

The Brownian cactus I. Scaling limits of discrete cactuses

Nicolas Curien^a, Jean-François Le Gall^b and Grégory Miermont^b

^a*DMA – Ecole Normale Supérieure 45, rue d'Ulm, 75230 Paris Cedex 05, France. E-mail: nicolas.curien@ens.fr*

^b*Département de Mathématiques – Bâtiment 425, Université Paris-Sud, 91405 Orsay Cedex, France.
E-mail: jean-francois.legall@math.u-psud.fr; gregory.miermont@math.u-psud.fr*

Received 4 March 2011; revised 16 October 2011; accepted 18 October 2011

Abstract. The cactus of a pointed graph is a discrete tree associated with this graph. Similarly, with every pointed geodesic metric space E , one can associate an \mathbb{R} -tree called the continuous cactus of E . We prove under general assumptions that the cactus of random planar maps distributed according to Boltzmann weights and conditioned to have a fixed large number of vertices converges in distribution to a limiting space called the Brownian cactus, in the Gromov–Hausdorff sense. Moreover, the Brownian cactus can be interpreted as the continuous cactus of the so-called Brownian map.

Résumé. Le cactus d'un graphe pointé est un certain arbre discret associé à ce graphe. De façon similaire, à tout espace métrique géodésique pointé E , on peut associer un \mathbb{R} -arbre appelé cactus continu de E . Sous des hypothèses générales, nous montrons que le cactus de cartes planaires aléatoires – dont la loi est déterminée par des poids de Boltzmann, et qui sont conditionnées à avoir un grand nombre fixé de sommets – converge en loi vers un espace limite appelé cactus brownien, au sens de la topologie de Gromov–Hausdorff. De plus, le cactus brownien peut être interprété comme le cactus continu de la carte brownienne.

MSC: 60F17; 60D05

Keywords: Random planar maps; Scaling limit; Brownian map; Brownian cactus; Hausdorff dimension

1. Introduction

In this work, we associate with every pointed graph a discrete tree called the cactus of the graph. Assuming that the pointed graph is chosen at random in a certain class of planar maps with a given number of vertices, and letting this number tend to infinity, we show that, modulo a suitable rescaling, the associated cactus converges to a universal object, which we call the Brownian cactus.

In order to motivate our results, let us recall some basic facts about planar maps. A planar map is a proper embedding of a finite connected graph in the two-dimensional sphere, viewed up to orientation-preserving homeomorphisms of the sphere. The faces of the map are the connected components of the complement of edges, and the degree of a face counts the number of edges that are incident to it, with the convention that if both sides of an edge are incident to the same face, this edge is counted twice in the degree of the face. Special cases of planar maps are triangulations, where each face has degree 3, quadrangulations, where each face has degree 4 and more generally p -angulations where each face has degree p . Since the pioneering work of Tutte [25], planar maps have been thoroughly studied in combinatorics, and they also arise in other areas of mathematics: See in particular the book of Lando and Zvonkin [11] for algebraic and geometric motivations. Large random planar graphs are of interest in theoretical physics, where they serve as models of random geometry [2].

A lot of recent work has been devoted to the study of scaling limits of large random planar maps viewed as compact metric spaces. The vertex set of the (random) planar map is equipped with the graph distance, and one is interested in the convergence in distribution of the (suitably rescaled) resulting metric space when the number of vertices tends to infinity, in the sense of the Gromov–Hausdorff distance. This convergence has been proved very recently [15,21] both for uniformly distributed random $2p$ -angulations and for uniformly distributed random triangulations (the problem in the case of triangulations had been stated by Schramm [24]). In the present work, we treat a similar problem, but we replace the metric space associated with a planar map by a simpler metric space called the cactus of the map. We are then able to prove, in a very general setting, the existence of a universal scaling limit, which we call the Brownian cactus. This result gives another strong indication of the universality of scaling limits of random planar maps, in the spirit of the papers [17,19,22] which were concerned with the profile of distances from a particular point.

Let us briefly explain the definition of the discrete cactus (see Section 2.1 for more details). We start from a graph \mathbf{G} with a distinguished vertex ρ . If a and b are two vertices of \mathbf{G} , and if $a_0 = a, a_1, \dots, a_p = b$ is a path from a to b in the graph \mathbf{G} , we consider the quantity

$$d_{\text{gr}}(\rho, a) + d_{\text{gr}}(\rho, b) - 2 \min_{0 \leq i \leq p} d_{\text{gr}}(\rho, a_i),$$

where d_{gr} stands for the graph distance in \mathbf{G} . The cactus distance $d_{\text{Cac}}^{\mathbf{G}}(a, b)$ is then the minimum of the preceding quantities over all choices of a path from a to b . The cactus distance is in fact only a pseudo-distance: We have $d_{\text{Cac}}^{\mathbf{G}}(a, b) = 0$ if and only if $d_{\text{gr}}(\rho, a) = d_{\text{gr}}(\rho, b)$ and if there is a path from a to b that stays at distance at least $d_{\text{gr}}(\rho, a)$ from the point ρ . The cactus $\text{Cac}(\mathbf{G})$ associated with \mathbf{G} is the quotient space of the vertex set of \mathbf{G} for the equivalence relation \asymp defined by setting $a \asymp b$ if and only if $d_{\text{Cac}}^{\mathbf{G}}(a, b) = 0$. The set $\text{Cac}(\mathbf{G})$ is equipped by the distance induced by $d_{\text{Cac}}^{\mathbf{G}}$. It is easy to verify that $\text{Cac}(\mathbf{G})$ is a discrete tree (Proposition 2.2). Although much information is lost when going from \mathbf{G} to its cactus, $\text{Cac}(\mathbf{G})$ still has a rich structure, as we will see in the case of planar maps.

A continuous analogue of the cactus can be defined for a (compact) geodesic metric space \mathbf{E} having a distinguished point ρ . As in the discrete setting, the cactus distance between two points x and y is the infimum over all continuous paths γ from x to y of the difference between the sum of the distances from x and y to the distinguished point ρ and twice the minimal distance from a point of γ to ρ . Again this is only a pseudo-distance, and the continuous cactus $\text{Kac}(\mathbf{E})$ is defined as the corresponding quotient space of \mathbf{E} . One can then check that the mapping $\mathbf{E} \rightarrow \text{Kac}(\mathbf{E})$ is continuous, and even Lipschitz, with respect to the Gromov–Hausdorff distance between pointed metric spaces (Proposition 2.7). It follows that if a sequence of (rescaled) pointed graphs \mathbf{G}_n converges towards a pointed space \mathbf{E} in the Gromov–Hausdorff sense, the (rescaled) cactuses $\text{Cac}(\mathbf{G}_n)$ also converge to $\text{Kac}(\mathbf{E})$. In particular, this implies that $\text{Kac}(\mathbf{E})$ is an \mathbb{R} -tree (we refer to [6] for the definition and basic properties of \mathbb{R} -trees).

The preceding observations yield a first approach to the convergence of rescaled cactuses associated with random planar maps. Let $p \geq 2$ be an integer, and for every $n \geq 2$, let m_n be a random planar map that is uniformly distributed over the set of all rooted $2p$ -angulations with n faces (recall that a planar map is rooted if there is a distinguished edge, which is oriented and whose origin is called the root vertex). We view the vertex set $V(m_n)$ of m_n as a metric space for the graph distance d_{gr} , with a distinguished point which is the root vertex of the map. According to [15,21], the rescaled pointed metric spaces $(V(m_n), n^{-1/4}d_{\text{gr}})$ converge in distribution in the Gromov–Hausdorff sense to a limiting (random) pointed metric space called the Brownian map. From the continuity of the mapping $\mathbf{E} \rightarrow \text{Kac}(\mathbf{E})$, one then gets that the suitably rescaled discrete cactus of m_n converges in distribution to the cactus of the Brownian map, which we call the Brownian cactus.

Let us give a brief description of the Brownian cactus. The random \mathbb{R} -tree known as the CRT, which has been introduced and studied by Aldous [1] is denoted by (\mathcal{T}_e, d_e) . The notation \mathcal{T}_e refers to the fact that the CRT is conveniently viewed as the \mathbb{R} -tree coded by a normalized Brownian excursion $\mathbf{e} = (e_t)_{0 \leq t \leq 1}$ (see Section 3 for more details). Let $(Z_a)_{a \in \mathcal{T}_e}$ be Brownian labels on the CRT. Informally, we may say that, conditionally on \mathcal{T}_e , $(Z_a)_{a \in \mathcal{T}_e}$ is a centered Gaussian process which vanishes at the root of the CRT and satisfies $\mathbb{E}[(Z_a - Z_b)^2] = d_e(a, b)$ for every $a, b \in \mathcal{T}_e$. Let a_* be the (almost surely unique) vertex of \mathcal{T}_e with minimal label. For every $a, b \in \mathcal{T}_e$, let $\llbracket a, b \rrbracket$ stand for the geodesic segment between a and b in the tree \mathcal{T}_e , and set

$$d_{\text{KAC}}(a, b) = Z_a + Z_b - 2 \min_{c \in \llbracket a, b \rrbracket} Z_c.$$

Then d_{KAC} is a pseudo-distance on \mathcal{T}_e , and the Brownian cactus KAC is the quotient space of the CRT for this pseudo-distance.

The main result of the present work (Theorem 4.5) states that the Brownian cactus is the limit in distribution of the discrete cactuses associated with random planar maps that are much more general than the uniformly distributed $2p$ -angulations discussed above. To explain this more precisely, we need to introduce Boltzmann distributions on planar maps. For technical reasons, we consider rooted and pointed planar maps, meaning that in addition to the root edge there is a distinguished vertex. Let $\mathbf{q} = (q_1, q_2, \dots)$ be a sequence of non-negative weights satisfying general assumptions (we require that \mathbf{q} has finite support, that $q_k > 0$ for some $k \geq 3$, and that \mathbf{q} is critical in the sense of [17, 19] – the latter property can always be achieved by multiplying \mathbf{q} by a suitable positive constant). For every rooted and pointed planar map m , set

$$W_{\mathbf{q}}(m) = \prod_{f \in F(m)} q_{\deg(f)},$$

where $F(m)$ stands for the set of all faces of m and $\deg(f)$ is the degree of the face f . For every n , choose a random rooted and pointed planar map M_n with n vertices, in such a way that $\mathbb{P}(M_n = m)$ is proportional to $W_{\mathbf{q}}(m)$ (to be precise, we need to restrict our attention to those integers n such that there exists at least one planar map m with n vertices such that $W_{\mathbf{q}}(m) > 0$). View M_n as a graph pointed at the distinguished vertex of M_n . Then Theorem 4.5 gives the existence of a positive constant $B_{\mathbf{q}}$ such that

$$B_{\mathbf{q}} n^{-1/4} \cdot \text{Cac}(M_n) \xrightarrow[n \rightarrow \infty]{(d)} \text{KAC}$$

in the Gromov–Hausdorff sense. Here the notation $\lambda \cdot E$ means that distances in the metric space E are multiplied by the factor λ .

As in much of the previous work on asymptotics for large random planar maps, the proof of Theorem 4.5 relies on the existence [3] of “nice” bijections between planar maps and certain multitype labeled trees. It was observed in [17] (for the bipartite case) and in [19] that the tree associated with a random planar map following a Boltzmann distribution is a (multitype) Galton–Watson tree, whose offspring distributions are determined explicitly in terms of the Boltzmann weights, and which is equipped with labels that are uniformly distributed over admissible choices. This labeled tree can be conveniently coded by the two random functions called the contour process and the label process (see the end of Section 4.3). In the bipartite case, where $q_k = 0$ if k is odd, one can prove [17] that the contour process and the label process associated with the random planar map M_n converge as $n \rightarrow \infty$, modulo a suitable rescaling, towards the pair consisting of a normalized Brownian excursion and the (tip of the) Brownian snake driven by this excursion. This convergence is a key tool for studying the convergence of rescaled (bipartite) random planar maps towards the Brownian map. In our general non-bipartite setting, it is not known whether the preceding convergence still holds, but Miermont [19] observed that it does hold if the tree is replaced by a “shuffled” version. Fortunately for our purposes, although the convergence of the coding functions of the shuffled tree would not be effective to study the asymptotics of rescaled planar maps, it gives enough information to deal with the associated cactuses. This is one of the key points of the proof of Theorem 4.5 in Section 4.

The last two sections of the present work are devoted to some properties of the Brownian cactus. We first show that the Hausdorff dimension of the Brownian cactus is equal to 4 almost surely, and is therefore the same as that of the Brownian map computed in [13]. As a tool for the calculation of the Hausdorff dimension, we derive precise information on the volume of balls centered at a typical point of the Brownian cactus (Proposition 5.1). Finally, we apply ideas of the theory of the Brownian cactus to a problem about the geometry of the Brownian map. Precisely, given three “typical” points in the Brownian map, we study the existence and uniqueness of a cycle with minimal length that separates the first point from the second one and visits the third one. This is indeed a continuous version of a problem discussed by Bouttier and Guitter [4] in the discrete setting of large quadrangulations. In particular, we recover the explicit distribution of the volume of the connected components bounded by the minimizing cycle, which had been derived in [4] via completely different methods. The results of this section strongly rely on the study of geodesics in the Brownian map developed in [14].

The paper is organized as follows. In Section 2, we give the definitions and main properties of discrete and continuous cactuses, and establish connections between the discrete and the continuous case. In Section 3, after recalling the

construction and main properties of the Brownian map, we introduce the Brownian cactus and show that it coincides with the continuous cactus of the Brownian map. Section 4.5 contains the statement and the proof of our main result Theorem 4.5. As a preparation for the proof, we recall in Section 4.1 the construction and main properties of the bijections between planar maps and multitype labeled trees. Section 5 is devoted to the Hausdorff dimension of the Brownian cactus, and Section 6 deals with minimizing cycles in the Brownian map. An Appendix gathers some facts about planar maps with Boltzmann distributions, that are needed in Section 4.

