made of positive integers, and $|u| \geq 0$ stands for the length of u , sometimes also called the height of u . If $u, v \in \mathcal{U}$, uv denotes the concatenation of u and v . If $v \in \mathcal{U} \setminus \{\emptyset\}$ we can write $v = uj$, where $u \in \mathcal{U}$ and $j \in \mathbb{N}$, and we say that u is the *parent* of v or that v is a *child* of u . More generally, if v is of the form uw , for $u, w \in \mathcal{U}$, we say that u is an *ancestor* of v or that v is a *descendant* of u . A *plane tree* τ is a (finite or infinite) subset of \mathcal{U} such that

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

If the plane tree τ is finite, $|\tau| = \#\tau - 1$ denotes the number of edges of τ and is called the size of τ .

A *ray* of an infinite plane tree τ is an infinite sequence u_0, u_1, u_2, \dots in τ such that $u_0 = \emptyset$ and u_i is the parent of u_{i+1} for every $i \geq 0$. If an infinite tree τ has a unique ray, we call it the *spine* of τ and denote it by $S_\tau(0), S_\tau(1), S_\tau(2), \dots$

In what follows, we say tree rather than plane tree. For every integer $n \geq 0$, we let \mathbb{T}_n denote the set of all trees with size n . The cardinality of \mathbb{T}_n is the Catalan number of order n ,

$$\text{Cat}(n) = \frac{1}{n+1} \binom{2n}{n}.$$

For every $n \geq 0$, we let T_n be uniformly distributed over \mathbb{T}_n .

The uniform infinite tree T_∞ . We now introduce the infinite “local limit” of the random trees T_n as $n \rightarrow \infty$. If τ is a plane tree and $k \in \{0, 1, 2, \dots\}$, we let $[\tau]_k := \{v \in \tau : |v| \leq k\}$ be the plane tree obtained from τ by keeping only its first k generations. It follows from the work of Kesten [11] (see also [18]) that there exists an infinite random tree T_∞ with only one spine such that, for every $k \geq 0$, we have

$$[T_n]_k \xrightarrow[n \rightarrow \infty]{(d)} [T_\infty]_k. \tag{15}$$

The tree T_∞ is known as the critical geometric Galton-Watson tree conditioned to survive. It can be described informally as follows (a rigorous presentation using the preceding formalism for trees can be found in [9]): Start with a semi-infinite line of vertices that will be the spine of T_∞ and graft to the left and to the right of each vertex of the spine independent critical geometric Galton-Watson trees (with parameter $1/2$). The root of T_∞ is obviously the first vertex of the spine. See Fig. 2.

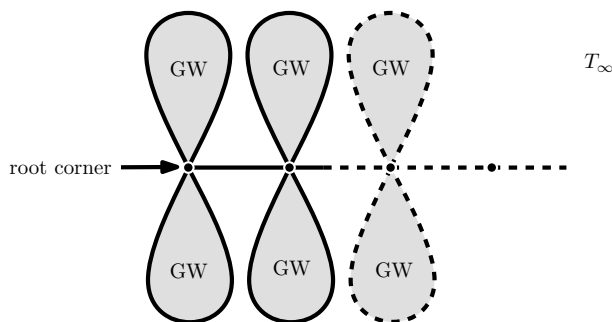


Figure 2: Construction of T_∞ . The trees that are grafted on the spine are independent Galton-Watson trees with geometric offspring distribution with parameter $1/2$.

Labeled trees. A *labeled tree* is a pair $\theta = (\tau, (\ell(u))_{u \in \tau})$ that consists of a tree τ and a collection of integer labels assigned to the vertices of τ , such that $\ell(\emptyset) = 0$ and $|\ell(u) - \ell(v)| \leq 1$ whenever $u, v \in \tau$ are neighbors (meaning that u is either the parent or a child of v). We denote the size of θ by $|\theta| = |\tau|$. We let \mathcal{T} stand for the set of all finite labeled trees and we let \mathcal{S} be the set of all infinite labeled trees $\theta = (\tau, \ell)$ such that τ has only one spine and $\inf_{i \geq 0} \ell(S_\tau(i)) = -\infty$.

If τ is a (finite or infinite) tree, we can assign labels to its vertices in a uniform way: Write E_τ for the set of all edges of τ and consider a collection $(\eta_e)_{e \in E_\tau}$ of independent random variables uniformly distributed over $\{-1, 0, +1\}$. For any vertex $u \in \tau$, the (random) label $\ell(u)$ of u is then defined as the sum of the variables η_e over all edges $e \in E_\tau$ belonging to the geodesic path from the root to u . In particular $\ell(\emptyset) = 0$. The resulting random labeled tree $\theta = (\tau, \ell)$ is called the uniform labeling of τ . If τ is random, we can still consider its uniform labeling by applying the preceding construction after conditioning on τ . This applies in particular to the random tree T_n , and the resulting random labeled tree $\Theta_n = (T_n, \ell_n)$ is uniformly distributed over the set

$$\mathcal{T}_n := \{\theta \in \mathcal{T} : |\theta| = n\}.$$

If we apply the construction to T_∞ , the resulting labeled tree $\Theta_\infty = (T_\infty, \ell_\infty)$ is called the uniform infinite labeled tree. Notice that $\Theta_\infty \in \mathcal{S}$ almost surely.

4.2.2 The construction in the finite case

We will describe Schaeffer's bijection between $\mathcal{T}_n \times \{0, 1\}$ and the set of all rooted and pointed quadrangulations with n faces. See Section 5 of [16] for more details.

A rooted and pointed quadrangulation is a pair (q, ∂) consisting of a rooted quadrangulation q together with a distinguished vertex ∂ of q . We start from a finite labeled tree $\theta = (\tau, \ell) \in \mathcal{T}_n$ and write $\min \ell$ for the minimal label on the tree. In order to associate a rooted and pointed quadrangulation with θ , we first embed the tree τ in the plane, in the way suggested by Fig.3 (on this figure the edges of τ are the dotted lines, the figures are the labels of vertices, the root vertex is obviously at the bottom and the lexicographical order between children of a given vertex corresponds to listing the edges from the left to the right). We also add an extra vertex (outside the embedded tree) denoted by ∂ . A corner c of the (embedded) tree is an angular sector between two adjacent edges. The label $\ell(c)$ of the corner c is the label of the associated vertex $\mathcal{V}(c)$. Corners of the tree have a cyclic ordering given by the clockwise contour of the tree in the plane (if we imagine a particle that moves around the embedded tree along its edges

in clockwise order, it will visit every corner exactly once before coming back to the corner it started from).




Figure 3: Construction of a rooted and pointed quadrangulation from a labeled tree.

To obtain the edges of the quadrangulation q associated with θ , we proceed as follows. For every corner c of the tree such that $\ell(c) > \min \ell$, we draw an edge between c and the first corner after c with label $\ell(c) - 1$. This corner is called the successor of c and denoted by $\mathcal{S}(c)$. If $\ell(c) = \min \ell$ (so that the definition of $\mathcal{S}(c)$ does not make sense), we draw an edge between c and ∂ . This construction can be made in such a way that the edges do not cross each other and do not cross the edges of the tree τ (see Fig.3). After erasing the embedding of the tree, the resulting map q is a quadrangulation with n faces and whose vertices are exactly the vertices of τ plus the extra vertex ∂ . Furthermore, the labeling has the following interpretation: For every $u \in \tau$, we have

$$d_{\text{gr}}(\partial, u) = \ell(u) - \min \ell + 1, \quad (16)$$

where d_{gr} stands for the graph distance in the quadrangulation.

The quadrangulation q is pointed at the vertex ∂ and its root edge is the edge of q which is drawn from the root corner c_0 of τ (the root corner is the corner “below” the root \emptyset of τ). In order to specify the orientation of the root edge, we need an extra variable $\epsilon \in \{0, 1\}$: The origin of the root edge is the vertex \emptyset if $\epsilon = 0$ and the other end of the root edge if $\epsilon = 1$.

Call $\Phi(\theta, \epsilon)$ the rooted and pointed quadrangulation that is obtained by the preceding construction. Then Φ is a bijection between $\mathcal{T}_n \times \{0, 1\}$ and the set of all rooted and pointed quadrangulations with n faces. Consequently, if Θ_n is a uniform labeled tree with n edges and η is an independent Bernoulli variable of parameter $1/2$, then the random rooted quadrangulation derived from $\Phi(\Theta_n, \eta)$ by “forgetting” the distinguished vertex has the same distribution as Q_n .

4.2.3 The construction in the infinite case

The preceding construction can easily be extended to the case when the tree $\theta = (\tau, \ell)$ is an infinite labeled tree in \mathcal{S} . We first embed the infinite tree properly in the plane and then draw




Figure 4: Construction of a rooted infinite quadrangulation from a tree of \mathcal{S} .

an edge between each corner c of the tree and the first corner with label $\ell(c) - 1$ following c in the clockwise contour order. We again denote this corner by $\mathcal{S}(c)$ and call it the successor of c . Because of our assumption $\theta \in \mathcal{S}$, there is no vertex of minimal label and thus unlike the finite case there is no need to add an extra vertex ∂ . The construction should be clear from Fig.4. The resulting planar map in an infinite quadrangulation, see [9]. We root this quadrangulation at the edge connecting the root corner of τ to its successor. As in the finite case, the orientation of this edge is given by an extra variable $\epsilon \in \{0, 1\}$. This infinite rooted quadrangulation is denoted by $\Phi'(\theta, \epsilon)$. Note that the vertex set of $\Phi'(\theta, \epsilon)$ is precisely the vertex set of τ .