2. Discrete and continuous cactuses

2.1. The discrete cactus

Throughout this section, we consider a graph $G = (V, \mathcal{E})$, meaning that V is a finite set called the vertex set and \mathcal{E} is a subset of the set of all (unordered) pairs $\{v, v'\}$ of distinct elements of V .

If $v, v' \in V$, a *path* from v to v' in G is a finite sequence $\gamma = (v_0, \dots, v_n)$ in V , such that $v_0 = v$, $v_n = v'$ and $\{v_{i-1}, v_i\} \in \mathcal{E}$, for every $1 \leq i \leq n$. The integer $n \geq 0$ is called the length of γ . We assume that G is connected, so that a path from v to v' exists for every choice of v and v' . The *graph distance* $d_{\text{gr}}^G(v, v')$ is the minimal length of a path from v to v' in G . A path with minimal length is called a geodesic from v to v' in G .

In order to define the cactus distance (see Fig. 1 for an illustration) we consider also a distinguished point ρ in V . The triplet $\mathbf{G} = (V, \mathcal{E}, \rho)$ is then called a pointed graph. With this pointed graph we associate the cactus (pseudo-)distance defined by setting for every $v, v' \in V$,

$$d_{\text{Cac}}^{\mathbf{G}}(v, v') := d_{\text{gr}}^G(\rho, v) + d_{\text{gr}}^G(\rho, v') - 2 \max_{\gamma: v \rightarrow v'} \min_{a \in \gamma} d_{\text{gr}}^G(\rho, a),$$

where the maximum is over all paths γ from v to v' in G .

Proposition 2.1. *The mapping $(v, v') \rightarrow d_{\text{Cac}}^{\mathbf{G}}(v, v')$ is a pseudo-distance on V taking integer values. Moreover, for every $v, v' \in V$,*

$$d_{\text{gr}}^G(v, v') \geq d_{\text{Cac}}^{\mathbf{G}}(v, v') \tag{1}$$

and

$$d_{\text{Cac}}^{\mathbf{G}}(\rho, v) = d_{\text{gr}}^G(\rho, v). \tag{2}$$

Proof. It is obvious that $d_{\text{Cac}}^{\mathbf{G}}(v, v) = 0$ and $d_{\text{Cac}}^{\mathbf{G}}(v, v') = d_{\text{Cac}}^{\mathbf{G}}(v', v)$. Let us verify the triangle inequality. Let $v, v', v'' \in V$ and choose two paths $\gamma_1: v \rightarrow v'$ and $\gamma_2: v' \rightarrow v''$ such that $\min_{a \in \gamma_1} d_{\text{gr}}^G(\rho, a)$ is maximal among all




Fig. 1. A planar map and on the right side the same planar map represented so that the height of every vertex coincides with its distance from the distinguished vertex ρ . We see a tree structure emerging from this picture, which corresponds to the associated cactus.

paths $\gamma : v \rightarrow v'$ in G and a similar property holds for γ_2 . The concatenation of γ_1 and γ_2 gives a path $\gamma_3 : v \rightarrow v''$ and we easily get

$$d_{\text{Cac}}^{\mathbf{G}}(v, v'') \leq d_{\text{gr}}^{\mathbf{G}}(\rho, v) + d_{\text{gr}}^{\mathbf{G}}(\rho, v'') - 2 \min_{a \in \gamma_3} d_{\text{gr}}^{\mathbf{G}}(\rho, a) \leq d_{\text{Cac}}^{\mathbf{G}}(v, v') + d_{\text{Cac}}^{\mathbf{G}}(v', v'').$$

In order to get the bound (1), let $v, v' \in V$, and choose a geodesic path γ from v to v' . Let w be a point on the path γ whose distance to ρ is minimal. Then,

$$\begin{aligned} d_{\text{gr}}^{\mathbf{G}}(v, v') &= d_{\text{gr}}^{\mathbf{G}}(v, w) + d_{\text{gr}}^{\mathbf{G}}(w, v') \geq d_{\text{gr}}^{\mathbf{G}}(\rho, v) + d_{\text{gr}}^{\mathbf{G}}(\rho, v') - 2d_{\text{gr}}^{\mathbf{G}}(\rho, w) \\ &= d_{\text{gr}}^{\mathbf{G}}(\rho, v) + d_{\text{gr}}^{\mathbf{G}}(\rho, v') - 2 \min_{a \in \gamma} d_{\text{gr}}^{\mathbf{G}}(\rho, a) \\ &\geq d_{\text{Cac}}^{\mathbf{G}}(v, v'). \end{aligned}$$

Property (2) is immediate from the definition. □

As usual, we introduce the equivalence relation $\overset{\mathbf{G}}{\sim}$ defined on V by setting $v \overset{\mathbf{G}}{\sim} v'$ if and only if $d_{\text{Cac}}^{\mathbf{G}}(v, v') = 0$. Note that $v \overset{\mathbf{G}}{\sim} v'$ if and only if $d_{\text{gr}}^{\mathbf{G}}(\rho, v) = d_{\text{gr}}^{\mathbf{G}}(\rho, v')$ and there exists a path from v to v' that stays at distance at least $d_{\text{gr}}^{\mathbf{G}}(\rho, v)$ from ρ .

The corresponding quotient space is denoted by $\text{Cac}(\mathbf{G}) = V / \overset{\mathbf{G}}{\sim}$. The pseudo-distance $d_{\text{Cac}}^{\mathbf{G}}$ induces a distance on $\text{Cac}(\mathbf{G})$, and we keep the notation $d_{\text{Cac}}^{\mathbf{G}}$ for this distance.

Proposition 2.2. *Consider the graph G° whose vertex set is $V^\circ = \text{Cac}(\mathbf{G})$ and whose edges are all pairs $\{a, b\}$ such that $d_{\text{Cac}}^{\mathbf{G}}(a, b) = 1$. Then this graph is a tree, and the graph distance $d_{\text{gr}}^{\mathbf{G}^\circ}$ on V° coincides with the cactus distance $d_{\text{Cac}}^{\mathbf{G}}$ on $\text{Cac}(\mathbf{G})$.*

Proof. Let us first verify that the graph G° is a tree. If $u \in V$ we use the notation \bar{u} for the equivalence class of u in the quotient $\text{Cac}(\mathbf{G})$. We argue by contradiction and assume that there exists a (non-trivial) cycle in $\text{Cac}(\mathbf{G})$. We can then find an integer $n \geq 3$ and vertices $x_0, x_1, x_2, \dots, x_n \in V$ such that

$$\begin{cases} \bar{x}_0 = \bar{x}_n \text{ and } \bar{x}_0, \bar{x}_1, \dots, \bar{x}_{n-1} \text{ are distinct,} \\ d_{\text{Cac}}^{\mathbf{G}}(x_i, x_{i+1}) = 1 \text{ for every } 0 \leq i \leq n-1. \end{cases}$$

Without loss of generality, we may assume that $d_{\text{gr}}^{\mathbf{G}}(\rho, x_0) = \max\{d_{\text{gr}}^{\mathbf{G}}(\rho, x_i), 0 \leq i \leq n\}$. By (2), we have $|d_{\text{gr}}^{\mathbf{G}}(\rho, x_0) - d_{\text{gr}}^{\mathbf{G}}(\rho, x_1)| \leq d_{\text{Cac}}^{\mathbf{G}}(x_0, x_1) = 1$. If $d_{\text{gr}}^{\mathbf{G}}(\rho, x_0) = d_{\text{gr}}^{\mathbf{G}}(\rho, x_1)$ then it follows from the definition of $d_{\text{Cac}}^{\mathbf{G}}$ that $d_{\text{Cac}}^{\mathbf{G}}(x_0, x_1)$ is even and thus different from 1. So we must have

$$d_{\text{gr}}^{\mathbf{G}}(\rho, x_1) = d_{\text{gr}}^{\mathbf{G}}(\rho, x_0) - 1.$$

Combining this equality with the property $d_{\text{Cac}}^{\mathbf{G}}(x_0, x_1) = 1$, we obtain that there exists a path from x_0 to x_1 that stays at distance at least $d_{\text{gr}}^{\mathbf{G}}(\rho, x_1)$ from ρ .

Using the same arguments and the equality $d_{\text{Cac}}^{\mathbf{G}}(x_0, x_{n-1}) = 1$, we obtain similarly that $d_{\text{gr}}^{\mathbf{G}}(\rho, x_{n-1}) = d_{\text{gr}}^{\mathbf{G}}(\rho, x_0) - 1 = d_{\text{gr}}^{\mathbf{G}}(\rho, x_1)$ and that there exists a path from x_{n-1} to x_0 that stays at distance at least $d_{\text{gr}}^{\mathbf{G}}(\rho, x_1)$ from ρ .

Considering the concatenation of the two paths we have constructed, we get $d_{\text{Cac}}^{\mathbf{G}}(x_1, x_{n-1}) = 0$ or equivalently $\bar{x}_1 = \bar{x}_{n-1}$. This gives the desired contradiction, and we have proved that G° is a tree.

We still have to verify the equality of the distances $d_{\text{gr}}^{\mathbf{G}^\circ}$ and $d_{\text{Cac}}^{\mathbf{G}}$ on $\text{Cac}(\mathbf{G})$. The bound $d_{\text{Cac}}^{\mathbf{G}} \leq d_{\text{gr}}^{\mathbf{G}^\circ}$ is immediate from the triangle inequality for $d_{\text{Cac}}^{\mathbf{G}}$ and the existence of a geodesic between any pair of vertices of G° . Conversely, let $a, b \in \text{Cac}(\mathbf{G})$. We can find a path (y_0, y_1, \dots, y_n) in G such that $\bar{y}_0 = a, \bar{y}_n = b$ and

$$d_{\text{Cac}}^{\mathbf{G}}(a, b) = d_{\text{gr}}^{\mathbf{G}}(\rho, y_0) + d_{\text{gr}}^{\mathbf{G}}(\rho, y_n) - 2 \min_{0 \leq j \leq n} d_{\text{gr}}^{\mathbf{G}}(\rho, y_j).$$

Set $m = \min_{0 \leq j \leq n} d_{\text{gr}}^G(\rho, y_j)$, $p = d_{\text{gr}}^G(\rho, y_0)$ and $q = d_{\text{gr}}^G(\rho, y_n)$ to simplify notation. Then set, for every $0 \leq i \leq p - m$,

$$k_i = \min\{j \in \{0, 1, \dots, n\}: d_{\text{gr}}^G(\rho, y_j) = p - i\}$$

and, for every $0 \leq i \leq q - m$,

$$\ell_i = \max\{j \in \{0, 1, \dots, n\}: d_{\text{gr}}^G(\rho, y_j) = q - i\}.$$

Then $\bar{y}_{k_0}, \bar{y}_{k_1}, \dots, \bar{y}_{k_{p-m}} = \bar{y}_{\ell_{q-m}}, \bar{y}_{\ell_{q-m-1}}, \dots, \bar{y}_{\ell_1}, \bar{y}_{\ell_0}$ is a path from a to b in G° . It follows that

$$d_{\text{gr}}^{G^\circ}(a, b) \leq p + q - 2m = d_{\text{Cac}}^G(a, b),$$

which completes the proof. □

Remark 2.3. *The notion of the cactus associated with a pointed graph strongly depends on the choice of the distinguished point ρ .*

In the next sections, we will be interested in rooted planar maps, which will even be pointed in Section 4. With such a planar map, we can associate a pointed graph in the preceding sense: just say that V is the vertex set of the map, \mathcal{E} is the set of all pairs $\{v, v'\}$ of distinct points of V such that there exists (at least) one edge of the map between v and v' , and the vertex ρ is either the root vertex, for a map that is only rooted, or the distinguished point for a map that is rooted and pointed. Note that the graph distance corresponding to this pointed graph (obviously) coincides with the usual graph distance on the vertex set of the map. Later, when we speak about the cactus of a planar map, we will always refer to the cactus of the associated pointed graph. In agreement with the notation of this section, we will use bold letters \mathbf{m}, \mathbf{M} to denote the pointed graphs associated with the planar maps m, M .

2.2. The continuous cactus

Let us recall some basic notions from metric geometry. If (E, d) is a metric space and $\gamma : [0, T] \rightarrow E$ is a continuous curve 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 $[0, T]$. Obviously $L(\gamma) \geq d(\gamma(0), \gamma(T))$.

We say that (E, d) is a geodesic space if for every $a, b \in E$ there exists a continuous curve $\gamma : [0, d(a, b)] \rightarrow E$ such that $\gamma(0) = a$, $\gamma(d(a, b)) = b$ and $d(\gamma(s), \gamma(t)) = t - s$ for every $0 \leq s \leq t \leq d(a, b)$. Such a curve γ is then called a geodesic from a to b in E . Obviously, $L(\gamma) = d(a, b)$. A pointed geodesic metric space is a geodesic space with a distinguished point ρ .

Let $\mathbf{E} = (E, d, \rho)$ be a pointed geodesic compact metric space. We define the (continuous) cactus associated with (E, d, ρ) in a way very similar to what we did in the discrete setting. We first define for every $a, b \in E$,

$$d_{\text{Kac}}^{\mathbf{E}}(a, b) = d(\rho, a) + d(\rho, b) - 2 \sup_{\gamma: a \rightarrow b} \left(\min_{0 \leq t \leq 1} d(\rho, \gamma(t)) \right),$$

where the supremum is over all continuous curves $\gamma : [0, 1] \rightarrow E$ such that $\gamma(0) = a$ and $\gamma(1) = b$.

The next proposition is then analogous to Proposition 2.1.

Proposition 2.4. *The mapping $(a, b) \rightarrow d_{\text{Kac}}^{\mathbf{E}}(a, b)$ is a pseudo-distance on E . Furthermore, for every $a, b \in E$,*

$$d_{\text{Kac}}^{\mathbf{E}}(a, b) \leq d(a, b)$$

and

$$d_{\text{Kac}}^{\mathbf{E}}(\rho, a) = d(\rho, a).$$

The proof is exactly similar to that of Proposition 2.1, and we leave the details to the reader. Note that in the proof of the bound $d_{\text{Kac}}^{\mathbf{E}}(a, b) \leq d(a, b)$ we use the existence of a geodesic from a to b .

If $a, b \in E$, we set $a \succ b$ if $d_{\text{Kac}}^{\mathbf{E}}(a, b) = 0$. We define the cactus of (E, d, ρ) as the quotient space $\text{Kac}(\mathbf{E}) := E / \succ$, which is equipped with the quotient distance $d_{\text{Kac}}^{\mathbf{E}}$. Then $\text{Kac}(\mathbf{E})$ is a compact metric space, which is pointed at the equivalence class of ρ .

Remark 2.5. *It is natural to ask whether the supremum in the definition of $d_{\text{Kac}}^{\mathbf{E}}(a, b)$ is achieved, or equivalently whether there is a continuous path γ from a to b such that*

$$d_{\text{Kac}}^{\mathbf{E}}(a, b) = d(\rho, a) + d(\rho, b) - \min_{0 \leq t \leq 1} d(\rho, \gamma(t)).$$

We will return to this question later.

2.3. Continuity properties of the cactus

Let us start by recalling the definition of the Gromov–Hausdorff distance between two pointed compact metric spaces (see [5], Section 7.4, and [9] for more details).

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

$$d_{\text{H}}^{\mathbf{E}}(A, B) := \inf\{\varepsilon > 0: A \subset B^{\varepsilon} \text{ and } B \subset A^{\varepsilon}\},$$

where $X^{\varepsilon} := \{x \in E: d(x, X) \leq \varepsilon\}$ denotes the ε -neighborhood of a subset X of E .

Definition 2.6. *If $\mathbf{E} = (E, d, \rho)$ and $\mathbf{E}' = (E', d', \rho')$ are two pointed compact metric spaces, the Gromov–Hausdorff distance between \mathbf{E} and \mathbf{E}' is*

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

where the infimum is taken over all choices of the metric space (F, δ) and 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 of isometry classes of pointed compact metric spaces. An alternative definition of this distance uses *correspondences*. A correspondence between two pointed metric spaces (E, d, ρ) and (E', d', ρ') is a subset \mathcal{R} of $E \times E'$ containing (ρ, ρ') , such that, for every $x_1 \in E$, there exists at least one point $x_2 \in E'$ such that $(x_1, x_2) \in \mathcal{R}$ and conversely, for every $y_2 \in E'$, there exists at least one point $y_1 \in E$ such that $(y_1, y_2) \in \mathcal{R}$. The distortion of the correspondence \mathcal{R} is defined by

$$\text{dis}(\mathcal{R}) := \sup\{|d(x_1, y_1) - d'(x_2, y_2)|: (x_1, x_2), (y_1, y_2) \in \mathcal{R}\}.$$

The Gromov–Hausdorff distance can be expressed in terms of correspondences by the formula

$$d_{\text{GH}}(\mathbf{E}, \mathbf{E}') = \frac{1}{2} \inf\{\text{dis}(\mathcal{R})\}, \tag{3}$$

where the infimum is over all correspondences \mathcal{R} between \mathbf{E} and \mathbf{E}' . See [5], Theorem 7.3.25, for a proof in the non-pointed case, which is easily adapted.

Proposition 2.7. *Let \mathbf{E} and \mathbf{E}' be two pointed geodesic compact metric spaces. Then,*

$$d_{\text{GH}}(\text{Kac}(\mathbf{E}), \text{Kac}(\mathbf{E}')) \leq 6d_{\text{GH}}(\mathbf{E}, \mathbf{E}').$$

Proof. It is enough to verify that, for any correspondence \mathcal{R} between \mathbf{E} and \mathbf{E}' with distortion D , we can find a correspondence \mathcal{R} between $\text{Kac}(\mathbf{E})$ and $\text{Kac}(\mathbf{E}')$ whose distortion is bounded above by $6D$. We define \mathcal{R} as the set of all pairs (a, a') such that there exists (at least) one representative x of a in E and one representative x' of a' in E' , such that $(x, x') \in \mathcal{R}$.

Let $(x, x') \in \mathcal{R}$ and $(y, y') \in \mathcal{R}$. We need to verify that

$$|d_{\text{Kac}}^{\mathbf{E}}(x, y) - d_{\text{Kac}}^{\mathbf{E}'}(x', y')| \leq 6D.$$

Fix $\varepsilon > 0$. We can find a continuous curve $\gamma : [0, 1] \rightarrow E$ such that $\gamma(0) = x$, $\gamma(1) = y$ and

$$d(\rho, x) + d(\rho, y) - 2 \min_{0 \leq t \leq 1} d(\rho, \gamma(t)) \leq d_{\text{Kac}}^{\mathbf{E}}(x, y) + \varepsilon.$$

By continuity, we may find a subdivision $0 = t_0 < t_1 < \dots < t_p = 1$ of $[0, 1]$ such that $d(\gamma(t_i), \gamma(t_{i+1})) \leq D$ for every $0 \leq i \leq p-1$. For every $0 \leq i \leq p$, set $x_i = \gamma(t_i)$, and choose $x'_i \in E'$ such that $(x_i, x'_i) \in \mathcal{R}$. We may and will take $x'_0 = x'$ and $y'_0 = y'$. Now note that, for $0 \leq i \leq p-1$,

$$d'(x'_i, x'_{i+1}) \leq d(x_i, x_{i+1}) + D \leq 2D.$$

Since \mathbf{E}' is a geodesic space, we can find a curve $\gamma' : [0, 1] \rightarrow E'$ such that $\gamma'(t_i) = x'_i$, for every $0 \leq i \leq p$, and any point $\gamma'(t)$, $0 \leq t \leq 1$, lies within distance at most D from one of the points $\gamma'(t_i)$. It follows that

$$\min_{0 \leq t \leq 1} d'(\rho', \gamma'(t)) \geq \min_{0 \leq i \leq p} d'(\rho', \gamma'(t_i)) - D \geq \min_{0 \leq i \leq p} d(\rho, \gamma(t_i)) - 2D.$$

Hence,

$$\begin{aligned} d_{\text{Kac}}^{\mathbf{E}'}(x', y') &\leq d'(\rho', x') + d'(\rho', y') - 2 \min_{0 \leq t \leq 1} d'(\rho', \gamma'(t)) \\ &\leq d(\rho, x) + d(\rho, y) - 2 \min_{0 \leq t \leq 1} d(\rho, \gamma(t)) + 6D \\ &\leq d_{\text{Kac}}^{\mathbf{E}}(x, y) + 6D + \varepsilon. \end{aligned}$$

The desired result follows since ε was arbitrary and we can interchange the roles of \mathbf{E} and \mathbf{E}' . \square

2.4. Convergence of discrete cactuses

Let $\mathbf{G} = (V, \mathcal{E}, \rho)$ be a pointed graph (and write $G = (V, \mathcal{E})$ for the non-pointed graph as previously). We can identify \mathbf{G} with the pointed (finite) metric space $(V, d_{\text{gr}}^{\mathbf{G}}, \rho)$. For any real $r > 0$, we then denote the “rescaled graph” $(V, rd_{\text{gr}}^{\mathbf{G}}, \rho)$ by $r \cdot \mathbf{G}$.

Similarly, we defined $\text{Cac}(\mathbf{G})$ as a pointed finite metric space. The space $r \cdot \text{Cac}(\mathbf{G})$ is then obtained by multiplying the distance on $\text{Cac}(\mathbf{G})$ by the factor r .

Proposition 2.8. *Let $(\mathbf{G}_n)_{n \geq 0}$ be a sequence of pointed graphs, and let $(r_n)_{n \geq 0}$ be a sequence of positive real numbers converging to 0. Suppose that $r_n \cdot \mathbf{G}_n$ converges to a pointed compact metric space \mathbf{E} , in the sense of the Gromov–Hausdorff distance. Then, $r_n \cdot \text{Cac}(\mathbf{G}_n)$ also converges to $\text{Kac}(\mathbf{E})$, in the sense of the Gromov–Hausdorff distance.*

Remark 2.9. *The cactus $\text{Kac}(\mathbf{E})$ is well defined because \mathbf{E} must be a geodesic space. The latter property can be derived from [5], Theorem 7.5.1, using the fact that the graphs $r_n \cdot \mathbf{G}_n$ can be approximated by geodesic spaces as explained in the forthcoming proof.*

Proof of Proposition 2.8. This is essentially a consequence of Proposition 2.7. We start with some simple observations. Let $\mathbf{G} = (V, \mathcal{E}, \rho)$ be a pointed graph. By considering the union of a collection $(I_{\{u,v\}})_{\{u,v\} \in \mathcal{E}}$ of unit segments indexed by \mathcal{E} (such that this union is a metric graph in the sense of [5], Section 3.2.2), we can construct a pointed

geodesic compact metric space $(\Lambda(\mathbf{G}), d_{\Lambda(\mathbf{G})}, \tilde{\rho})$, such that the graph \mathbf{G} (viewed as a pointed metric space) is embedded isometrically in $\Lambda(\mathbf{G})$, and the Gromov–Hausdorff distance between \mathbf{G} and $\Lambda(\mathbf{G})$ is bounded above by 1.

A moment’s thought shows that $\text{Cac}(\mathbf{G})$ is also embedded isometrically in $\text{Kac}(\Lambda(\mathbf{G}))$, and the Gromov–Hausdorff distance between $\text{Cac}(\mathbf{G})$ and $\text{Kac}(\Lambda(\mathbf{G}))$ is still bounded above by 1.

We apply these observations to the graphs \mathbf{G}_n . By scaling, we get that the Gromov–Hausdorff distance between the metric spaces $r_n \cdot \mathbf{G}_n$ and $r_n \cdot \Lambda(\mathbf{G}_n)$ is bounded above by r_n , so that the sequence $r_n \cdot \Lambda(\mathbf{G}_n)$ also converges to \mathbf{E} in the sense of the Gromov–Hausdorff distance. From Proposition 2.7, we now get that $\text{Kac}(r_n \cdot \Lambda(\mathbf{G}_n))$ converges to $\text{Kac}(\mathbf{E})$. On the other hand, the Gromov–Hausdorff distance between $\text{Kac}(r_n \cdot \Lambda(\mathbf{G}_n)) = r_n \cdot \text{Kac}(\Lambda(\mathbf{G}_n))$ and $r_n \cdot \text{Cac}(\mathbf{G}_n)$ is bounded above by r_n , so that the convergence of the proposition follows. \square

Corollary 2.10. *Let \mathbf{E} be a pointed geodesic compact metric space. Then $\text{Kac}(\mathbf{E})$ is a compact \mathbb{R} -tree.*

Proof. As a simple consequence of Proposition 7.5.5 in [5], we can find a sequence $(r_n)_{n \geq 0}$ of positive real numbers converging to 0 and a sequence $(\mathbf{G}_n)_{n \geq 0}$ of pointed graphs, such that the rescaled graphs $r_n \cdot \mathbf{G}_n$ converge to \mathbf{E} in the Gromov–Hausdorff sense. By Proposition 2.8, $r_n \cdot \text{Cac}(\mathbf{G}_n)$ converges to $\text{Kac}(\mathbf{E})$ in the Gromov–Hausdorff sense. Using the notation of the preceding proof, it also holds that $r_n \cdot \Lambda(\text{Cac}(\mathbf{G}_n))$ converges to $\text{Kac}(\mathbf{E})$. Proposition 2.2 then implies that $r_n \cdot \Lambda(\text{Cac}(\mathbf{G}_n))$ is a (compact) \mathbb{R} -tree. The desired result follows since the set of all compact \mathbb{R} -trees is known to be closed for the Gromov–Hausdorff topology (see e.g. [7], Lemma 2.1). \square

2.5. Another approach to the continuous cactus

In this section, we present an alternative definition of the continuous cactus, which gives a different perspective on the previous results, and in particular on Corollary 2.10. Let $\mathbf{E} = (E, d, \rho)$ be a pointed geodesic compact metric space, and for $r \geq 0$, let

$$\mathbf{B}(r) = \{x \in E: d(\rho, x) < r\}, \quad \bar{\mathbf{B}}(r) = \{x \in E: d(\rho, x) \leq r\},$$

be respectively the open and the closed ball of radius r centered at ρ . We let $\text{Kac}'(\mathbf{E})$ be the set of all subsets of E that are (non-empty) connected components of the closed set $\mathbf{B}(r)^c$, for some $r \geq 0$ (here, A^c denotes the complement of the set A). Note that all elements of $\text{Kac}'(\mathbf{E})$ are themselves closed subsets of E .

For every $C \in \text{Kac}'(\mathbf{E})$, we let

$$h(C) = d(\rho, C) = \inf\{d(\rho, x): x \in C\}.$$

Since E is path-connected, $h(C)$ is also the unique real $r \geq 0$ such that C is a connected component of $\mathbf{B}(r)^c$.

Note that $\text{Kac}'(\mathbf{E})$ is partially ordered by the relation

$$C \preceq C' \iff C' \subseteq C$$

and has a unique minimal element E . Every totally ordered subset of $\text{Kac}'(\mathbf{E})$ has a supremum, given by the intersection of all its elements. To see this, observe that if $(C_i)_{i \in I}$ is a totally ordered subset of $\text{Kac}'(\mathbf{E})$ then we can choose a sequence $(i_n)_{n \geq 1}$ taking values in I such that the sequence $(h(C_{i_n}))_{n \geq 1}$ is non-decreasing and converges to $r_{\max} := \sup\{h(C_i): i \in I\}$. Then the intersection

$$\bigcap_{n=1}^{\infty} C_{i_n}$$

is non-empty, closed and connected as the intersection of a decreasing sequence of non-empty closed connected sets in a compact space, and it easily follows that this intersection is a connected component of $\mathbf{B}(r_{\max})^c$ and coincides with the intersection of all C_i , $i \in I$. At this point, it is crucial that elements of $\text{Kac}'(\mathbf{E})$ are closed, and this is one of the reasons why one considers complements of *open* balls in the definition of $\text{Kac}'(\mathbf{E})$.