Finally, from [9, Theorem 1] we know that if Θ_∞ is a uniform infinite labeled tree and η is an independent Bernoulli variable of parameter $1/2$ then

$$\Phi'(\Theta_\infty, \eta) \stackrel{(d)}{=} Q_\infty.$$

4.3 Comparison lemmas

4.3.1 Comparison of trees

One of the key ingredients in the proof of Theorem 2 is an improvement of the convergence (15). Roughly speaking, we will see that the fixed integer $k \geq 0$ in (15) can be replaced by a sequence $k = k(n)$ as soon as $k(n) = o(\sqrt{n})$ as $n \rightarrow \infty$. For technical reasons, it is more convenient to deal with *pointed* trees.

A (finite) pointed tree is a pair $\tau = (\tau, \xi)$ where τ is a finite tree and ξ is a distinguished vertex of τ . Let $\tau = (\tau, \xi)$ be a pointed tree. For every integer h such that $0 \leq h < |\xi|$, we let $\mathcal{P}(\tau, h)$ stand for the subtree of τ consisting of all vertices $u \in \tau$ such that the height of the most recent common ancestor of u and ξ is strictly less than h , together with the ancestor of ξ at height exactly h , which is denoted by $[\xi]_h$. We furthermore point this tree at $[\xi]_h$. See Fig.5. By convention when $h \geq |\xi|$, we declare that $\mathcal{P}(\tau, h) = (\{\emptyset\}, \emptyset)$.

If τ is an infinite tree with only one spine (denoted as above by $S_\tau(0), S_\tau(1), \dots$), we can extend the former definition by considering τ as pointed at infinity. Formally, for every integer

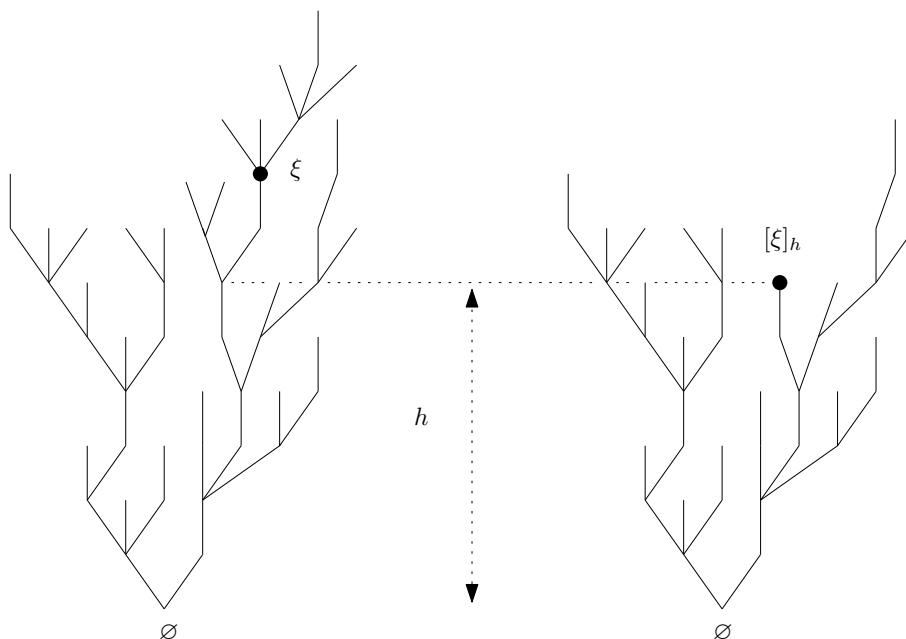


Figure 5: A pointed tree $\tau = (\tau, \xi)$ and the resulting pruned tree $\mathcal{P}(\tau, h)$ at height h .

$h \geq 0$, we let $\mathcal{P}(\tau, h)$ be the subtree of τ consisting of the vertices $S_\tau(0), \dots, S_\tau(h)$ of the spine, together with the subtrees grafted to the left and the right of $S_\tau(i)$ for $0 \leq i \leq h-1$. We point this tree at $S_\tau(h)$.

If τ is a pointed tree, and h, h' are integers such that $0 \leq h \leq h'$, then one immediately checks that

$$\mathcal{P}(\mathcal{P}(\tau, h'), h) = \mathcal{P}(\tau, h). \quad (17)$$

The same property holds if we replace τ by an infinite tree having only one spine.

In what follows, $\mathbf{T}_n = (T_n, \xi_n)$ is uniformly distributed over the set of all pointed trees with size n . This is consistent with our previous notation, since T_n is then uniformly distributed over \mathbb{T}_n . Note that, conditionally on T_n , the vertex ξ_n is uniformly distributed over the vertex set of T_n . In particular if τ_0 is a fixed pointed tree with size n , we have

$$P(\mathbf{T}_n = \tau_0) = \frac{1}{(n+1) \text{Cat}(n)} = \frac{1}{\binom{2n}{n}}.$$

Recall that T_∞ denotes the uniform infinite tree. The following proposition relates \mathbf{T}_n to T_∞ .

Proposition 7. *For every $\varepsilon > 0$, there exists $\delta > 0$ such that, for every sufficiently large integer n , the bound*

$$\left| P\left(\mathcal{P}(\mathbf{T}_n, \lfloor \delta n^{1/2} \rfloor) \in A\right) - P\left(\mathcal{P}(T_\infty, \lfloor \delta n^{1/2} \rfloor) \in A\right) \right| \leq \varepsilon,$$

holds for any set A of pointed trees. Consequently, if k_n is a sequence of positive integers such that $k_n = o(n^{1/2})$ as $n \rightarrow \infty$, then the total variation distance between the law of $\mathcal{P}(\mathbf{T}_n, k_n)$ and that of $\mathcal{P}(T_\infty, k_n)$ tends to 0 as n goes to ∞ .

Remark. It is easy to see that Proposition 7 does imply (15) and is in fact much stronger. A similar result has been proved by Aldous [1, Theorem 2] for Poisson Galton-Watson trees.

Proof. Let $h \geq 0$ be an integer. We first identify the distribution of $\mathcal{P}(T_\infty, h)$. Recall that $\mathcal{P}(T_\infty, h)$ consists of the fragment of the spine of T_∞ up to height h , together with the subtrees grafted to the left and to the right of the spine up to height $h - 1$, and that this tree is pointed at $S_{T_\infty}(h)$. The subtrees grafted on the spine of T_∞ are independent Galton-Watson trees with geometric offspring distribution of parameter $1/2$. If \mathbb{Q}_{GW} stands for the distribution of one of these trees, a standard formula for Galton-Watson trees gives, for every finite tree τ ,

$$\mathbb{Q}_{GW}(\tau) = \prod_{u \in \tau} 2^{-k_u(\tau)-1} = 2^{-2|\tau|-1} = \frac{1}{2} 4^{-|\tau|}. \quad (18)$$

Let $\tau_0 = (\tau_0, \xi)$ be a pointed tree, such that ξ is a vertex of τ_0 at height h with no child. Write $\tau_{(i)}$, $0 \leq i \leq h - 1$, resp. $\tau'_{(i)}$, $0 \leq i \leq h - 1$, for the successive subtrees that branch from the left side, resp. from the right side, of the ancestral line of ξ in τ_0 . It follows from the previous observations that

$$P(\mathcal{P}(T_\infty, h) = \tau_0) = \prod_{i=0}^{h-1} \left(\frac{1}{2} 4^{-|\tau_{(i)}|}\right) \left(\frac{1}{2} 4^{-|\tau'_{(i)}|}\right) = 4^{-|\tau_0|}. \quad (19)$$

We also notice that the size of the random pointed tree $\mathcal{P}(T_\infty, h)$ can be written in the form

$$|\mathcal{P}(T_\infty, h)| = h + \sum_{i=0}^{h-1} (\mathbf{N}_{i,\ell} + \mathbf{N}_{i,r}),$$

where the random variables $\mathbf{N}_{i,\ell}$, $\mathbf{N}_{i,r}$, for $0 \leq i \leq h - 1$ are independent and distributed according to the size of a Galton-Watson tree with geometric offspring distribution of parameter $1/2$. For every integer $n \geq 0$, (18) gives

$$P(\mathbf{N}_{0,\ell} = n) = \frac{1}{2} \text{Cat}(n) 4^{-n} \underset{n \rightarrow \infty}{\sim} \frac{n^{-3/2}}{2\sqrt{\pi}}.$$

Standard facts about domains of attraction (see for example [5, p. 343-350]) thus imply that $n^{-2} |\mathcal{P}(T_\infty, n)|$ converges in distribution towards a stable law with parameter $1/2$ (we could also derive this from the fact that $2\mathbf{N}_{0,\ell} + 1$ is distributed as the hitting time of 1 for simple random walk on \mathbb{Z} started from 0). In particular, if $\varepsilon \in (0, 1)$ is fixed, there exists a constant $c_\varepsilon > 1$ such that for every integer $k \geq 0$ we have

$$P\left(|\mathcal{P}(T_\infty, k)| \leq c_\varepsilon k^2\right) \geq 1 - \varepsilon. \quad (20)$$