In particular, for every $C, C' \in \text{Kac}'(\mathbf{E})$, the infimum $C \wedge C'$ makes sense as the supremum of all $C'' \in \text{Kac}'(\mathbf{E})$ such that $C'' \preceq C$ and $C'' \preceq C'$, and $h(C \wedge C')$ is the maximal value of r such that C and C' are contained in the same connected component of $\mathbf{B}(r)^c$.

Moreover, if $C \in \text{Kac}'(\mathbf{E})$, the set $\{C' \in \text{Kac}'(\mathbf{E}) : C' \preceq C\}$ is isomorphic as an ordered set to the segment $[0, h(C)]$, because for every $t \in [0, h(C)]$ there is a unique $C' \in \text{Kac}'(\mathbf{E})$ with $h(C') = t$ and $C \subset C'$.

Finally, $h : \text{Kac}'(\mathbf{E}) \rightarrow \mathbb{R}_+$ is an increasing function, inducing a bijection from every segment of the partially ordered set $\text{Kac}'(\mathbf{E})$ to a real segment. It follows from general results (see Proposition 3.10 in [8]) that the set $\text{Kac}'(\mathbf{E})$ equipped with the distance

$$d_{\text{Kac}'}^{\mathbf{E}}(C, C') = h(C) + h(C') - 2h(C \wedge C')$$

is an \mathbb{R} -tree rooted at $E = \mathbf{B}(0)^c$. Note that $d_{\text{Kac}'}^{\mathbf{E}}(E, C) = h(C)$ for every $C \in \text{Kac}'(\mathbf{E})$.

Proposition 2.11. *The spaces $\text{Kac}'(\mathbf{E})$ and $\text{Kac}(\mathbf{E})$ are isometric pointed metric spaces.*

Proof. We consider the mapping from E to $\text{Kac}'(\mathbf{E})$, which maps x to the connected component C_x of $\mathbf{B}(d(\rho, x))^c$ containing x . This mapping is clearly onto: if $C \in \text{Kac}'(\mathbf{E})$, we have $C = C_x$ for any $x \in C$ such that $d(\rho, x) = d(\rho, C)$. Let us show that this mapping is an isometry from the pseudo-metric space $(E, d_{\text{Kac}}^{\mathbf{E}})$ onto $(\text{Kac}'(\mathbf{E}), d_{\text{Kac}'}^{\mathbf{E}})$.

Let $x, y \in E$ be given, and $\gamma : [0, 1] \rightarrow E$ be a path from x to y . Let t_0 be such that $d(\rho, \gamma(t_0)) \leq d(\rho, \gamma(t))$ for every $t \in [0, 1]$. Then the path γ lies in a single path-connected component of $\mathbf{B}(d(\rho, \gamma(t_0)))^c$, entailing that x and y are in the same connected component of this set. Consequently, $h(C_x \wedge C_y) \geq d(\rho, \gamma(t_0))$, and since obviously $h(C_x) = d(x, \rho)$,

$$d_{\text{Kac}'}^{\mathbf{E}}(C_x, C_y) \leq d(\rho, x) + d(\rho, y) - 2 \inf_{t \in [0, 1]} d(\rho, \gamma(t)).$$

Taking the infimum over all γ gives

$$d_{\text{Kac}'}^{\mathbf{E}}(C_x, C_y) \leq d_{\text{Kac}}^{\mathbf{E}}(x, y). \tag{4}$$

Let us verify that the reverse inequality also holds. If $h(C_x \wedge C_y) > 0$ and $\varepsilon \in (0, h(C_x \wedge C_y))$, the infimum $C_x \wedge C_y$ is contained in some connected component of $\overline{\mathbf{B}}(h(C_x \wedge C_y) - \varepsilon)^c$. Since the latter set is open, and E is a geodesic space, hence locally path-connected, we deduce that this connected component is in fact path-connected, and since it contains x and y , we can find a path γ from x to y that remains in $\overline{\mathbf{B}}(h(C_x \wedge C_y) - \varepsilon)^c$. This entails that

$$d_{\text{Kac}}^{\mathbf{E}}(x, y) \leq d_{\text{Kac}'}^{\mathbf{E}}(C_x, C_y) + \varepsilon,$$

and letting $\varepsilon \rightarrow 0$ yields the bound $d_{\text{Kac}'}^{\mathbf{E}}(C_x, C_y) \geq d_{\text{Kac}}^{\mathbf{E}}(x, y)$. The latter bound remains true when $h(C_x \wedge C_y) = 0$, since in that case $C_x \wedge C_y = E$ and $d_{\text{Kac}'}^{\mathbf{E}}(C_x, C_y) = h(C_x) + h(C_y) = d(\rho, x) + d(\rho, y)$.

From the preceding observations, we directly obtain that $x \mapsto C_x$ induces a quotient mapping from $\text{Kac}(\mathbf{E})$ onto $\text{Kac}'(\mathbf{E})$, which is an isometry and maps (the class of) ρ to E . \square

Remark 2.12. *The discrete cactus of a graph can be defined in an analogous way as above, using the notion of graph connectedness instead of connectedness in metric spaces.*

Let us return to Remark 2.5 about the existence, for given $x, y \in E$, of a minimizing path $\gamma : [0, 1] \rightarrow E$ going from x to y , such that

$$d_{\text{Kac}}^{\mathbf{E}}(x, y) = d(\rho, x) + d(\rho, y) - 2 \min_{0 \leq t \leq 1} d(\rho, \gamma(t)).$$

With the notation of the previous proof, it may happen that the closed set $C_x \wedge C_y$ is connected without being path-connected: Fig. 2 suggests an example of this phenomenon. In that event, if x and y cannot be connected by a continuous path that stays in $C_x \wedge C_y$, there exists no minimizing path.




Fig. 2. An example of a geodesic compact metric space E , such that the complement of the open ball of radius 1 centered at the distinguished point ρ is connected but not path-connected. Here E is a compact subset of \mathbb{R}^3 and is equipped with the intrinsic distance associated with the L^∞ -metric $\delta((x_1, x_2, x_3), (y_1, y_2, y_3)) = \sup\{|x_i - y_i|, i = 1, 2, 3\}$. For this distance, the sphere of radius 1 centered at ρ , which coincides with the complement of the open ball of radius 1, consists of the closure of the union of the bold lines at the top of the figure. It is not path-connected, by the same argument as the one showing that the closure of the graph of the function $x \mapsto \sin(1/x)$ on $(0, \infty)$ is a subset of the plane which is not path-connected.

3. The Brownian cactus

In this section, we define the Brownian cactus and we show that it is the continuous cactus associated with the (random) compact metric space called the Brownian map. We first recall some basic facts about the Brownian map.

We let $\mathbf{e} = (\mathbf{e}_t)_{0 \leq t \leq 1}$ be a Brownian excursion with duration 1. For our purposes it is crucial to view \mathbf{e} as the coding function for the random continuous tree known as the CRT. Precisely, we define a pseudo-distance $d_{\mathbf{e}}$ on $[0, 1]$ by setting for every $s, t \in [0, 1]$,

$$d_{\mathbf{e}}(s, t) = \mathbf{e}_s + \mathbf{e}_t - 2 \min_{s \wedge t \leq r \leq s \vee t} \mathbf{e}_r$$

and we set $s \sim_{\mathbf{e}} t$ iff $d_{\mathbf{e}}(s, t) = 0$. The CRT is defined as the quotient metric space $\mathcal{T}_{\mathbf{e}} := [0, 1] / \sim_{\mathbf{e}}$, and is equipped with the induced metric $d_{\mathbf{e}}$. Then $(\mathcal{T}_{\mathbf{e}}, d_{\mathbf{e}})$ is a random (compact) \mathbb{R} -tree. We write $p_{\mathbf{e}} : [0, 1] \rightarrow \mathcal{T}_{\mathbf{e}}$ for the canonical projection, and we define the mass measure (or volume measure) Vol on the CRT as the image of Lebesgue measure on $[0, 1]$ under $p_{\mathbf{e}}$. For every $a, b \in \mathcal{T}_{\mathbf{e}}$, we let $[[a, b]]$ be the range of the geodesic path from a to b in $\mathcal{T}_{\mathbf{e}}$: This is the line segment between a and b in the tree $\mathcal{T}_{\mathbf{e}}$. We will need the following simple fact, which is easily checked from the definition of $d_{\mathbf{e}}$. Let $a, b \in \mathcal{T}_{\mathbf{e}}$, and let $s, t \in [0, 1]$ be such that $p_{\mathbf{e}}(s) = a$ and $p_{\mathbf{e}}(t) = b$. Assume for definiteness that $s \leq t$. Then $[[a, b]]$ exactly consists of the points c that can be written as $c = p_{\mathbf{e}}(r)$, with $r \in [s, t]$ satisfying

$$\mathbf{e}_r = \max\left(\min_{u \in [s, r]} \mathbf{e}_u, \min_{u \in [r, t]} \mathbf{e}_u\right).$$

Conditionally given \mathbf{e} , we introduce the centered Gaussian process $(Z_t)_{0 \leq t \leq 1}$ with continuous sample paths such that

$$\text{cov}(Z_s, Z_t) = \min_{s \wedge t \leq r \leq s \vee t} \mathbf{e}_r.$$

It is easy to verify that a.s. for every $s, t \in [0, 1]$ the condition $s \sim_{\mathbf{e}} t$ implies that $Z_s = Z_t$. Therefore we may and will view Z as indexed by the CRT $\mathcal{T}_{\mathbf{e}}$. In fact, it is natural to interpret Z as Brownian motion indexed by the CRT. We will write indifferently $Z_a = Z_t$ when $a \in \mathcal{T}_{\mathbf{e}}$ and $t \in [0, 1]$ are such that $a = p_{\mathbf{e}}(t)$.

We set $\underline{Z} := \min_{t \in [0,1]} Z_t$. One can then prove [18] that a.s. there exists a unique $s_* \in [0, 1]$ such that $Z_{s_*} = \underline{Z}$. We let $a_* = p_e(s_*)$.

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

$$D^\circ(s, t) = Z_s + Z_t - 2 \max\left(\min_{r \in [s,t]} Z_r, \min_{r \in [t,s]} Z_r\right),$$

where we make the convention that when $s > t$, the notation $r \in [s, t]$ means $r \in [s, 1] \cup [0, t]$. We then define D° on $\mathcal{T}_e \times \mathcal{T}_e$ by setting for $a, b \in \mathcal{T}_e$,

$$D^\circ(a, b) = \min\{D^\circ(s, t) : s, t \in [0, 1], p_e(s) = a, p_e(t) = b\}.$$

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

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

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 such that $a_0 = a$ and $a_p = b$. It is not hard to verify that D is a pseudo-distance on \mathcal{T}_e , and we introduce the associated equivalence relation

$$a \approx b \quad \text{if and only if} \quad D(a, b) = 0.$$

The Brownian map is now defined as the quotient space

$$m_\infty := \mathcal{T}_e / \approx$$

which is equipped with the distance induced by D . We will view the Brownian map as a (random) pointed metric space with distinguished point $\rho_* = \Pi(a_*)$, where $\Pi : \mathcal{T}_e \rightarrow m_\infty$ is the canonical projection. We also let λ be the image of Vol under Π , and we interpret λ as the volume measure on m_∞ . For every $x \in m_\infty$, we set $Z_x = Z_a$, where $a \in \mathcal{T}_e$ is such that $\Pi(a) = x$ (this definition does not depend on the choice of a). It then easily follows from the definition of D that, for every $x \in m_\infty$,

$$D(\rho_*, x) = Z_x - \underline{Z}. \tag{5}$$

It is proved in [15,21] that the Brownian map is the limit in distribution, in the Gromov–Hausdorff sense, of rescaled uniformly distributed rooted $2p$ -angulations with n faces, for any integer $p \geq 2$ (the result in fact also holds for triangulations). By the argument in Remark 2.9, it follows that the metric space (m_∞, D) is a geodesic space a.s.

We now turn to the definition of the Brownian cactus.

Definition 3.1. *The Brownian cactus KAC is the random metric space defined as the quotient space of \mathcal{T}_e for the equivalence relation*

$$a \asymp b \quad \text{iff} \quad Z_a = Z_b = \min_{c \in \llbracket a, b \rrbracket} Z_c$$

and equipped with the distance induced by

$$d_{\text{KAC}}(a, b) = Z_a + Z_b - 2 \min_{c \in \llbracket a, b \rrbracket} Z_c \quad \text{for every } a, b \in \mathcal{T}_e.$$

We view KAC as a pointed metric space whose root is the equivalence class of a_* .

It is an easy matter to verify that d_{KAC} is a pseudo-distance on \mathcal{T}_e , and that \asymp is the associated equivalence relation. We write \mathbf{m}_∞ for the pointed metric space (m_∞, D, ρ_*) .