We now compute the distribution of $\mathcal{P}(\mathbf{T}_n, h)$. Let the pointed tree $\tau_0 = (\tau_0, \xi)$ be as previously, with $|\xi| = h$. The event $\{\mathcal{P}(\mathbf{T}_n, h) = \tau_0\}$ holds if and only if the tree T_n is obtained from the tree τ_0 by grafting at ξ a subtree \mathfrak{t} having $n - |\tau_0|$ edges and if furthermore the distinguished point ξ_n of T_n is in \mathfrak{t} but is different from its root. Hence a direct counting argument shows that

$$P(\mathcal{P}(\mathbf{T}_n, h) = \tau_0) = \mathbf{1}_{\{|\tau_0| < n\}} \frac{(n - |\tau_0|) \text{Cat}(n - |\tau_0|)}{(n + 1) \text{Cat}(n)}.$$

Recall that we have fixed $\varepsilon \in (0, 1)$. Using asymptotics for Catalan numbers, we can find an integer $N_\varepsilon \geq 1$, which does not depend on h nor on the choice of the pointed tree $\tau_0 = (\tau_0, \xi)$ satisfying the preceding properties, such that, for every integer $n \geq |\tau_0| + N_\varepsilon$,

$$(1 - \varepsilon) 4^{-|\tau_0|} \left(1 - \frac{|\tau_0|}{n}\right)^{-1/2} \leq P(\mathcal{P}(\mathbf{T}_n, h) = \tau_0) \leq (1 + \varepsilon) 4^{-|\tau_0|} \left(1 - \frac{|\tau_0|}{n}\right)^{-1/2},$$

and therefore, using (19),

$$(1 - \varepsilon) \left(1 - \frac{|\tau_0|}{n}\right)^{-1/2} \leq \frac{P(\mathcal{P}(\mathbf{T}_n, h) = \tau_0)}{P(\mathcal{P}(T_\infty, h) = \tau_0)} \leq (1 + \varepsilon) \left(1 - \frac{|\tau_0|}{n}\right)^{-1/2}. \quad (21)$$

Recall the constant c_ε from (20). We choose $\delta > 0$ small enough so that $c_\varepsilon \delta^2 < 1/2$ and $(1 - c_\varepsilon \delta^2)^{-1/2} < 1 + \varepsilon$. We apply (21) with $h = \lfloor \delta n^{1/2} \rfloor$, and we get that, for every $n \geq 2N_\varepsilon$, and for every pointed tree $\tau_0 = (\tau_0, \xi)$ with $|\tau_0| \leq c_\varepsilon \delta^2 n$, such that $|\xi| = \lfloor \delta n^{1/2} \rfloor$ and ξ has no child,

$$1 - \varepsilon \leq \frac{P(\mathcal{P}(\mathbf{T}_n, \lfloor \delta n^{1/2} \rfloor) = \tau_0)}{P(\mathcal{P}(T_\infty, \lfloor \delta n^{1/2} \rfloor) = \tau_0)} \leq (1 + \varepsilon)^2.$$

By (20), we have $P(|\mathcal{P}(T_\infty, \lfloor \delta n^{1/2} \rfloor)| \leq c_\varepsilon \delta^2 n) \geq 1 - \varepsilon$, and it then follows from the preceding bounds that, for every $n \geq 2N_\varepsilon$,

$$P(|\mathcal{P}(\mathbf{T}_n, \lfloor \delta n^{1/2} \rfloor)| \leq c_\varepsilon \delta^2 n) \geq (1 - \varepsilon)^2.$$

Finally, if $n \geq 2N_\varepsilon$ and A is any set of pointed trees,

$$\begin{aligned} & \left| P(\mathcal{P}(\mathbf{T}_n, \lfloor \delta n^{1/2} \rfloor) \in A) - P(\mathcal{P}(T_\infty, \lfloor \delta n^{1/2} \rfloor) \in A) \right| \\ & \leq P(|\mathcal{P}(T_\infty, \lfloor \delta n^{1/2} \rfloor)| > c_\varepsilon \delta^2 n) + P(|\mathcal{P}(\mathbf{T}_n, \lfloor \delta n^{1/2} \rfloor)| > c_\varepsilon \delta^2 n) \\ & \quad + \sum_{\tau_0 \in A, |\tau_0| \leq c_\varepsilon \delta^2 n} P(\mathcal{P}(T_\infty, \lfloor \delta n^{1/2} \rfloor) = \tau_0) \left| \frac{P(\mathcal{P}(\mathbf{T}_n, \lfloor \delta n^{1/2} \rfloor) = \tau_0)}{P(\mathcal{P}(T_\infty, \lfloor \delta n^{1/2} \rfloor) = \tau_0)} - 1 \right| \\ & \leq \varepsilon + (1 - (1 - \varepsilon)^2) + (1 + \varepsilon)^2 - 1 \\ & \leq 5\varepsilon. \end{aligned}$$

This completes the proof of the first assertion of the proposition. The second assertion easily follows from the first one by using the composition property (17). \square

4.3.2 Comparison of maps

In this section we use Proposition 7 to derive a version of the convergence (14) where the radius r can tend to infinity with n . Recall that if Q is a (finite or infinite) rooted quadrangulation, the vertex set $V(Q)$ of Q is viewed as a pointed metric space. To simplify notation, we write $B_r(Q) = B_r(V(Q))$. Note that $B_r(Q)$ is again a metric space pointed at the root vertex of Q .

Lemma 8. *Let $\theta = (\tau, \ell)$ be a finite labeled tree and let $\theta' = (\tau', \ell')$ be an infinite labeled tree in \mathcal{S} . Let $\eta \in \{0, 1\}$ and let $Q = \Phi(\theta, \eta)$, resp. $Q' = \Phi(\theta', \eta)$, be the rooted quadrangulation, resp. the rooted infinite quadrangulation, constructed from the pair (θ, η) , resp. from the pair*

(θ', η) , via Schaeffer's bijection. Assume that there exist $\xi \in \tau$ and an integer $h \geq 1$ such that $\mathcal{P}((\tau, \xi), h) = \mathcal{P}(\tau', h)$ and $\ell(u) = \ell'(u)$ for every $u \in \mathcal{P}((\tau, \xi), h)$. Set

$$r = - \min_{0 \leq i \leq h} \ell(S_{\tau'}(i)) \geq 0.$$

Assume that $r \geq 3$ and set $r' = \frac{1}{2}(r - 3)$. Then,

$$B_{r'}(Q) = B_{r'}(Q').$$

Remark. The conclusion of the lemma should be understood in the sense that $B_{r'}(Q)$ and $B_{r'}(Q')$ are isometric as pointed metric spaces.

Proof. We use the notation d_{gr}^Q , resp. $d_{\text{gr}}^{Q'}$ for the graph distance on $V(Q)$, resp. on $V(Q')$. Recall that $V(Q)$ can be identified with the tree τ (plus an extra vertex that plays no role in this proof) and similarly $V(Q')$ is identified with τ' . If $u, v \in \tau$, write $[[u, v]]$ for the geodesic path between u and v in the tree τ . A simple consequence of the construction of edges in Schaeffer's bijection is the discrete cactus bound (see formula (4) in [9] and compare with (4) above) stating that

$$d_{\text{gr}}^Q(u, v) \geq \ell(u) + \ell(v) - 2 \min_{w \in [[u, v]]} \ell(w).$$

The same bound holds for $d_{\text{gr}}^{Q'}(u, v)$ when $u, v \in \tau'$, replacing ℓ by ℓ' .

Let $k \in \{0, 1, \dots, h\}$ be such that

$$\ell(S_{\tau'}(k)) = \min_{0 \leq i \leq h} \ell(S_{\tau'}(i)).$$

Note that we have also $\mathcal{P}((\tau, \xi), k) = \mathcal{P}(\tau', k)$ by our assumption and (17). We then observe that, if $u \in \tau \setminus \mathcal{P}(\tau, k)$, the ancestral line of u coincides with the ancestral line of ξ at least up to level k , and thus must contain the vertex $S_{\tau'}(k)$ (which belongs to the latter ancestral line by our assumption $\mathcal{P}((\tau, \xi), h) = \mathcal{P}(\tau', h)$). The cactus bound then gives

$$d_{\text{gr}}^Q(\emptyset, u) \geq - \min_{w \in [[\emptyset, u]]} \ell(w) \geq -\ell(S_{\tau'}(k)) = r.$$

If $u \in \tau' \setminus \mathcal{P}(\tau', k)$, the same argument gives

$$d_{\text{gr}}^{Q'}(\emptyset, u) \geq r.$$

Now let $u \in \mathcal{P}((\tau, \xi), k) = \mathcal{P}(\tau', k)$ such that $d_{\text{gr}}^Q(\emptyset, u) \leq r - 1$. Then any vertex v that belongs to a geodesic path from \emptyset to u in Q must satisfy $d_{\text{gr}}^Q(\emptyset, v) \leq r - 1$ and therefore also belong to $\mathcal{P}((\tau, \xi), k)$. We claim that any edge of Q that appears on this geodesic path must correspond to an edge of Q' with the same endpoints. This follows from the construction of edges in Schaeffer's bijection (both in the finite and in the infinite case), except that we must rule out the possibility of an edge starting from the left side of the ancestral line of $[\xi]_k = S_{\tau'}(k)$ in $\mathcal{P}((\tau, \xi), k)$ and ending on the right side of this ancestral line (obviously such edges do not appear in Q'). However, such an edge would start from a corner c and end at a corner c' such that the set of all corners between c and c' (in cyclic ordering) would contain a corner of the vertex $S_{\tau'}(k)$. But then the label of both endpoints of the edge would be smaller than $\ell(S_{\tau'}(k)) = -r$. This is absurd since labels on the geodesic path from \emptyset to u must be greater than or equal to $-r + 1$ since $d_{\text{gr}}^Q(\emptyset, u) \leq r - 1$.