Proposition 3.2. *Almost surely, $\text{Kac}(\mathbf{m}_\infty)$ is isometric to KAC .*

Proof. We first need to identify the pseudo-distance $d_{\text{Kac}}^{\mathbf{m}_\infty}$ (see Section 2.2). Let $x, y \in m_\infty$ and choose $a, b \in \mathcal{T}_e$ such that $x = p_e(a)$ and $y = p_e(b)$. If $\gamma : [0, 1] \rightarrow m_\infty$ is a continuous path such that $\gamma(0) = x$ and $\gamma(1) = y$, Proposition 3.1 in [14] ensures that

$$\min_{0 \leq t \leq 1} Z_{\gamma(t)} \leq \min_{c \in \llbracket a, b \rrbracket} Z_c.$$

Using (5), it follows that

$$\min_{0 \leq t \leq 1} D(\rho_*, \gamma(t)) \leq \min_{c \in \llbracket a, b \rrbracket} (Z_c - \underline{Z}).$$

Since this holds for any continuous curve γ from x to y in m_∞ , we get from the definition of $d_{\text{Kac}}^{\mathbf{m}_\infty}$ that

$$d_{\text{Kac}}^{\mathbf{m}_\infty}(x, y) \geq (Z_a - \underline{Z}) + (Z_b - \underline{Z}) - 2 \min_{c \in \llbracket a, b \rrbracket} (Z_c - \underline{Z}) = d_{\text{KAC}}(a, b).$$

The corresponding upper bound is immediately obtained by letting γ be the image under Π of the (rescaled) geodesic path from a to b in the tree \mathcal{T}_e . Note that the resulting path from x to y in m_∞ is continuous because the projection Π is so. Summarizing, we have obtained that, for every $a, b \in \mathcal{T}_e$,

$$d_{\text{Kac}}^{\mathbf{m}_\infty}(\Pi(a), \Pi(b)) = d_{\text{KAC}}(a, b). \tag{6}$$

In particular, the property $a \asymp b$ holds if and only if $\Pi(a) \overset{\mathbf{m}_\infty}{\asymp} \Pi(b)$. Hence, the composition of the canonical projections from \mathcal{T}_e onto m_∞ and from m_∞ onto $\text{Kac}(\mathbf{m}_\infty)$ induces a one-to-one mapping from $\text{KAC} = \mathcal{T}_e / \asymp$ onto $\text{Kac}(\mathbf{m}_\infty)$. By (6) this mapping is an isometry, which completes the proof. \square

Remark 3.3. *From (6) and Proposition 2.4, we have*

$$D(\Pi(a), \Pi(b)) \geq d_{\text{KAC}}(a, b)$$

for every $a, b \in \mathcal{T}_e$ (cf. Corollary 3.2 in [14]).

As a corollary of the preceding proposition, the results of [15,21] and Proposition 2.8, we immediately get that the (suitably rescaled) discrete cactus associated with uniformly distributed rooted $2p$ -angulations, or triangulations, with n faces converges in distribution as $n \rightarrow \infty$ towards the Brownian cactus. We refrain from stating this corollary in a precise form, since we will get the same result for much more general random planar maps in the next section.

4. Convergence of cactuses associated with random planar maps

4.1. Planar maps and bijections with trees

We denote the set of all rooted and pointed planar maps by $\mathcal{M}_{r,p}$. As in [19], it is convenient for technical reasons to make the convention that $\mathcal{M}_{r,p}$ contains the ‘‘vertex map,’’ denoted by \dagger , which has no edge and only one vertex ‘‘bounding’’ a face of degree 0. With the exception of \dagger , a planar map in $\mathcal{M}_{r,p}$ has at least one edge. An element of $\mathcal{M}_{r,p}$ other than \dagger consists of a planar map m together with an oriented edge e (the root edge) and a distinguished vertex ρ . We write e_- and e_+ for the origin and the target of the root edge e . Note that we may have $e_- = e_+$ if e is a loop.

As previously, we denote the graph distance on the vertex set $V(m)$ of m by d_{gr}^m . We say that the rooted and pointed planar map (m, e, ρ) is positive, respectively negative, respectively null if $d_{\text{gr}}^m(\rho, e_+) = d_{\text{gr}}^m(\rho, e_-) + 1$, resp. $d_{\text{gr}}^m(\rho, e_+) = d_{\text{gr}}^m(\rho, e_-) - 1$, resp. $d_{\text{gr}}^m(\rho, e_+) = d_{\text{gr}}^m(\rho, e_-)$. We make the convention that the vertex map \dagger is positive. We write $\mathcal{M}_{r,p}^+$, resp. $\mathcal{M}_{r,p}^-$, resp. $\mathcal{M}_{r,p}^0$ for the set of all positive, resp. negative, resp. null, rooted and pointed planar

maps. Reversing the orientation of the root edge yields an obvious bijection between the sets $\mathcal{M}_{r,p}^+$ and $\mathcal{M}_{r,p}^-$, and for this reason we will mainly discuss $\mathcal{M}_{r,p}^+$ and $\mathcal{M}_{r,p}^0$ in what follows.

We will make use of the Bouttier–Di Francesco–Guitter bijection [3] between $\mathcal{M}_{r,p}^+ \cup \mathcal{M}_{r,p}^0$ and a certain set of multitype labeled trees called mobiles. In order to describe this bijection, we use the standard formalism for plane trees, as found in Section 1.1 of [12] for instance. In this formalism, vertices are elements of the set

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

of all finite sequences of positive integers, including the empty sequence \emptyset that serves as the root vertex of the tree. A plane tree τ is a finite subset of \mathcal{U} that satisfies the following three conditions:

1. $\emptyset \in \tau$.
2. For every $u = (i_1, \dots, i_k) \in \tau \setminus \{\emptyset\}$, the sequence (i_1, \dots, i_{k-1}) (the “parent” of u) also belongs to τ .
3. For every $u = (i_1, \dots, i_k) \in \tau$, there exists an integer $k_u(\tau) \geq 0$ (the “number of children” of u) such that the vertex (i_1, \dots, i_k, j) belongs to τ if and only if $1 \leq j \leq k_u(\tau)$.

The generation of $u = (i_1, \dots, i_k)$ is denoted by $|u| = k$. The notions of an ancestor and a descendant in the tree τ are defined in an obvious way.

We will be interested in four-type plane trees, meaning that each vertex is assigned a type which can be 1, 2, 3 or 4.

We next introduce mobiles following the presentation in [19], with a few minor modifications. We consider a four-type plane tree τ satisfying the following properties:

- (i) The root vertex \emptyset is of type 1 or of type 2.
- (ii) The children of any vertex of type 1 are of type 3.
- (iii) Each individual of type 2 and which is not the root vertex of the tree has exactly one child of type 4 and no other child. If the root vertex is of type 2, it has exactly two children, both of type 4.
- (iv) The children of individuals of type 3 or 4 can only be of type 1 or 2.

Let $\tau_{(1,2)}$ be the set of all vertices of τ at even generation (these are exactly the vertices of type 1 or 2). An admissible labeling of τ is a collection of integer labels $(\ell_u)_{u \in \tau_{(1,2)}}$ assigned to the vertices of type 1 or 2, such that the following properties hold:

- (a) $\ell_{\emptyset} = 0$.
- (b) Let u be a vertex of type 3 or 4, let $u_{(1)}, \dots, u_{(k)}$ be the children of u (in lexicographical order) and let $u_{(0)}$ be the parent of u . Then, for every $i = 0, 1, \dots, k$,

$$\ell_{u_{(i+1)}} \geq \ell_{u_{(i)}} - 1$$

with the convention $u_{(k+1)} = u_{(0)}$. Moreover, for every $i = 0, 1, \dots, k$ such that $u_{(i+1)}$ is of type 2, we have

$$\ell_{u_{(i+1)}} \geq \ell_{u_{(i)}}.$$

By definition, a *mobile* is a pair $(\tau, (\ell_u)_{u \in \tau_{(1,2)}})$ consisting of a four-type plane tree satisfying the preceding conditions (i)–(iv), and an admissible labeling of τ . We let \mathbb{T}_+ be the set of all mobiles such that the root vertex of τ is of type 1. We also let \mathbb{T}_0 be the set of all mobiles such that the root vertex is of type 2.

Remark 4.1. Our definition of admissible labelings is slightly different from the ones that are used in [19] or [22]. To recover the definitions of [19] or [22], just subtract 1 from the label of each vertex of type 2. Because of this difference, our construction of the bijections between maps and trees will look slightly different from the ones in [19] or [22].

The Bouttier–Di Francesco–Guitter construction provides bijections between the set \mathbb{T}_+ and the set $\mathcal{M}_{r,p}^+$ on one hand, between the set \mathbb{T}_0 and the set $\mathcal{M}_{r,p}^0$ on the other hand. Let us describe this construction in the first case.

We start from a mobile $(\tau, (\ell_u)_{u \in \tau_{(1,2)}}) \in \mathbb{T}_+$. In the case when $\tau = \{\emptyset\}$, we decide by convention that the associated planar map is the vertex map \dagger . Otherwise, let $p \geq 1$ be the number of edges of τ ($p = \#\tau - 1$). The contour sequence

of τ is the sequence v_0, v_1, \dots, v_{2p} of vertices of τ defined inductively as follows. First $v_0 = \emptyset$. Then, for every $i \in \{0, 1, \dots, 2p - 1\}$, v_{i+1} is either the first child of v_i that has not yet appeared among v_0, v_1, \dots, v_i , or if there is no such child, the parent of v_i . It is easy to see that this definition makes sense and $v_{2p} = \emptyset$. Moreover all vertices of τ appear in the sequence v_0, v_1, \dots, v_{2p} , and more precisely the number of occurrences of a vertex u of τ is equal to the multiplicity of u in τ . In fact, each index i such that $v_i = u$ corresponds to one corner of the vertex u in the tree τ : We will abusively call it the corner v_i . We also introduce the modified contour sequence of τ as the sequence u_0, u_1, \dots, u_p defined by

$$u_i = v_{2i} \quad \forall i = 0, 1, \dots, p.$$

By construction, the vertices appearing in the modified contour sequence are exactly the vertices of $\tau_{(1,2)}$. We extend the modified contour sequence periodically by setting $u_{p+i} = u_i$ for $i = 1, \dots, p$. Note that the properties of labels entail $\ell_{u_{i+1}} \geq \ell_{u_i} - 1$ for $i = 0, 1, \dots, 2p - 1$.

To construct the edges of the rooted and pointed planar map (m, e, ρ) associated with the mobile $(\tau, (\ell_u)_{u \in \tau_{(1,2)}}) \in \mathbb{T}_+$ we proceed as follows. We first embed the tree τ in the plane in a way consistent with the planar order. We then add an extra vertex of type 1, which we call ρ . Then, for every $i = 0, 1, \dots, p - 1$:

(i) If

$$\ell_{u_i} = \min_{0 \leq k \leq p} \ell_{u_k}$$

we draw an edge between the corner u_i and ρ .

(ii) If

$$\ell_{u_i} > \min_{0 \leq k \leq p} \ell_{u_k}$$

we draw an edge between the corner u_i and the corner u_j , where $j = \min\{k \in \{i + 1, \dots, i + p - 1\} : \ell_{u_k} = \ell_{u_i} - 1\}$. Because of property (b) of the labeling, the vertex u_j must be of type 1.

The construction can be made in such a way that edges do not intersect, and do not intersect the edges of the tree τ . Furthermore each face of the resulting planar map contains exactly one vertex of type 3 or 4, and both the parent and the children of this vertex are incident to this face. See Fig. 3 for an example.

The resulting planar map is bipartite with vertices either of type 1 or of type 2. Furthermore, the fact that in the tree τ each vertex of type 2 has exactly one child, and the labeling rules imply that each vertex of type 2 is incident to exactly two edges of the map, which connect it to two vertices of type 1, which may be the same (these vertices of type 1 will be said to be associated with the vertex of type 2 we are considering). Each of these edges corresponds in the preceding construction to one of the two corners of the vertex of type 2 that we consider. To complete the construction, we just erase all vertices of type 2 and for each of these we merge its two incident edges into a single edge connecting the two associated vertices of type 1. In this way we get a (non-bipartite in general) planar map m . Finally we decide that the root edge e of the map is the first edge drawn in the construction, oriented in such a way that $e_+ = \emptyset$, and we let the distinguished vertex of the map be the vertex ρ . Note that vertices of the map m that are different from the distinguished vertex ρ are exactly the vertices of type 1 in the tree τ . In other words, the vertex set $V(m)$ is identified with the set $\tau_{(1)} \cup \{\rho\}$, where $\tau_{(1)}$ denotes the set of all vertices of τ of type 1.

The mapping $(\tau, (\ell_u)_{u \in \tau_{(1,2)}}) \longrightarrow (m, e, \rho)$ that we have just described is indeed a bijection from \mathbb{T}_+ onto $\mathcal{M}_{r,p}^+$. We can construct a similar bijection from \mathbb{T}_+ onto $\mathcal{M}_{r,p}^-$ by the same construction, with the minor modification that we orient the root edge in such a way that $e_- = \emptyset$.

Furthermore we can also adapt the preceding construction in order to get a bijection from \mathbb{T}_0 onto $\mathcal{M}_{r,p}^0$. The construction of edges of the map proceeds in the same way, but the root edge is now obtained as the edge resulting of the merging of the two edges incident to \emptyset (recall that for a tree in \mathbb{T}_0 the root \emptyset is a vertex of type 2 that has exactly two children, hence also two corners). The orientation of the root edge is chosen according to some convention: For instance, one may decide that the ‘‘half-edge’’ coming from the first corner of \emptyset corresponds to the origin of the root edge.




Fig. 3. A mobile $(\tau, (\ell_u)_{u \in \tau_{(1,2)}})$ in \mathbb{T}_+ and its image m under the BDG bijection. Vertices of type 1 are represented by big circles, vertices of type 2 by lozanges, vertices of type 3 by small circles and vertices of type 4 by small black disks. The edges of the tree τ are represented by thin lines, and the edges of the planar map m by thick curves. In order to get the planar map m one needs to erase the vertices of type 2 and, for each of these vertices, to merge its two incident edges into a single one. The root edge is at the bottom left.

In all three cases, distances in the planar map m satisfy the following key property: For every vertex $u \in \tau_{(1)}$, we have

$$d_{\text{gr}}^m(\rho, u) = \ell_u - \min \ell + 1, \tag{7}$$

where $\min \ell$ denotes the minimal label on the tree τ . In the left-hand side u is viewed as a vertex of the map m , in agreement with the preceding construction.