The preceding discussion entails that if $u \in \mathcal{P}((\tau, \xi), k) = \mathcal{P}(\tau', k)$ is such that $d_{\text{gr}}^Q(\emptyset, u) \leq r - 1$ then $d_{\text{gr}}^{Q'}(\emptyset, u) \leq d_{\text{gr}}^Q(\emptyset, u)$. But the same argument (in fact easier since we do not have to rule out the possibility of edges going from the left side of the spine to its right side) also shows that if $u \in \mathcal{P}(\tau', k)$ is such that $d_{\text{gr}}^{Q'}(\emptyset, u) \leq r - 1$ then $d_{\text{gr}}^Q(\emptyset, u) \leq d_{\text{gr}}^{Q'}(\emptyset, u)$. Hence vertices that are at distance less than $r - 1$ from \emptyset are the same in Q and in Q' .

We next observe that, if u and v are two vertices of $\mathcal{P}((\tau, \xi), k)$ such that $d_{\text{gr}}^Q(\emptyset, u) \leq \frac{1}{2}(r - 1)$ and $d_{\text{gr}}^Q(\emptyset, v) \leq \frac{1}{2}(r - 1)$, we have

$$d_{\text{gr}}^Q(u, v) = d_{\text{gr}}^{Q'}(u, v). \quad (22)$$

Indeed, any vertex w on a geodesic path from u to v in Q must be at d_{gr}^Q -distance at most $\frac{1}{2}(r - 1)$ from either u or v , and thus at $d_{\text{gr}}^{Q'}$ -distance at most $r - 1$ from \emptyset . By the same argument as previously, any edge on a geodesic path from u to v in Q corresponds to an edge in Q' , and the converse also holds. The equality (22) follows.

Finally (22) and the preceding considerations show that the ball of radius $\frac{1}{2}(r - 1)$ centered at \emptyset in Q is isometric to the same ball in Q' . Since the root vertex of both Q and Q' is either \emptyset (if $\eta = 0$) or the successor of the first corner of \emptyset , which is at graph distance 1 from \emptyset , the conclusion of the lemma follows. \square

Recall that Q_n is uniformly distributed over the set of all rooted quadrangulations with n edges and that Q_∞ is the uniform infinite planar quadrangulation.

Proposition 9. *For every $\varepsilon > 0$, there exists $\alpha > 0$ such that, for every sufficiently large integer n and every $m \in \{n + 1, n + 2, \dots\}$, we can construct Q_n , Q_m and Q_∞ on the same probability space, in such a way that the equalities*

$$B_{\alpha n^{1/4}}(Q_n) = B_{\alpha n^{1/4}}(Q_m) = B_{\alpha n^{1/4}}(Q_\infty)$$

hold with probability at least $1 - \varepsilon$.

Proof. Let $\varepsilon > 0$. From the first assertion of Proposition 7, we can find $\delta > 0$ and an integer $n_0 \geq 0$ such that the following holds. If $n \geq n_0$ and $m \in \{n + 1, n + 2, \dots\}$ are fixed, we can construct on the same probability space a uniformly distributed pointed labeled tree with n edges (T_n, ξ_n, ℓ_n) , a uniformly distributed pointed labeled tree with m edges (T_m, ξ_m, ℓ_m) , and a uniform infinite labeled tree $\Theta_\infty = (T_\infty, \ell_\infty)$, in such a way that the event

$$E_{m,n} = \left\{ \mathcal{P}(T_\infty, \lfloor \delta n^{1/2} \rfloor) = \mathcal{P}((T_m, \xi_m), \lfloor \delta n^{1/2} \rfloor) = \mathcal{P}((T_n, \xi_n), \lfloor \delta n^{1/2} \rfloor), \right. \\ \left. \ell_\infty \Big|_{\mathcal{P}(T_\infty, \lfloor \delta n^{1/2} \rfloor)} = \ell_m \Big|_{\mathcal{P}((T_m, \xi_m), \lfloor \delta n^{1/2} \rfloor)} = \ell_n \Big|_{\mathcal{P}((T_n, \xi_n), \lfloor \delta n^{1/2} \rfloor)} \right\}$$

holds with probability at least $1 - \varepsilon/2$. Set

$$M_n = - \min_{0 \leq i \leq \lfloor \delta n^{1/2} \rfloor} \ell_\infty(S_{T_\infty}(i)).$$

Since M_n is distributed as the maximal value of a random walk started at 0 with increments uniformly distributed over $\{-1, 0, +1\}$ and stopped after $\lfloor \delta n^{1/2} \rfloor$ steps, Donsker's invariance principle shows that there exists a constant $\delta' > 0$, which does not depend on n , such that the event $F_n := \{M_n \geq \delta' n^{1/4}\}$ holds with probability at least $1 - \varepsilon/2$.

Let η be a Bernoulli variable of parameter $1/2$ independent of the triplet $(\Theta_n, \Theta_m, \Theta_\infty)$. We set $Q_n = \Phi((T_n, \ell_n), \eta)$, $Q_m = \Phi((T_m, \ell_m), \eta)$ and $Q_\infty = \Phi'(\Theta_\infty, \eta)$. By Lemma 8, the equalities

$$B_{\frac{1}{2}(\lfloor \delta' n^{1/4} \rfloor - 3)}(Q_n) = B_{\frac{1}{2}(\lfloor \delta' n^{1/4} \rfloor - 3)}(Q_\infty) = B_{\frac{1}{2}(\lfloor \delta' n^{1/4} \rfloor - 3)}(Q_m)$$

hold on the event $E_{m,n} \cap F_n$ whose probability is at least $1 - \varepsilon$. We just have to take $\alpha = \delta'/4$ to complete the proof. \square

4.4 Proof of Theorem 2

Let $(k_n)_{n \geq 1}$ be a sequence of non-negative real numbers converging to ∞ such that $k_n = o(n^{1/4})$ as $n \rightarrow \infty$. We will prove simultaneously that, for every $r > 0$,

$$B_r(k_n^{-1} \cdot \mathbf{Q}_n) \xrightarrow[n \rightarrow \infty]{(d)} B_r(\mathcal{P}), \quad (23)$$

$$B_r(k_n^{-1} \cdot \mathbf{Q}_\infty) \xrightarrow[n \rightarrow \infty]{(d)} B_r(\mathcal{P}), \quad (24)$$

where the convergence holds in distribution in \mathbb{K} . Both assertions of Theorem 2 follow from these convergences (see also the comments following the statement of the theorem).

We take $r = 1$ to simplify notation. We first notice that when proving (23), we may replace $B_1(k_n^{-1} \cdot \mathbf{Q}_n)$ by $k_n^{-1} \cdot B_{k_n}(Q_n)$, simply because the Gromov-Hausdorff distance between $B_1(k_n^{-1} \cdot \mathbf{Q}_n) = k_n^{-1} \cdot B_{k_n}(\mathbf{Q}_n)$ and $k_n^{-1} \cdot B_{k_n}(Q_n)$ is bounded above by $1/k_n$ – see the beginning of this section. Similarly, when proving (24), we may replace $B_1(k_n^{-1} \cdot \mathbf{Q}_\infty)$ by $k_n^{-1} \cdot B_{k_n}(Q_\infty)$.

Let $\varepsilon > 0$. By Proposition 4, we may find a constant $\lambda_0 > 0$ such that, for every $\lambda \geq \lambda_0$, we can construct the Brownian plane \mathcal{P} and the Brownian map \mathbf{m}_∞ simultaneously on the same probability space, in such a way that the equality

$$B_1(\lambda \cdot \mathbf{m}_\infty) = B_1(\mathcal{P}) \quad (25)$$

holds with probability at least $1 - \varepsilon$.

Then choose $\alpha > 0$ such that the conclusion of Proposition 9 holds. We may assume that $\alpha < (2\lambda_0)^{-1}$. Without loss of generality we may also assume that $k_n \leq \alpha \lfloor n^{1/4} \rfloor$ for every n . We set $m_n = \lceil \alpha^{-1} k_n \rceil^4$. Notice that $m_n \leq n$ and $k_n \leq \alpha m_n^{1/4}$. Using Proposition 9 (with n replaced by m_n and m replaced by n) and the notation of this proposition, we can for every sufficiently large n construct Q_n , Q_{m_n} and Q_∞ simultaneously on the same probability space, in such a way that the equalities

$$B_{k_n}(Q_n) = B_{k_n}(Q_\infty) = B_{k_n}(Q_{m_n}), \quad (26)$$

hold with probability at least $1 - \varepsilon$. From the convergence of uniform quadrangulations towards the Brownian map (6), we have

$$(V(Q_{m_n}), (\alpha^{-1} k_n)^{-1} d_{\text{gr}}, \rho_{(m_n)}) \xrightarrow[n \rightarrow \infty]{(d)} \left(\frac{8}{9}\right)^{1/4} \cdot \mathbf{m}_\infty.$$

where $\rho_{(m_n)}$ is the root vertex of Q_{m_n} as in (6). This implies

$$k_n^{-1} \cdot B_{k_n}(Q_{m_n}) \xrightarrow[n \rightarrow \infty]{(d)} B_1(\lambda \cdot \mathbf{m}_\infty). \quad (27)$$

where $\lambda := \left(\frac{8}{9}\right)^{1/4} \alpha^{-1}$. Notice that $\lambda \geq \lambda_0$ since $\alpha < (2\lambda_0)^{-1}$. Hence, as already noted, we can construct the Brownian plane \mathcal{P} and the Brownian map \mathbf{m}_∞ simultaneously on the same probability space in such a way that (25) holds with probability at least $1 - \varepsilon$. Thus, for any bounded continuous function $F : \mathbb{K} \rightarrow \mathbb{R}$, we have

$$\begin{aligned} E \left[\left| F(k_n^{-1} \cdot B_{k_n}(Q_n)) - F(B_1(\mathcal{P})) \right| \right] &\leq E \left[\left| F(k_n^{-1} \cdot B_{k_n}(Q_n)) - F(k_n^{-1} \cdot B_{k_n}(Q_{m_n})) \right| \right] \\ &+ E \left[\left| F(k_n^{-1} \cdot B_{k_n}(Q_{m_n})) - F(B_1(\lambda \cdot \mathbf{m}_\infty)) \right| \right] \\ &+ E \left[\left| F(B_1(\lambda \cdot \mathbf{m}_\infty)) - F(B_1(\mathcal{P})) \right| \right]. \end{aligned}$$

By (26) and (25), the first and the third terms in the right-hand side of the last display are bounded by $2\varepsilon \sup |F|$, whereas the convergence (27) entails that the second term tends to 0 as $n \rightarrow \infty$. We have proved that $k_n^{-1} \cdot B_{k_n}(Q_n) \rightarrow B_1(\mathcal{P})$ in distribution in the Gromov-Hausdorff sense, and as noted at the beginning of the proof, this suffices to get (23). The proof of (24) is similar: just replace $B_{k_n}(Q_n)$ by $B_{k_n}(Q_\infty)$ in the last display, still using (26). This completes the proof of the theorem. \square

4.5 An application to separating cycles

Proposition 9 can be used to relate the large scale properties of the UIPQ to those of random quadrangulations and of the Brownian map. The sphericity of the Brownian map was already applied in [17] to prove the non-existence of small bottlenecks in large uniform quadrangulations. We present here a similar property of the UIPQ, which partially answers a question raised by Krikun [12].

Recall from subsection 4.1 our notation $\text{Ball}_r(Q)$ for the ‘‘combinatorial ball’’ of radius r associated with a finite or infinite quadrangulation Q . Notice that $\text{Ball}_r(Q)$ does not determine the metric ball $B_r(Q)$ (and the converse does not hold either) because the knowledge of $\text{Ball}_r(Q)$ does not determine distances between vertices of $B_r(Q)$. However, a minor modification in the proof of Lemma 8 shows that the assumptions of this lemma also imply that $\text{Ball}_{r'}(Q) = \text{Ball}_{r'}(Q')$ – just notice that all edges of the submap $\text{Ball}_{r'}(Q)$ will start and end at a corner of the pruned tree $\mathcal{P}((\tau, \xi), h)$. Hence the statement of Proposition 9 also remains valid if we replace the metric balls $B_{\alpha n^{1/4}}(Q_n)$, $B_{\alpha n^{1/4}}(Q_m)$ and $B_{\alpha n^{1/4}}(Q_\infty)$ by the combinatorial balls of the same radius.

Let Q be a finite or infinite rooted quadrangulation. Let $p \geq 2$, let e_1, \dots, e_p be *oriented* edges of Q and let x_1, \dots, x_p be the respective origins of e_1, \dots, e_p . Also set $x_{p+1} = x_1$. We say that (e_1, \dots, e_p) is an *injective cycle* of Q of length p if x_1, \dots, x_p are distinct and if x_{i+1} is the target of e_i , for every $1 \leq i \leq p$. When $p = 2$, we also exclude the case when e_2 is the same edge as e_1 with reverse orientation. If $\mathcal{C} = (e_1, \dots, e_p)$ is an injective cycle of Q , the concatenation of e_1, \dots, e_p gives a closed loop on the sphere, whose complement has exactly two connected components by Jordan’s theorem. Suppose that Q is infinite. We say that the cycle \mathcal{C} separates the origin from infinity if the root vertex of Q is not a vertex of \mathcal{C} and if the connected component containing the root vertex contains finitely many vertices of Q . In what follows, we say cycle instead of injective cycle.

Consider the special case when Q is the UIPQ Q_∞ . For every integer $n \geq 2$, Krikun [12, Section 3.5] constructs a cycle \mathcal{C}_n separating the origin from infinity in Q_∞ , which is contained in $B_{2n}(Q_\infty) \setminus B_n(Q_\infty)$, and such that its expected length grows linearly in n when $n \rightarrow \infty$. Krikun [12, Conjecture 1] also conjectures that the minimal length of a cycle contained in the

complement of $B_n(Q_\infty)$ and separating the origin from infinity must grow linearly in n , a.s. The following corollary is a first step towards this conjecture.

Corollary 10. *Let $\kappa > 1$ be an integer and let $f : \mathbb{N} \rightarrow \mathbb{R}_+$ be a function such that $f(n) = o(n)$ as $n \rightarrow \infty$. The probability that there exists an injective cycle of Q_∞ with length less than $f(n)$, which separates the origin from infinity and whose vertices belong to $B_{\kappa n}(Q_\infty) \setminus B_n(Q_\infty)$, tends to 0 as $n \rightarrow \infty$.*

Proof. Fix $\varepsilon > 0$, and choose $\alpha > 0$ small enough so that the conclusion of Proposition 9 (or rather of the variant of this proposition for combinatorial balls, as explained above) holds for every sufficiently large n . Set

$$m = \lceil \left(\frac{2\kappa n}{\alpha} \right)^4 \rceil$$

and notice that $\alpha m^{1/4} \geq 2\kappa n$. If n is sufficiently large, we can then construct the UIPQ Q_∞ and a uniform rooted quadrangulation Q_m with m faces, in such a way that the equality

$$\text{Ball}_{2\kappa n}(Q_\infty) = \text{Ball}_{2\kappa n}(Q_m) \tag{28}$$

holds with probability at least $1 - \varepsilon$. If the event considered in the proposition holds, there exists a cycle \mathcal{C} of Q_∞ with length less than $f(n)$ and whose vertices belong to $B_{\kappa n}(Q_\infty) \setminus B_n(Q_\infty)$, and a vertex v of Q_∞ at graph distance $2\kappa m$ from the root vertex, such that v and the root vertex of Q_∞ belong to distinct components of the complement of the cycle (choose v such that there exists a path from v to infinity that stays outside $B_{2\kappa n-1}(Q_\infty)$). Assuming that the identity (28) holds, we see that on the event of the proposition there exists a cycle \mathcal{C}' of Q_m , with length less than $f(n)$ and whose vertices belong to $B_{\kappa n}(Q_m) \setminus B_n(Q_m)$, and a vertex v' at graph distance $2\kappa m$ from the root vertex of Q_m , such that v' and the root vertex of Q_m belong to distinct components of the complement of the cycle \mathcal{C}' . Now note that the set of all vertices of Q_m lying in the connected component containing v' must have diameter at least κn for the graph distance, whereas the set of all vertices lying in the connected component of the root vertex must have diameter at least n . By Corollary 1.2 in [17], we get that the probability that both (28) and the event of the proposition hold tends to 0 as $n \rightarrow \infty$. This completes the proof. \square

5 Properties of the Brownian plane

5.1 Topology and Hausdorff dimension of the Brownian plane

The following proposition is analogous to a result of [13]. It gives a more explicit form to the definition of the equivalence relation \approx in Section 1.

Proposition 11. *Almost surely, for every $a, b \in \mathcal{T}_\infty$, the property $a \approx b$ holds if and only if $D_\infty^\circ(a, b) = 0$.*

Proof. By definition, $a \approx b$ if and only if $D_\infty(a, b) = 0$. Since $D_\infty(a, b) \leq D_\infty^\circ(a, b)$, it is obvious that the property $D_\infty^\circ(a, b) = 0$ implies $a \approx b$. We need to prove the converse. We fix $r > 0$ and $\delta \in (0, 1)$. It will be enough to prove that, with probability at least $1 - \delta$, for every $a, b \in \mathcal{T}_\infty$ such that $D_\infty(\rho_\infty, a) \leq r$ and $D_\infty(\rho_\infty, b) \leq r$, the property $D_\infty(a, b) = 0$ implies $D_\infty^\circ(a, b) = 0$.