The three bijections we have described are called the BDG bijections. In the remaining part of this section, we fix a mobile $(\tau, (\ell_u)_{u \in \tau_{(1,2)}})$ belonging to \mathbb{T}_+ (or to \mathbb{T}_0) and its image (m, e, ρ) under the relevant BDG bijection.

Remark 4.2. We could have defined the BDG bijections without distinguishing between types 3 and 4. However, this distinction will be important in the next section when we consider random planar maps and the associated (random) trees. We will see that these random trees are Galton–Watson trees with a different offspring distribution for vertices of type 3 than for vertices of type 4.

If $u, v \in \tau_{(1,2)}$, we denote by $\llbracket u, v \rrbracket$ the set of all vertices of type 1 or 2 that lie on the geodesic path from u to v in the tree τ .

Proposition 4.3. For every $u, v \in V(m) \setminus \{\rho\} = \tau_{(1)}$, and every path $\gamma = (\gamma(0), \gamma(1), \dots, \gamma(k))$ in m such that $\gamma(0) = u$ and $\gamma(k) = v$, we have

$$\min_{0 \leq i \leq k} d_{\text{gr}}^m(\rho, \gamma(i)) \leq \min_{w \in \llbracket u, v \rrbracket} \ell_w - \min \ell + 1.$$

Proof. We may assume that the path γ does not visit ρ , since otherwise the result is trivial. Using (7), the statement reduces to

$$\min_{0 \leq i \leq k} \ell_{\gamma(i)} \leq \min_{w \in \llbracket u, v \rrbracket} \ell_w.$$

So we fix $w \in \llbracket u, v \rrbracket$ and we verify that $\ell_{\gamma(i)} \leq \ell_w$ for some $i \in \{0, 1, \dots, k\}$. We may assume that $w \neq u$ and $w \neq v$. The removal of the vertex w (and of the edges incident to w) disconnects the tree τ in several connected components. Write C for the connected component containing v , and note that this component does not contain u . Then let $j \geq 1$ be the first integer such that $\gamma(j)$ belongs to C . Thus $\gamma(j-1) \notin C$, $\gamma(j) \in C$ and the vertices $\gamma(j-1)$ and $\gamma(j)$ are linked by an edge of the map m . From (7), we have $|\ell_{\gamma(j)} - \ell_{\gamma(j-1)}| \leq 1$. Now we use the fact that the edge between $\gamma(j-1)$ and $\gamma(j)$ is produced by the BDG bijection. Suppose first that $\gamma(j-1)$ and $\gamma(j)$ have a different label. In that case, noting that the modified contour sequence must visit w between any visit of $\gamma(j-1)$ and any visit of $\gamma(j)$, we easily get that $\min\{\ell_{\gamma(j)}, \ell_{\gamma(j-1)}\} \leq \ell_w$ (otherwise our construction could not produce an edge from $\gamma(j-1)$ to $\gamma(j)$). A similar argument applies to the case when $\gamma(j-1)$ and $\gamma(j)$ have the same label. In that case, the edge between $\gamma(j-1)$ and $\gamma(j)$ must come from the merging of two edges originating from a vertex of τ of type 2. This vertex of type 2 has to belong to the set $\llbracket \gamma(j-1), \gamma(j) \rrbracket$ (which contains w), because otherwise the two associated vertices of type 1 could not be $\gamma(j-1)$ and $\gamma(j)$. It again follows from our construction that we must have $\min\{\ell_{\gamma(j)}, \ell_{\gamma(j-1)}\} \leq \ell_w$. This completes the proof. \square

In the next corollary, we write \mathbf{m} for the graph associated with the map m (in the sense of Section 2.1), which is pointed at the distinguished vertex ρ . The notation $d_{\text{Cac}}^{\mathbf{m}}$ then refers to the cactus distance for this pointed graph.

Corollary 4.4. *Suppose that the degree of all faces of m is bounded above by $D \geq 1$. Then, for every $u, v \in V(m) \setminus \{\rho\}$, we have*

$$\left| d_{\text{Cac}}^{\mathbf{m}}(u, v) - \left(\ell_u + \ell_v - 2 \min_{w \in \llbracket u, v \rrbracket} \ell_w \right) \right| \leq 2D + 2.$$

Proof. From the definition of the cactus distance $d_{\text{Cac}}^{\mathbf{m}}$ and the preceding proposition, we immediately get the lower bound

$$d_{\text{Cac}}^{\mathbf{m}}(u, v) \geq d_{\text{gr}}^m(\rho, u) + d_{\text{gr}}^m(\rho, v) - 2 \left(\min_{w \in \llbracket u, v \rrbracket} \ell_w - \min \ell + 1 \right) = \ell_u + \ell_v - 2 \min_{w \in \llbracket u, v \rrbracket} \ell_w,$$

by (7). In order to get a corresponding upper bound, let $\eta(0) = u, \eta(1), \dots, \eta(k) = v$ be the vertices of type 1 or 2 belonging to the geodesic path from u to v in the tree τ , enumerated in their order of appearance on this path. Set $\tilde{\eta}(i) = \eta(i)$ if $\eta(i)$ is of type 1, and if $\eta(i)$ is of type 2, let $\tilde{\eta}(i)$ be one of the two (possibly equal) vertices of type 1 that are associated with $\eta(i)$ in the BDG bijection. Then the properties of the BDG bijection ensure that, for every $i = 0, 1, \dots, k-1$, the two vertices $\eta(i)$ and $\eta(i+1)$ lie on the boundary of the same face of m (the point is that, in the BDG construction, edges of the map m are drawn in such a way that they do not cross edges of the tree τ). From our assumption we have thus $d_{\text{gr}}^m(\tilde{\eta}(i), \tilde{\eta}(i+1)) \leq D$ for every $i = 0, 1, \dots, k-1$. Hence, we can find a path γ in m starting from u and ending at v , such that

$$\min_j d_{\text{gr}}^m(\rho, \gamma(j)) \geq \min_{0 \leq i \leq k} d_{\text{gr}}^m(\rho, \tilde{\eta}(i)) - D = \min_{0 \leq i \leq k} \ell_{\tilde{\eta}(i)} - \min \ell + 1 - D \geq \min_{0 \leq i \leq k} \ell_{\eta(i)} - \min \ell - D.$$

It follows that

$$d_{\text{Cac}}^{\mathbf{m}}(u, v) \leq d_{\text{gr}}^m(\rho, u) + d_{\text{gr}}^m(\rho, v) - 2 \left(\min_{w \in \llbracket u, v \rrbracket} \ell_w - \min \ell - D \right) = \ell_u + \ell_v - 2 \min_{w \in \llbracket u, v \rrbracket} \ell_w + 2D + 2.$$

This completes the proof. \square

4.2. Random planar maps

Following [18] and [19], we now discuss Boltzmann distributions on the space $\mathcal{M}_{r,p}$. We consider a sequence $\mathbf{q} = (q_1, q_2, \dots)$ of non-negative real numbers. We assume that the sequence \mathbf{q} has finite support ($q_k = 0$ for all sufficiently large k), and is such that $q_k > 0$ for some $k \geq 3$. We will then split our study according to the following two possibilities:

(A1) There exists an odd integer k such that $q_k > 0$.

(A2) The sequence \mathbf{q} is supported on even integers.

If $m \in \mathcal{M}_{r,p}$, we define

$$W_{\mathbf{q}}(m) = \prod_{f \in F(m)} q_{\deg(f)},$$

where $F(m)$ stands for the set of all faces of m and $\deg(f)$ is the degree of the face f . In the case when $m = \dagger$, we make the convention that $q_0 = 1$ and thus $W_{\mathbf{q}}(\dagger) = 1$.

By multiplying the sequence \mathbf{q} by a suitable positive constant, we may assume that this sequence is regular critical in the sense of [19], Definition 1, under Assumption (A1) or of [17], Definition 1, under Assumption (A2). We refer the reader to the Appendix below for details. In particular, the measure $W_{\mathbf{q}}$ is then finite, and we can define a probability measure $P_{\mathbf{q}}$ on $\mathcal{M}_{r,p}$ by setting

$$P_{\mathbf{q}} = Z_{\mathbf{q}}^{-1} W_{\mathbf{q}},$$

where $Z_{\mathbf{q}} = W_{\mathbf{q}}(\mathcal{M}_{r,p})$.

For every integer n such that $W_{\mathbf{q}}(\#V(m) = n) > 0$, we consider a random planar map M_n distributed according to the conditional measure

$$\frac{P_{\mathbf{q}}(\cdot \cap \{\#V(m) = n\})}{P_{\mathbf{q}}(\#V(m) = n)}.$$

Throughout the remaining part of Section 4, we restrict our attention to values of n such that $W_{\mathbf{q}}(\#V(m) = n) > 0$, so that M_n is well defined. We write ρ_n for the distinguished vertex of M_n .

We now state the main result of this section. In this result, \mathbf{M}_n stands for the graph (pointed at ρ_n) associated with M_n , as explained at the end of Section 2.1.

Theorem 4.5. *There exists a positive constant $B_{\mathbf{q}}$ such that*

$$B_{\mathbf{q}} n^{-1/4} \cdot \text{Cac}(\mathbf{M}_n) \xrightarrow[n \rightarrow \infty]{(d)} \text{KAC}$$

in the Gromov–Hausdorff sense.

The proof of Theorem 4.5 relies on the asymptotic study of the random trees associated with planar maps distributed under Boltzmann distributions via the BDG bijection. The distribution of these random trees was identified in [17] (in the bipartite case) and in [19]. We set

$$Z_{\mathbf{q}}^+ = W_{\mathbf{q}}(\mathcal{M}_{r,p}^+) \geq 1, \quad Z_{\mathbf{q}}^0 = W_{\mathbf{q}}(\mathcal{M}_{r,p}^0).$$

Note that, under Assumption (A2), $W_{\mathbf{q}}$ is supported on bipartite maps and thus $Z_{\mathbf{q}}^0 = 0$. We also set

$$P_{\mathbf{q}}^+ = P_{\mathbf{q}}(\cdot | \mathcal{M}_{r,p}^+), \quad P_{\mathbf{q}}^- = P_{\mathbf{q}}(\cdot | \mathcal{M}_{r,p}^-), \quad P_{\mathbf{q}}^0 = P_{\mathbf{q}}(\cdot | \mathcal{M}_{r,p}^0).$$

Note that the definition of $P_{\mathbf{q}}^0$ only makes sense under Assumption (A1).

The next proposition gives the distribution of the tree associated with a random planar map distributed according to $P_{\mathbf{q}}^+$. Before stating this proposition, let us recall that the notion of a four-type Galton–Watson tree is defined analogously to the case of a single type. The distribution of such a random tree is determined by the type of the ancestor, and four offspring distributions v_i , $i = 1, 2, 3, 4$, which are probability distributions on \mathbb{Z}_+^4 ; for every $i = 1, 2, 3, 4$, v_i corresponds to the law of the number of children (having each of the four possible types) of an individual of type i ; furthermore, given the numbers of children of each type of an individual, these children are ordered in the tree with the same probability for each possible ordering. See [19], Section 2.2.1, for more details, noting that we consider only the case of “uniform ordering” in the terminology of [19].

Proposition 4.6. *Suppose that M^+ is a random planar map distributed according to $P_{\mathbf{q}}^+$, and let $(\theta, (\mathcal{L}_u)_{u \in \theta(1,2)})$ be the four-type labeled tree associated with M^+ via the BDG bijection between \mathbb{T}_+ and $\mathcal{M}_{r,p}^+$. Then the distribution of $(\theta, (\mathcal{L}_u)_{u \in \theta(1,2)})$ is characterized by the following properties:*

(i) *The random tree θ is a four-type Galton–Watson tree, such that the root \emptyset has type 1 and the offspring distributions ν_1, \dots, ν_4 are determined as follows:*

- ν_1 is supported on $\{0\} \times \{0\} \times \mathbb{Z}_+ \times \{0\}$, and for every $k \geq 0$,

$$\nu_1(0, 0, k, 0) = \frac{1}{Z_{\mathbf{q}}^+} \left(1 - \frac{1}{Z_{\mathbf{q}}^+} \right)^k.$$

- $\nu_2(0, 0, 0, 1) = 1$.
- ν_3 and ν_4 are supported on $\mathbb{Z}_+ \times \mathbb{Z}_+ \times \{0\} \times \{0\}$, and for every integers $k, k' \geq 0$,

$$\nu_3(k, k', 0, 0) = c_{\mathbf{q}} (Z_{\mathbf{q}}^+)^k (Z_{\mathbf{q}}^0)^{k'/2} \binom{2k+k'+1}{k+1} \binom{k+k'}{k} q_{2+2k+k'},$$

$$\nu_4(k, k', 0, 0) = c'_{\mathbf{q}} (Z_{\mathbf{q}}^+)^k (Z_{\mathbf{q}}^0)^{k'/2} \binom{2k+k'}{k} \binom{k+k'}{k} q_{1+2k+k'},$$

where $c_{\mathbf{q}}$ and $c'_{\mathbf{q}}$ are the appropriate normalizing constants.

(ii) *Conditionally given θ , $(\mathcal{L}_u)_{u \in \theta(1,2)}$ is uniformly distributed over all admissible labelings.*

Remark 4.7. *The definition of ν_4 does not make sense under Assumption (A2) (because $Z_{\mathbf{q}}^0 = 0$ in that case, $\nu_4(k, k', 0, 0)$ can be non-zero only if $k' = 0$, but then $q_{1+2k+k'} = 0$). This is however irrelevant since under Assumption (A2) the property $Z_{\mathbf{q}}^0 = 0$ entails that ν_3 is supported on $\mathbb{Z}_+ \times \{0\} \times \{0\} \times \{0\}$, and thus the Galton–Watson tree will have no vertices of type 2 or 4.*

We refer to [19], Proposition 3, for the proof of Proposition 4.6 under Assumption (A1) and to [17], Proposition 7, for the case of Assumption (A2). In fact, [19] assumes that $q_k > 0$ for some odd integer $k \geq 3$, but the results in that paper do cover the situation considered in the present work.