To this end, we rely on the coupling argument of Proposition 4, and we use the notation of the proof of this proposition. In particular, we fix $\lambda \geq \lambda_0$, and we argue on the event \mathcal{F}_λ

(recall that the probability of this event is bounded below by $1 - \delta$). We also use the notation $T = \lambda^4$. Let $a, b \in \mathcal{T}_\infty$ such that $D_\infty(\rho_\infty, a) \leq r$ and $D_\infty(\rho_\infty, b) \leq r$, and assume that $D_\infty(a, b) = 0$. Write $a = p_\infty(s')$ and $b = p_\infty(t')$ for some $s', t' \in \mathbb{R}$. By Lemma 5 (ii), we must have $s', t' \in (-\eta_\infty(A), \gamma_\infty(A))$. Set $s = s'$ if $s' \geq 0$ and $s = T + s'$ if $s' < 0$ and define t similarly from t' . Note that $s, t \in [0, \gamma_\lambda(A)) \cup (\eta_\lambda(A), T]$. By Lemma 6,

$$D_\lambda^*(p_{(\lambda)}(s), p_{(\lambda)}(t)) = D_\infty(p_\infty(s'), p_\infty(t')) = 0.$$

By [13, Theorem 3.4], the condition $D_\lambda^*(p_{(\lambda)}(s), p_{(\lambda)}(t)) = 0$ implies $D_\lambda^\circ(p_{(\lambda)}(s), p_{(\lambda)}(t)) = 0$.

Consider first the case when $p_{(\lambda)}(s) = p_{(\lambda)}(t)$, meaning that

$$\mathbf{e}_s^\lambda = \mathbf{e}_t^\lambda = \min_{s \wedge t \leq r \leq s \vee t} \mathbf{e}_r^\lambda.$$

Since $s, t \in [0, \gamma_\lambda(A)) \cup (\eta_\lambda(A), T] \subset [0, \alpha] \cup [T - \alpha, T]$, and we know that

$$\min_{\alpha \leq r \leq T - \alpha} \mathbf{e}_r^\lambda = \inf_{r \geq \alpha} R_r \wedge \inf_{r \geq \alpha} R'_r > A^4,$$

it easily follows that $X_{s'} = X_{t'} = m_X(s', t')$, and consequently $a = p_\infty(s') = p_\infty(t') = b$, so that obviously $D_\infty^\circ(a, b) = 0$.

Consider then the case when $p_{(\lambda)}(s) \neq p_{(\lambda)}(t)$. Then, the fact that $D_\lambda^\circ(p_{(\lambda)}(s), p_{(\lambda)}(t)) = 0$ implies that $p_{(\lambda)}(s)$ and $p_{(\lambda)}(t)$ are both leaves of $\mathcal{T}_{(\lambda)}$ (see Lemma 3.2 in [17]). It follows that

$$D_\lambda^\circ(s, t) = D_\lambda^\circ(p_{(\lambda)}(s), p_{(\lambda)}(t)) = 0.$$

Hence we have

$$\min_{r \in [s \wedge t, s \vee t]} Z_r^\lambda = Z_s^\lambda = Z_t^\lambda \tag{29}$$

or

$$\min_{r \in [0, s \wedge t] \cup [s \vee t, T]} Z_r^\lambda = Z_s^\lambda = Z_t^\lambda. \tag{30}$$

Suppose first that $s, t \in [0, \gamma_\lambda(A))$. Then since

$$\min_{r \in [\gamma_\lambda(A), \eta_\lambda(A)]} Z_r^\lambda < -10r$$

and $Z_s^\lambda = Z_{s'} \geq -r$ (because $D_\infty(\rho_\infty, a) \leq r$), it is obvious that (30) cannot hold, so that (29) holds. Similarly, if $s, t \in (\eta_\lambda(A), T]$, we get that (29) holds. Finally, if $s \in [0, \gamma_\lambda(A))$ and $t \in (\eta_\lambda(A), T]$ (or conversely), we obtain that (30) holds.

In all three cases, we can now verify that

$$\min_{r \in [s' \wedge t', s' \vee t']} Z_r = Z_{s'} = Z_{t'}$$

so that $D_\infty^\circ(s', t') = 0$, and $D_\infty^\circ(a, b) = 0$. This completes the proof. \square

Proposition 12. *The Hausdorff dimension of \mathcal{P} is almost surely equal to 4.*

This statement immediately follows from Proposition 4, together with the fact that the Hausdorff dimension of the Brownian map, or of any nontrivial ball in the Brownian map, is equal to 4, see Theorem 6.1 in [13]. We leave the details to the reader.

Proposition 13. *Almost surely, the space \mathcal{P} is homeomorphic to the plane.*

Proof. We adapt the arguments of [17] to our setting. We write $\widehat{\mathcal{P}}$ for the Alexandroff compactification of \mathcal{P} . We will prove that $\widehat{\mathcal{P}}$ is almost surely homeomorphic to the sphere \mathbb{S}^2 . Recall the definition of the equivalence relation \sim_X in the first section, and with a slight abuse of notation, define the relation \approx on \mathbb{R} by setting $s \approx t$ if and only if

$$Z_s = Z_t = \min_{r \in [s \wedge t, s \vee t]} Z_r$$

or equivalently if $D_\infty^\circ(s, t) = 0$. Note that \approx is also an equivalence relation on \mathbb{R} . Furthermore, it follows from Lemma 3.2 in [17] (and the coupling argument of the proof of Proposition 4) that, almost surely for every $s, t, u \in \mathbb{R}$, the properties $s \sim_X t$ and $s \approx u$ may hold simultaneously only if $s = t$ or $s = u$. It follows that, outside a set of probability zero which we discard from now on, we can define another equivalence relation \simeq on \mathbb{R} by setting $s \simeq t$ if and only if $s \sim_X t$ or $s \approx t$.

Write $\widehat{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$ for the Alexandroff compactification of \mathbb{R} , and extend both equivalence relations \sim_X and \simeq to $\widehat{\mathbb{R}}$ by declaring that the equivalence class of ∞ is a singleton. Equip $\widehat{\mathbb{R}}/\sim_X$ with the quotient topology (clearly the resulting space is the Alexandroff compactification of \mathcal{T}_∞). Let $\Pi : \mathbb{R}/\sim_X = \mathcal{T}_\infty \rightarrow \mathcal{P}$ be the canonical projection, and extend it to a projection from $\widehat{\mathbb{R}}/\sim_X$ onto $\widehat{\mathcal{P}}$ by mapping the equivalence class of ∞ to the point at infinity in $\widehat{\mathcal{P}}$. Then Π is continuous. The only point to check is the continuity at ∞ , which follows from the easy fact that $\Pi^{-1}(B_r(\mathcal{P}))$ is a compact subset of \mathcal{T}_∞ , for every $r \geq 0$. The projection Π then factorizes through a bijection from $\widehat{\mathbb{R}}/\simeq$ onto $\widehat{\mathcal{P}}$, which is also continuous hence provides a homeomorphism from $\widehat{\mathbb{R}}/\simeq$ onto $\widehat{\mathcal{P}}$.

We then use Proposition 2.4 in [17] to see that the quotient space $\widehat{\mathbb{R}}/\simeq$ is a.s. homeomorphic to the sphere \mathbb{S}^2 . Our setting is slightly different since [17] deals with a quotient space of the circle \mathbb{S}^1 . This makes no real difference since $\widehat{\mathbb{R}}$ is homeomorphic to \mathbb{S}^1 . Another difference is the fact that the random functions X and Z used to define the equivalence relations \sim_X and \approx are not defined at the point ∞ , whereas [17] considers functions that are (continuous and) defined everywhere on the circle. Nonetheless one can easily check that the arguments of the proof of Proposition 2.4 in [17] go through without change, provided we can verify that the local minima of X , respectively of Z , are distinct. In the case of X , this is a standard fact that follows from the connections between linear Brownian motion and the three-dimensional Bessel process. The case of Z is treated by arguments similar to the proof of Lemma 3.1 in [17].

We can now conclude that $\widehat{\mathcal{P}}$ is (almost surely) homeomorphic to \mathbb{S}^2 , and the statement of the proposition follows. \square

5.2 Geodesic rays in the Brownian plane

The results of this section are analogous to the discrete results proved in [9, Section 3.2] for the uniform infinite planar quadrangulation, and are also closely related to the study of geodesics in the Brownian map [14].

Let $x \in \mathcal{P}$. A geodesic ray from x is an infinite continuous path $\omega : [0, \infty) \rightarrow \mathcal{P}$ such that $\omega(0) = x$ and $D_\infty(\omega(s), \omega(t)) = |t - s|$ for every $s, t \geq 0$.

Recall our notation \mathbf{p}_∞ for the canonical projection from \mathbb{R} onto \mathcal{P} . Fix $x \in \mathcal{P}$ and let $t \in \mathbb{R}$ such that $\mathbf{p}_\infty(t) = x$. For every $r \geq 0$, set

$$\gamma_t(r) = \inf\{s \geq t : Z_s = Z_t - r\}.$$

It is clear that $\gamma_t(r) < \infty$ for every $r \geq 0$, a.s. Then set $\omega_t(r) = \mathbf{p}_\infty(\gamma_t(r))$, for every $r \geq 0$.

Proposition 14. *For every $x \in \mathcal{P}$ and $t \in \mathbb{R}$ such that $\mathbf{p}_\infty(t) = x$, ω_t is a geodesic ray from x . Such a geodesic ray will be called a simple geodesic ray.*

Proof. It is obvious that $\gamma_t(0) = t$ and thus $\omega_t(0) = x$. Then, by definition $Z_{\gamma_t(r)} = Z_t - r$, for every $r \geq 0$. It follows that, for every $r, s \geq 0$,

$$D_\infty(\omega_t(r), \omega_t(s)) \geq |Z_{\gamma_t(r)} - Z_{\gamma_t(s)}| = |r - s|.$$

On the other hand, it is also immediate from the definition of $\gamma_t(r)$, that, for every $r, s \geq 0$,

$$D_\infty^\circ(\gamma_t(r), \gamma_t(s)) = |r - s|.$$

The equality $D_\infty(\omega_t(r), \omega_t(s)) = |r - s|$ now follows. \square

Proposition 15. *Almost surely, ω_0 is the unique geodesic ray from ρ_∞ .*

Proof. By Proposition 4, for every $\delta > 0$, we can find $\lambda > 0$ large enough so that, with probability at least $1 - \delta$, the ball $B_1(\mathcal{P})$ is isometric to the ball $B_1(\lambda \cdot \mathbf{m}_\infty)$. By Corollary 7.7 in [14] (and the invariance of the Brownian map under re-rooting, see Theorem 8.1 in [14]), we can find a (random) $\varepsilon > 0$ such that all geodesic paths from ρ to a point outside $B_1(\lambda \cdot \mathbf{m}_\infty)$ coincide over the interval $[0, \varepsilon]$. Now, let \mathcal{R} be the set of all geodesic rays from ρ_∞ , and set

$$\tau = \inf\{t \geq 0 : \exists \omega \in \mathcal{R} : \omega(t) \neq \omega_0(t)\}.$$

By the previous considerations, $\tau > 0$ a.s. However, the scaling invariance of \mathcal{P} guarantees that τ has the same distribution as $\lambda\tau$, for every $\lambda > 0$. It follows that $\tau = \infty$ a.s., which completes the proof. \square

Our goal is to prove that all geodesic rays in \mathcal{P} are simple geodesic rays. We will rely on the preceding proposition and on the following invariance property of the Brownian plane under re-rooting.

Proposition 16. *Let $t \in \mathbb{R}$. The pointed space $(\mathcal{P}, D_\infty, \mathbf{p}_\infty(t))$ has the same distribution as $\mathcal{P} = (\mathcal{P}, D_\infty, \rho_\infty)$.*

Proof. Suppose that $t > 0$ for the definiteness. Define two processes \tilde{R} and \tilde{R}' by setting, for every $s \geq 0$,

$$\tilde{R}_s = R_t + R_{t+s} - 2 \min_{t \leq r \leq t+s} R_r$$

and

$$\tilde{R}'_s = \begin{cases} R_t + R_{t-s} - 2 \min_{t-s \leq r \leq t} R_r & \text{if } s \leq t, \\ R'_{s-t} + R_t - 2 \min_{r \in (-\infty, t-s] \cup [t, \infty)} X_r & \text{if } s > t. \end{cases}$$

With the notation of Section 1, we have $\tilde{R}_s = d_X(t, t+s)$ and $\tilde{R}'_s = d_X(t, t-s)$ for every $s \geq 0$.

We also set $\tilde{Z}_s = Z_{t+s} - Z_t$ for every $s \in \mathbb{R}$. Then, we claim that the two triplets $(\tilde{R}, \tilde{R}', \tilde{Z})$ and (R, R', Z) have the same distribution. In the case when \mathcal{T}_∞ is replaced by the CRT, a similar “re-rooting invariance” identity can be found as Corollary 4.9 in Marckert and Mokkadem [19], and our claim then follows from a suitable passage to the limit: Just apply the Marckert-Mokkadem result to the (scaled) CRT coded by a Brownian excursion of duration T and let T tend to ∞ in the spirit of Section 3 above. Alternatively, the reader who is familiar with properties of the three-dimensional Bessel process will be able to verify that the pairs (R, R') and (\tilde{R}, \tilde{R}') have the same distribution, from which our claim also follows in a straightforward

manner. Notice that, in the same way as we defined \mathcal{T}_∞ in Section 1, we can associate a random tree $\tilde{\mathcal{T}}_\infty$ with the pair (\tilde{R}, \tilde{R}') , and $\tilde{\mathcal{T}}_\infty$ is easily identified to \mathcal{T}_∞ “re-rooted” at $p_\infty(t)$.

To complete the proof, just note that the pointed space $(\mathcal{P}, D_\infty, \mathbf{p}_\infty(t))$ can be obtained from the triplet $(\tilde{R}, \tilde{R}', \tilde{Z})$ by the same construction that we used to obtain $(\mathcal{P}, D_\infty, \rho_\infty)$ from (R, R', Z) . \square

Remark. By combining the last proposition with the preceding one, we obtain that almost surely for every rational t the simple geodesic ray ω_t is the unique geodesic ray from $\mathbf{p}_\infty(t)$. This will be useful in the proof of Theorem 18 below.

The next proposition shows that the quantities Z_a , $a \in \mathcal{T}_\infty$ can be interpreted as measuring the relative distances from the point at infinity in the Brownian plane. Recall that Π stands for the canonical projection from \mathcal{T}_∞ onto \mathcal{P} . The following proposition should be compared to [9, Theorem 1].

Proposition 17. *Almost surely, for every $a, b \in \mathcal{T}_\infty$,*

$$Z_a - Z_b = \lim_{x \rightarrow \infty} (D_\infty(\Pi(a), x) - D_\infty(\Pi(b), x))$$

where the limit holds when x tends to the point at infinity in the Alexandroff compactification of \mathcal{P} . Consequently, if ω is any geodesic ray in \mathcal{P} , we have

$$Z_{\omega(r)} = Z_{\omega(0)} - r$$

for every $r \geq 0$.

Proof. We fix $s, t \in \mathbb{R}$ and take $a = p_\infty(s)$ and $b = p_\infty(t)$. To simplify notation, we set

$$m_Z(s, t) = \min_{r \in [s \wedge t, s \vee t]} Z_r.$$

From our construction, it is clear that the simple geodesic rays ω_s and ω_t coalesce in finite time. More precisely, we have for every $r \geq 0$,

$$\omega_s(Z_s - m_Z(s, t) + r) = \omega_t(Z_t - m_Z(s, t) + r). \quad (31)$$

For every integer $n \geq 1$, set

$$S_n(a) = \{x \in \mathcal{P} : D_\infty(\Pi(a), x) = n\}.$$

Let $\text{Geo}_n(a)$ be the set of all geodesic paths from $\Pi(a)$ to a point of $S_n(a)$, and

$$\eta_n(a) = \inf\{r > 0 : \exists \omega, \omega' \in \text{Geo}_n(a) : \omega(r) \neq \omega'(r)\}.$$

By combining the invariance of the Brownian plane under re-rooting (Proposition 16) with the argument already used in the proof of Proposition 15, we get that $\eta_n(a) > 0$ a.s. On the other hand, the sequence $(\eta_n(a))_{n \geq 1}$ is clearly increasing and if $\eta_\infty(a)$ denotes its limit, the scale invariance of the Brownian plane implies that $\lambda \eta_\infty(a)$ has the same distribution as $\eta_\infty(a)$, for every $\lambda > 0$. It follows that $\eta_\infty(a) = \infty$ a.s.

Consequently, we can choose n sufficiently large so that $\eta_n(a) > Z_s - m_Z(s, t)$. Then, let $x \in \mathcal{P}$ such that $D_\infty(\Pi(a), x) \geq n$, and let $\bar{\omega}$ be any geodesic path from $\Pi(a)$ to x (such a geodesic exists because the Brownian plane is a boundedly compact length space, see subsection

2.1). By the definition of $\eta_n(a)$, we have $\bar{\omega}(r) = \omega_s(r)$ for every $r \in [0, \eta_n(a)]$. Recalling (31), we can construct a continuous path ω' from $\Pi(b)$ to x by concatenating the paths

$$(\omega_t(r), 0 \leq r \leq Z_t - m_Z(s, t))$$

and

$$(\bar{\omega}(Z_s - m_Z(s, t) + r), 0 \leq r \leq D_\infty(\Pi(a), x) - (Z_s - m_Z(s, t))).$$

The length of the path ω' is

$$Z_t - m_Z(s, t) + (D_\infty(\Pi(a), x) - (Z_s - m_Z(s, t))) = D_\infty(\Pi(a), x) + (Z_t - Z_s)$$

and so we have obtained the bound

$$D_\infty(\Pi(b), x) \leq D_\infty(\Pi(a), x) + (Z_t - Z_s)$$

which holds for any $x \in \mathcal{P}$ such that $D_\infty(\Pi(a), x) \geq n$. It follows that

$$\liminf_{x \rightarrow \infty} (D_\infty(\Pi(a), x) - D_\infty(\Pi(b), x)) \geq Z_s - Z_t = Z_a - Z_b.$$

We can now interchange the roles of a and b and we get that the convergence of the proposition holds for this particular choice of a and b , almost surely.

Finally, the convergence of the proposition holds outside a set of probability zero for all a, b of the form $a = p_\infty(s)$ and $b = p_\infty(t)$ with $s, t \in \mathbb{Q}$. A simple density argument now completes the proof of the first assertion.