In the next two subsections, we prove Theorem 4.5 under Assumption (A1). The case when Assumption (A2) holds is much easier and will be treated briefly in Section 4.5.

4.3. The shuffling operation

As already mentioned, we suppose in this section that Assumption (A1) holds. We consider the random four-type labeled tree $(\theta, (\mathcal{L}_v)_{v \in \theta(1,2)})$ associated with the planar map M^+ via the BDG bijection, as in Proposition 4.6.

Our goal is to investigate the asymptotic behavior, when n tends to ∞ , of the labeled tree $(\theta, (\mathcal{L}_v)_{v \in \theta(1,2)})$ conditioned to have $n - 1$ vertices of type 1 (this corresponds to conditioning M^+ on the event $\{\#V(M^+) = n\}$). As already observed in [19], a difficulty arises from the fact that the label displacements along the tree are not centered, and so the results of [20] cannot be applied immediately. To overcome this difficulty, we will use an idea of [19], which consists in introducing a “shuffled” version of the tree θ . In order to explain this, we need to introduce some notation.

Let τ be a plane tree and $u = (i_1, \dots, i_p) \in \tau$. The tree τ shifted at u is defined by

$$T_u \tau := \{v = (j_1, \dots, j_\ell) : (i_1, \dots, i_p, j_1, \dots, j_\ell) \in \tau\}.$$

Let $k = k_u(\tau)$ be the number of children of u in τ , and, for every $1 \leq i \leq k$, write $u_{(i)}$ for the i th child of u . The tree τ reversed at vertex u is the new tree τ^* characterized by the properties:

- Vertices of τ^* which are not descendants of u are the same as vertices of τ which are not descendants of u .
- $u \in \tau^*$ and $k_u(\tau^*) = k_u(\tau) = k$.
- For every $1 \leq i \leq k$, $T_{u_{(i)}} \tau^* = T_{u_{(k+1-i)}} \tau$.

Our (random) shuffling operation will consist in reversing the tree τ at every vertex of τ at an odd generation, with probability $1/2$ for every such vertex. We now give a more formal description, which will be needed in our applications. We keep on considering a (deterministic) plane tree τ . Let \mathcal{U}^o stand for the set of all $u \in \mathcal{U}$ such that $|u|$ is odd. We consider a collection $(\varepsilon_u)_{u \in \mathcal{U}^o}$ of independent Bernoulli variables with parameter $1/2$. We then define a (random) mapping $\sigma : \tau \rightarrow \mathcal{U}$ by setting, if $u = (i_1, i_2, \dots, i_p)$,

$$\sigma(u) = (j_1, j_2, \dots, j_p),$$

where, for every $1 \leq \ell \leq p$,

- if ℓ is odd, $j_\ell = i_\ell$,
- if ℓ is even,

$$j_\ell = \begin{cases} i_\ell & \text{if } \varepsilon_{(i_1, \dots, i_{\ell-1})} = 0, \\ k_{(i_1, \dots, i_{\ell-1})}(\tau) + 1 - i_\ell & \text{if } \varepsilon_{(i_1, \dots, i_{\ell-1})} = 1. \end{cases}$$

Then $\tilde{\tau} = \{\sigma(u) : u \in \tau\}$ is a (random) plane tree, called the tree derived from τ by the shuffling operation. If τ is a four-type tree, we also view $\tilde{\tau}$ as a four-type tree by assigning to the vertex $\sigma(u)$ of $\tilde{\tau}$ the type of the vertex u in τ .

For our purposes it is very important to note that the bijection $\sigma : \tau \rightarrow \tilde{\tau}$ preserves the genealogical structure, in the sense that u is an ancestor of v in τ if and only if $\sigma(u)$ is an ancestor of $\sigma(v)$ in $\tilde{\tau}$. Consequently, if u and v are any two vertices of $\tau_{(1,2)}$, $\llbracket \sigma(u), \sigma(v) \rrbracket$ is the image under σ of the set $\llbracket u, v \rrbracket$.

We can apply this shuffling operation to the random tree θ (of course we assume that the collection $(\varepsilon_u)_{u \in \mathcal{U}^o}$ is independent of $(\theta, (\mathcal{L}_v)_{v \in \theta_{(1,2)}})$). We write $\tilde{\theta}$ for the four-type tree derived from θ by the shuffling operation and we use the same notation σ as above for the “shuffling bijection” from θ onto $\tilde{\theta}$. We assign labels to the vertices of $\tilde{\theta}_{(1,2)}$ by setting for every $u \in \theta_{(1,2)}$,

$$\tilde{\mathcal{L}}_{\sigma(u)} = \mathcal{L}_u.$$

Note that the random tree $\tilde{\theta}$ has the same distribution as θ , and is therefore a four-type Galton–Watson tree as described in Proposition 4.6. On the other hand, the labeled trees $(\theta, (\mathcal{L}_v)_{v \in \theta_{(1,2)}})$ and $(\tilde{\theta}, (\tilde{\mathcal{L}}_v)_{v \in \tilde{\theta}_{(1,2)}})$ have a different distribution because the admissibility property of labels is not preserved under the shuffling operation. We can still describe the distribution of the labels in the shuffled tree in a simple way. To this end, write $\text{tp}(u)$ for the type of a vertex u . Then conditionally on $\tilde{\theta}$, for every vertex u of $\tilde{\theta}$ such that $|u|$ is odd, if $u_{(1)}, \dots, u_{(k)}$ are the children of u in lexicographical order, and if $u_{(0)}$ is the parent of u , the vector of label increments

$$(\tilde{\mathcal{L}}_{u_{(1)}} - \tilde{\mathcal{L}}_{u_{(0)}}, \dots, \tilde{\mathcal{L}}_{u_{(k)}} - \tilde{\mathcal{L}}_{u_{(0)}})$$

is with probability $1/2$ uniformly distributed over the set

$$\mathbb{A} := \{(i_1, \dots, i_k) \in \mathbb{Z}^k : i_{j+1} \geq i_j - \mathbf{1}_{\{\text{tp}(u_{(j+1)})=1\}} \text{ for all } 0 \leq j \leq k\},$$

and with probability $1/2$ uniformly distributed over the set

$$\mathbb{A}' := \{(i_1, \dots, i_k) \in \mathbb{Z}^k : i_j \geq i_{j+1} - \mathbf{1}_{\{\text{tp}(u_{(j)})=1\}} \text{ for all } 0 \leq j \leq k\}.$$

In the definition of both \mathbb{A} and \mathbb{A}' we make the convention that $i_0 = i_{k+1} = 0$ and $u_{(k+1)} = u_{(0)}$. Furthermore the vectors of label increments are independent (still conditionally on $\tilde{\theta}$) when u varies over vertices of $\tilde{\theta}$ at odd generations.

The preceding description of the distribution of labels in the shuffled tree is easy to establish. Note that the set \mathbb{A} corresponds to the admissibility property of labels, whereas \mathbb{A}' corresponds to a “reversed” version of this property.

For every $u \in \tilde{\theta}_{(1,2)}$, set

$$\tilde{\mathcal{L}}'_u = \tilde{\mathcal{L}}_u - \frac{1}{2} \mathbf{1}_{\{\text{tp}(u)=2\}}.$$

If we replace $\tilde{\mathcal{L}}_u$ by $\tilde{\mathcal{L}}'_u$, then the vectors of label increments in $\tilde{\theta}$ become centered. This follows from elementary arguments: See [19], Lemma 2, for a detailed proof. As in [19] or in [22], the fact that the label increments are centered allows us to use the asymptotic results of [20], noting that these results will apply to $\tilde{\mathcal{L}}_u$ as well as to $\tilde{\mathcal{L}}'_u$ since the additional term $\frac{1}{2}\mathbf{1}_{\{\text{tp}(u)=2\}}$ obviously plays no role in the scaling limit. Before we state the relevant result, we need to introduce some notation.

For $n \geq 2$, let $(\tilde{\theta}^n, (\tilde{\mathcal{L}}_v^n)_{v \in \tilde{\theta}^n_{(1,2)}})$ be distributed as the labeled tree $(\tilde{\theta}, (\tilde{\mathcal{L}}_v)_{v \in \tilde{\theta}_{(1,2)}})$ conditioned on the event $\{\#\tilde{\theta}_{(1)} = n - 1\}$ (recall that we restrict our attention to values of n such that the latter event has positive probability). Let $p_n = \#\tilde{\theta}^n - 1$ and let $u_0^n = \emptyset, u_1^n, \dots, u_{p_n}^n = \emptyset$ be the modified contour sequence of $\tilde{\theta}^n$. The contour process $C^n = (C_i^n)_{0 \leq i \leq p_n}$ is defined by

$$C_i^n = |u_i^n|$$

and the label process $V^n = (V_i^n)_{0 \leq i \leq p_n}$ by

$$V_i^n = \tilde{\mathcal{L}}_{u_i^n}^n.$$

We extend the definition of both processes C^n and V^n to the real interval $[0, p_n]$ by linear interpolation.

Recall the notation (\mathbf{e}, Z) from Section 3.

Proposition 4.8. *There exist two positive constants A_q and B_q such that*

$$\left(A_q \frac{C^n(p_n s)}{n^{1/2}}, B_q \frac{V^n(p_n s)}{n^{1/4}} \right)_{0 \leq s \leq 1} \xrightarrow[n \rightarrow \infty]{(d)} (\mathbf{e}_s, Z_s)_{0 \leq s \leq 1} \tag{8}$$

in the sense of weak convergence of the distributions on the space $C([0, 1], \mathbb{R}^2)$.

This follows from the more general results proved in [20] for spatial multitype Galton–Watson trees. One should note that the results of [20] are given for variants of the contour process and the label process (in particular the contour process is replaced by the so-called height process of the tree). However simple arguments show that the convergence in the proposition can be deduced from the ones in [20]: See in particular Section 1.6 of [12] for a detailed explanation of why convergence results for the height process imply similar results for the contour process. Proposition 4.8 is also equivalent to Theorem 3.1 in [22], where the contour and label processes are defined in a slightly different way.

4.4. *Proof of Theorem 4.5 under Assumption (A1)*

We keep assuming that Assumption (A1) holds. Let M_n^+ be distributed according to the probability measure $P_q^+(\cdot | \#V(m) = n)$, or equivalently as M^+ conditionally on the event $\{\#V(M^+) = n\}$. As above, ρ_n stands for the distinguished point of M_n^+ , and we will write \mathbf{M}_n^+ for the pointed graph associated with M_n^+ . Let $(\theta^n, (\mathcal{L}_v^n)_{v \in \theta^n_{(1,2)}})$ be the random labeled tree associated with M_n^+ via the BDG bijection between \mathbb{T}_+ and $\mathcal{M}_{r,p}^+$. Notice that $(\theta^n, (\mathcal{L}_v^n)_{v \in \theta^n_{(1,2)}})$ has the same distribution as $(\theta, (\mathcal{L}_v)_{v \in \theta_{(1,2)}})$ conditional on $\{\#\theta_{(1)} = n - 1\}$.

We write $(\tilde{\theta}^n, (\tilde{\mathcal{L}}_v^n)_{v \in \tilde{\theta}^n_{(1,2)}})$ for the tree derived from $(\theta^n, (\mathcal{L}_v^n)_{v \in \theta^n_{(1,2)}})$ by the shuffling operation, and σ_n for the shuffling bijection from θ^n onto $\tilde{\theta}^n$. The notation $(\tilde{\theta}^n, (\tilde{\mathcal{L}}_v^n)_{v \in \tilde{\theta}^n_{(1,2)}})$ is consistent with the end of the preceding subsection, since conditioning the tree on having $n - 1$ vertices of type 1 clearly commutes with the shuffling operation.

As previously, $u_0^n = \emptyset, u_1^n, \dots, u_{p_n}^n$ denotes the modified contour sequence of $\tilde{\theta}^n$. For every $j \in \{0, 1, \dots, p_n\}$, we set $v_j^n = \sigma_n^{-1}(u_j^n)$. Recall that by construction the type of u_j^n (in $\tilde{\theta}^n$) coincides with the type of v_j^n (in θ^n).

Using the Skorokhod representation theorem, we may assume that the convergence (8) holds almost surely. We will then prove that the convergence

$$B_q n^{-1/4} \cdot \text{Cac}(\mathbf{M}_n^+) \xrightarrow[n \rightarrow \infty]{} \text{KAC} \tag{9}$$

also holds almost surely, in the Gromov–Hausdorff sense.

We first define a correspondence \mathcal{R}_n^0 between \mathcal{T}_e and $V(M_n^+)$ by declaring that (a_*, ρ_n) belongs to \mathcal{R}_n^0 , and, for every $s \in [0, 1]$:

- if $v_{[p_n s]}^n$ is of type 1, $(p_e(s), v_{[p_n s]}^n)$ belongs to \mathcal{R}_n^0 ;
- if $v_{[p_n s]}^n$ is of type 2, then if w is any of the two (possibly equal) vertices of type 1 associated with $v_{[p_n s]}^n$, $(p_e(s), w)$ belongs to \mathcal{R}_n^0 .

We then write \mathcal{R}_n for the induced correspondence between the quotient spaces $\text{KAC} = \mathcal{T}_e / \simeq$ and $\text{Cac}(\mathbf{M}_n^+)$. A pair $(x, \alpha) \in \text{KAC} \times \text{Cac}(\mathbf{M}_n^+)$ belongs to \mathcal{R}_n if and only if there exists a representative a of x in \mathcal{T}_e and a representative u of α in $V(M_n^+)$ such that $(a, u) \in \mathcal{R}_n^0$.