The last assertion of the proposition is then immediate: If ω is any geodesic ray, we have for every $0 \leq r \leq r'$,

$$D_\infty(\omega(r), \omega(r')) - D_\infty(\omega(0), \omega(r')) = -r$$

and $\omega(r') \rightarrow \infty$ as $r' \rightarrow \infty$, so that we just need to apply the first assertion. \square

Remark. The preceding proof shows that, if $a = p_\infty(s)$ and $b = p_\infty(t)$ for some fixed $s, t \in \mathbb{R}$, we have indeed $D_\infty(\Pi(a), x) - D_\infty(\Pi(b), x) = Z_a - Z_b$ for all $x \in \mathcal{P}$ such that $D(\rho_\infty, x)$ is sufficiently large, almost surely.

Theorem 18. *All geodesic rays in \mathcal{P} are simple geodesic rays.*

In particular, this implies that any pair of geodesic rays coalesces in finite time.

Proof. Let $x \in \mathcal{P}$, and let ω be a geodesic ray from x . We first assume that $x = \Pi(a)$ where a is a leaf of \mathcal{T}_∞ (the case when a is not a leaf will be treated at the end of the proof). Then, there is a unique $t \in \mathbb{R}$ such that $a = p_\infty(t)$. Fix $u > 0$. We will prove that $\omega(u) = \omega_t(u)$. Recall the notation

$$\gamma_t(u) = \inf\{s \geq t : Z_s = Z_t - u\},$$

and also set

$$\gamma'_t(u) = \sup\{s \leq t : Z_s = Z_t - u\}.$$

Note that $\mathbf{p}_\infty(\gamma_t(u)) = \mathbf{p}_\infty(\gamma'_t(u))$, and that $p_\infty(\gamma_t(u))$ and $p_\infty(\gamma'_t(u))$ are both leaves of \mathcal{T}_∞ (the latter property is a consequence of the fact that, for $t_1, t_2, t_3 \in \mathbb{R}$, the properties $t_1 \sim_X t_2$ and $t_1 \approx t_3$ may hold simultaneously only if $t_1 = t_2$ or $t_1 = t_3$, as was mentioned in the proof of Proposition 13). Let c_0 be the unique vertex of \mathcal{T}_∞ such that

$$[[a, p_\infty(\gamma_t(u))] \cap [a, p_\infty(\gamma'_t(u))] = [[a, c_0]].$$

Fix a point $c_1 \in]a, c_0[$ such that the range of ω does not contain $\Pi(c_1)$. Such a point exists because, for some values of $r \in \mathbb{R}$, the set $]a, c_0[$ contains uncountably many vertices c such that $Z_c = r$, and, by the preceding proposition, the geodesic ray ω can visit $\Pi(c)$ for at most one such vertex c . We can also assume that c_1 is not a branching point of \mathcal{T}_∞ , because the set of all branching points is countable. Then let \mathcal{T}_1 be the connected component of $\mathcal{T}_\infty \setminus \{c_1\}$ that contains a , and note that \mathcal{T}_1 is bounded (otherwise this would contradict the definition of $\gamma_t(u)$ and $\gamma'_t(u)$).

Let $T = \inf\{s \geq 0 : \omega(s) \notin \Pi(\mathcal{T}_1)\}$. By a simple argument (see the beginning of Section 3 in [14]), there must exist two vertices b and b' such that $b \in \mathcal{T}_1$, $b' \in \mathcal{T}_\infty \setminus \mathcal{T}_1$ and $\omega(T) = \Pi(b) = \Pi(b')$, so that in particular $b \approx b'$. Notice that b and b' must be leaves of \mathcal{T}_∞ , and we can define $s_1, s'_1 \in \mathbb{R}$ by the conditions $p_\infty(s_1) = b$, $p_\infty(s'_1) = b'$. Since $b \approx b'$, we have $Z_s \geq Z_{s_1} = Z_{s'_1}$ for every $s \in [s_1 \wedge s'_1, s_1 \vee s'_1]$. We can now pick any rational s in $[s_1 \wedge s'_1, s_1 \vee s'_1]$, and we obtain from our definitions that the range of the geodesic ray ω_s must contain $\mathbf{p}_\infty(s_1 \vee s'_1) = \omega(T)$. So there exists $u_1 \geq 0$ such that $\omega_s(u_1) = \omega(T)$. Let $\tilde{\omega}$ be the infinite path obtained from the concatenation of $(\omega_s(r), 0 \leq r \leq u_1)$ and $(\omega(r), r \geq T)$. Then, it is easy to verify that $\tilde{\omega}$ is a geodesic ray: If $r \in [0, u_1]$ and $r' \in [u_1, \infty)$, the bound $D_\infty(\tilde{\omega}(r), \tilde{\omega}(r')) \leq r' - r$ is clear from the triangle inequality, but the reverse bound is also easy by writing $D_\infty(\tilde{\omega}(r), \tilde{\omega}(r')) \geq |Z_{\tilde{\omega}(r)} - Z_{\tilde{\omega}(r')}|$.

Finally, since $\tilde{\omega}$ is a geodesic ray starting from $\mathbf{p}_\infty(s)$ with a rational value of s , we know that $\tilde{\omega}$ must coincide with the simple geodesic ray ω_s . Since ω_s clearly visits $\omega_t(u)$ it follows that ω also visits $\omega_t(u)$, and finally $\omega_t(u) = \omega(u)$.

It remains to consider the case when a is not a leaf of \mathcal{T}_∞ . In that case, we can find arbitrarily small values of $r > 0$ such that $\omega(r) = \Pi(b)$, where b is leaf of \mathcal{T}_∞ (otherwise ω would have to visit an entire geodesic segment of \mathcal{T}_∞ , which is absurd). By the first part of the proof, there are arbitrarily small values of $r > 0$ such that $(\omega(r+u))_{u \geq 0}$ is a simple geodesic ray. The desired result easily follows. \square

Remark. If $x = \Pi(a)$ and a is not a leaf of \mathcal{T}_∞ , there are two (three if b is a branching point) distinct geodesic rays starting from x . This should be compared with Theorem 1.4 in [14].

References

- [1] ALDOUS, D., The continuum random tree I. *Ann. Probab.*, 19, 1-28 (1991)
- [2] ALDOUS, D., The continuum random tree III. *Ann. Probab.*, 21, 248-289 (1993)
- [3] ANGEL, O., Growth and percolation on the uniform infinite planar triangulation. *Geomet. Funct. Anal.* 3, 935-974 (2003)
- [4] ANGEL, O., SCHRAMM, O. Uniform infinite planar triangulations. *Comm. Math. Phys.* 241, 191-213 (2003)
- [5] BINGHAM, N.H., GOLDIE, C.M., TEUGELS, J.L. *Regular Variation*. Cambridge University Press, Cambridge 1989.
- [6] BURAGO, D., BURAGO, Y., IVANOV, S., *A Course in Metric Geometry*. Graduate Studies in Mathematics, vol. 33. Amer. Math. Soc., Boston, 2001.
- [7] CHASSAING, P., DURHUUS, B. Local limit of labeled trees and expected volume growth in a random quadrangulation. *Ann. Probab.* 34, 879-917 (2006)

- [8] CHASSAING, P., SCHAEFFER, G., Random planar lattices and integrated superBrownian excursion. *Probab. Th. Rel. Fields*, 128, 161-212 (2004)
- [9] CURIEN, N., MÉNARD, L., MIERMONT, G., A view from infinity of the uniform infinite quadrangulation. Preprint, arXiv:1201.1052
- [10] DUQUESNE, T., WINKEL, M. Growth of Lévy trees. *Probab. Th. Rel. Fields*, 139, 313-371 (2007)
- [11] KESTEN, H. Subdiffusive behavior of random walk on a random cluster. *Ann. Inst. H. Poincaré Probab. Stat.* 22, 425-487 (1986)
- [12] KRIKUN, M., Local structure of random quadrangulations. Preprint, math:PR/0512304
- [13] LE GALL, J.F., The topological structure of scaling limits of large planar maps. *Invent. Math.*, 169, 621-670 (2007)
- [14] LE GALL, J.F., Geodesics in large planar maps and in the Brownian map. *Acta Math.*, 205, 287-360 (2010)
- [15] LE GALL, J.F., Uniqueness and universality of the Brownian map. Preprint, arXiv:1105.4842
- [16] LE GALL, J.F., MIERMONT, G., Scaling limits of random trees and planar maps. To appear in the Clay Mathematics Proceedings, arXiv:1101.4856
- [17] LE GALL, J.F., PAULIN, F., Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere. *Geomet. Funct. Anal.*, 18 , 893-918 (2008)
- [18] LYONS, R., PERES, Y., PEMANTLE, R., Conceptual proofs of $L \log L$ criteria for mean behavior of branching processes. *Ann. Probab.* 23, 1125-1138 (1995)
- [19] MARCKERT, J.F., MOKKADEM, A., Limit of normalized quadrangulations: the Brownian map. *Ann. Probab.* 34, 2144-2202 (2006)
- [20] MÉNARD, L., The two uniform infinite quadrangulations of the plane have the same law. *Ann. Inst. H. Poincaré Probab. Stat.* 46, 190-208 (2010)
- [21] MIERMONT, G., The Brownian map is the scaling limit of uniform random plane quadrangulations. Preprint, arXiv:1104.1606
- [22] REVUZ, D., YOR, M., *Continuous Martingales and Brownian Motion*. Springer, Berlin, 1991.