Thanks to (3), the convergence (9) will be proved if we can verify that the distortion of \mathcal{R}_n , when KAC is equipped with the distance d_{KAC} and $\text{Cac}(\mathbf{M}_n^+)$ is equipped with $B_{\mathbf{q}} n^{-1/4} d_{\text{Cac}}^{\mathbf{M}_n^+}$, tends to 0 as $n \rightarrow \infty$, almost surely. To this end, it is enough to verify that

$$\lim_{n \rightarrow \infty} \sup_{0 \leq s \leq 1} |d_{\text{KAC}}(a_*, p_e(s)) - B_{\mathbf{q}} n^{-1/4} d_{\text{Cac}}^{\mathbf{M}_n^+}(\rho_n, \widehat{v}_{[p_n s]}^n)| = 0 \quad \text{a.s.} \quad (10)$$

and

$$\lim_{n \rightarrow \infty} \sup_{s, t \in [0, 1]} |d_{\text{KAC}}(p_e(s), p_e(t)) - B_{\mathbf{q}} n^{-1/4} d_{\text{Cac}}^{\mathbf{M}_n^+}(\widehat{v}_{[p_n s]}^n, \widehat{v}_{[p_n t]}^n)| = 0 \quad \text{a.s.} \quad (11)$$

In both (10) and (11), $\widehat{v}_{[p_n s]}^n = v_{[p_n s]}^n$ if $v_{[p_n s]}^n$ is of type 1, whereas, if $v_{[p_n s]}^n$ is of type 2, $\widehat{v}_{[p_n s]}^n$ stands for one of the vertices of type 1 associated with $v_{[p_n s]}^n$ (obviously the validity of (10) and (11) does not depend on the choice of this vertex).

The proof of (10) is easy. Note that

$$d_{\text{KAC}}(a_*, p_e(s)) = Z_{p_e(s)} - Z_{a_*} = Z_s - \underline{Z}$$

and, by (7),

$$d_{\text{Cac}}^{\mathbf{M}_n^+}(\rho_n, \widehat{v}_{[p_n s]}^n) = d_{\text{gr}}^{\mathbf{M}_n^+}(\rho_n, \widehat{v}_{[p_n s]}^n) = \mathcal{L}_{\widehat{v}_{[p_n s]}^n}^n - \min \mathcal{L}^n + 1$$

so that

$$|d_{\text{Cac}}^{\mathbf{M}_n^+}(\rho_n, \widehat{v}_{[p_n s]}^n) - (\mathcal{L}_{\widehat{v}_{[p_n s]}^n}^n - \min \mathcal{L}^n)| \leq 1.$$

Since $\mathcal{L}_{\widehat{v}_{[p_n s]}^n}^n - \min \mathcal{L}^n = \tilde{\mathcal{L}}_{\widehat{v}_{[p_n s]}^n}^n - \min \tilde{\mathcal{L}}^n = V_{[p_n s]}^n - \min V^n$, our claim (10) follows from the (almost sure) convergence (8).

It remains to establish (11). It suffices to prove that almost surely, for every choice of the sequences (s_n) and (t_n) in $[0, 1]$, we have

$$\lim_{n \rightarrow \infty} |d_{\text{KAC}}(p_e(s_n), p_e(t_n)) - B_{\mathbf{q}} n^{-1/4} d_{\text{Cac}}^{\mathbf{M}_n^+}(\widehat{v}_{[p_n s_n]}^n, \widehat{v}_{[p_n t_n]}^n)| = 0.$$

We will prove that the preceding convergence holds for all choices of the sequences (s_n) and (t_n) , on the set of full probability measure where the convergence (8) holds. From now on we argue on the latter set.

By a compactness argument, we may assume that the sequences (s_n) and (t_n) converge to s and t respectively as $n \rightarrow \infty$. The proof then reduces to checking that

$$\lim_{n \rightarrow \infty} B_{\mathbf{q}} n^{-1/4} d_{\text{Cac}}^{\mathbf{M}_n^+}(\widehat{v}_{[p_n s_n]}^n, \widehat{v}_{[p_n t_n]}^n) = d_{\text{KAC}}(p_e(s), p_e(t)) = Z_s + Z_t - 2 \min_{c \in \llbracket p_e(s), p_e(t) \rrbracket} Z_c.$$

From Corollary 4.4 (and the fact that the sequence \mathbf{q} is finitely supported), this will follow if we can verify that

$$\lim_{n \rightarrow \infty} B_{\mathbf{q}} n^{-1/4} \left(\mathcal{L}_{\widehat{v}_{[p_n s_n]}^n}^n + \mathcal{L}_{\widehat{v}_{[p_n t_n]}^n}^n - 2 \min_{w \in \llbracket \widehat{v}_{[p_n s_n]}^n, \widehat{v}_{[p_n t_n]}^n \rrbracket} \mathcal{L}_w^n \right) = Z_s + Z_t - 2 \min_{c \in \llbracket p_e(s), p_e(t) \rrbracket} Z_c.$$

Observe that

$$\left| \mathcal{L}_{\widehat{v}_{[p_n s_n]}^n}^n - \mathcal{L}_{v_{[p_n s_n]}^n}^n \right| \leq 1$$

and $\mathcal{L}_{v_{[p_n s_n]}^n}^n = \widetilde{\mathcal{L}}_{u_{[p_n s_n]}^n}^n$. From the convergence (8), we have

$$\lim_{n \rightarrow \infty} B_{\mathbf{q}} n^{-1/4} \mathcal{L}_{\widehat{v}_{[p_n s_n]}^n}^n = \lim_{n \rightarrow \infty} B_{\mathbf{q}} n^{-1/4} \widetilde{\mathcal{L}}_{u_{[p_n s_n]}^n}^n = \lim_{n \rightarrow \infty} B_{\mathbf{q}} n^{-1/4} V_{[p_n s_n]}^n = Z_s$$

and similarly if the sequence (s_n) is replaced by (t_n) . Finally, we need to verify that

$$\lim_{n \rightarrow \infty} \left(B_{\mathbf{q}} n^{-1/4} \min_{w \in \llbracket \widehat{v}_{[p_n s_n]}^n, \widehat{v}_{[p_n t_n]}^n \rrbracket} \mathcal{L}_w^n \right) = \min_{c \in \llbracket p_{\mathbf{e}}(s), p_{\mathbf{e}}(t) \rrbracket} Z_c. \quad (12)$$

In proving (12), we may replace $\widehat{v}_{[p_n s_n]}^n$ and $\widehat{v}_{[p_n t_n]}^n$ by $v_{[p_n s_n]}^n$ and $v_{[p_n t_n]}^n$ respectively. The point is that if u is a vertex of θ^n of type 2 and v is an associated vertex of type 1, our definitions imply that $\min_{w \in \llbracket u, v \rrbracket} \mathcal{L}_w^n = \mathcal{L}_v^n$. Without loss of generality we can also assume that $s \leq t$.

Since $\llbracket u_{[p_n s_n]}^n, u_{[p_n t_n]}^n \rrbracket$ is the image under σ_n of $\llbracket v_{[p_n s_n]}^n, v_{[p_n t_n]}^n \rrbracket$, (12) will hold if we can prove that

$$\lim_{n \rightarrow \infty} \left(B_{\mathbf{q}} n^{-1/4} \min_{w \in \llbracket u_{[p_n s_n]}^n, u_{[p_n t_n]}^n \rrbracket} \widetilde{\mathcal{L}}_w^n \right) = \min_{c \in \llbracket p_{\mathbf{e}}(s), p_{\mathbf{e}}(t) \rrbracket} Z_c. \quad (13)$$

Let us first prove the upper bound

$$\limsup_{n \rightarrow \infty} \left(B_{\mathbf{q}} n^{-1/4} \min_{w \in \llbracket u_{[p_n s_n]}^n, u_{[p_n t_n]}^n \rrbracket} \widetilde{\mathcal{L}}_w^n \right) \leq \min_{c \in \llbracket p_{\mathbf{e}}(s), p_{\mathbf{e}}(t) \rrbracket} Z_c. \quad (14)$$

Let us pick $c \in \llbracket p_{\mathbf{e}}(s), p_{\mathbf{e}}(t) \rrbracket$. We may assume that $c \neq p_{\mathbf{e}}(s)$ and $c \neq p_{\mathbf{e}}(t)$ (otherwise the desired lower bound immediately follows from the convergence (8)). Then, we can find $r \in (s, t)$ such that $c = p_{\mathbf{e}}(r)$ and either $\mathbf{e}_u > \mathbf{e}_r$, for every $u \in [s, r)$, or $\mathbf{e}_u > \mathbf{e}_r$ for every $u \in (r, t]$. Consider only the first case, since the second one can be treated in a similar manner. The convergence of the rescaled contour processes then guarantees that we can find a sequence (k_n) of positive integers such that $k_n/p_n \rightarrow r$ as $n \rightarrow \infty$, and

$$C_k^n > C_{k_n}^n \quad \text{for every } k \in \{[p_n s_n], [p_n s_n] + 1, \dots, k_n - 1\}$$

for all sufficiently large n . The latter property, and the construction of the contour sequence of the tree θ^n , ensure that $u_{k_n}^n \in \llbracket u_{[p_n s_n]}^n, u_{[p_n t_n]}^n \rrbracket$, for all sufficiently large n . However, by the convergence of the rescaled label processes, we have

$$\lim_{n \rightarrow \infty} B_{\mathbf{q}} n^{-1/4} \widetilde{\mathcal{L}}_{u_{k_n}^n}^n = Z_r = Z_c.$$

Consequently,

$$\limsup_{n \rightarrow \infty} \left(B_{\mathbf{q}} n^{-1/4} \min_{w \in \llbracket u_{[p_n s_n]}^n, u_{[p_n t_n]}^n \rrbracket} \widetilde{\mathcal{L}}_w^n \right) \leq Z_c$$

and since this holds for every choice of c the upper bound (14) follows.

Let us turn to the lower bound

$$\liminf_{n \rightarrow \infty} \left(B_{\mathbf{q}} n^{-1/4} \min_{w \in \llbracket u_{[p_n s_n]}^n, u_{[p_n t_n]}^n \rrbracket} \widetilde{\mathcal{L}}_w^n \right) \geq \min_{c \in \llbracket p_{\mathbf{e}}(s), p_{\mathbf{e}}(t) \rrbracket} Z_c. \quad (15)$$

For every n , let $w_n \in \llbracket u_{[p_n s_n]}^n, u_{[p_n t_n]}^n \rrbracket$ be such that

$$\min_{w \in \llbracket u_{[p_n s_n]}^n, u_{[p_n t_n]}^n \rrbracket} \widetilde{\mathcal{L}}_w^n = \widetilde{\mathcal{L}}_{w_n}^n.$$

We can write $w_n = u_{j_n}^n$ where $j_n \in \{[p_n s_n], [p_n s_n] + 1, \dots, [p_n t_n]\}$ is such that

$$C_{j_n}^n = \min_{[p_n s_n] \leq j \leq j_n} C_j^n \quad \text{or} \quad C_{j_n}^n = \min_{j_n \leq j \leq [p_n t_n]} C_j^n. \quad (16)$$

We need to verify that

$$\liminf_{n \rightarrow \infty} B_{\mathbf{q}} n^{-1/4} \tilde{\mathcal{L}}_{w_n}^n \geq \min_{c \in \llbracket p_{\mathbf{e}}(s), p_{\mathbf{e}}(t) \rrbracket} Z_c.$$

We argue by contradiction and suppose that there exist $\varepsilon > 0$ and a subsequence (n_k) such that, for every k ,

$$B_{\mathbf{q}} n_k^{-1/4} \tilde{\mathcal{L}}_{w_{n_k}}^{n_k} \leq \min_{c \in \llbracket p_{\mathbf{e}}(s), p_{\mathbf{e}}(t) \rrbracket} Z_c - \varepsilon.$$

By extracting another subsequence if necessary, we may assume furthermore that $j_{n_k}/p_{n_k} \rightarrow r \in [s, t]$ as $k \rightarrow \infty$, and that the first equality in (16) holds with $n = n_k$ for every k (the case when the other equality holds is treated in a similar manner). Then, from the convergence of rescaled contour processes, we have

$$\mathbf{e}_r = \min_{s \leq u \leq r} \mathbf{e}_r,$$

which implies that $p_{\mathbf{e}}(r) \in \llbracket p_{\mathbf{e}}(s), p_{\mathbf{e}}(t) \rrbracket$. Furthermore, from the convergence of rescaled label processes,

$$Z_{p_{\mathbf{e}}(r)} = Z_r = \lim_{k \rightarrow \infty} B_{\mathbf{q}} n_k^{-1/4} \tilde{\mathcal{L}}_{w_{n_k}}^{n_k} \leq \min_{c \in \llbracket p_{\mathbf{e}}(s), p_{\mathbf{e}}(t) \rrbracket} Z_c - \varepsilon.$$

This contradiction completes the proof of (15) and of the convergence (9).

In order to complete the proof of Theorem 4.5 under Assumption (A1), it suffices to verify that the convergence (9) also holds (in distribution) if M_n^+ is replaced by a random planar map M_n^- distributed according to $P_{\mathbf{q}}^-(\cdot \mid \#V(m) = n)$, or by a random planar map M_n^0 distributed according to $P_{\mathbf{q}}^0(\cdot \mid \#V(m) = n)$. The first case is trivial since M_n^- can be obtained from M_n^+ simply by reversing the orientation of the root edge. The case of M_n^0 is treated by a similar method as the one we used for M_n^+ . We first need an analogue of Proposition 4.6, which is provided by the last statement of Proposition 3 in [19]. In this analogue, the random labeled tree associated with a planar map distributed according to $P_{\mathbf{q}}^0$ is described as the concatenation (at the root vertex) of two independent labeled Galton–Watson trees whose root is of type 2, with the same offspring distributions as in Proposition 4.6. The results of [20] can be used to verify that Proposition 4.8 still holds with the same constants $A_{\mathbf{q}}$ and $B_{\mathbf{q}}$, and the remaining part of the argument goes through without change. This completes the proof of Theorem 4.5 under Assumption (A1).

4.5. The bipartite case

In this section, we briefly discuss the proof of Theorem 4.5 under Assumption (A2). In that case, since $W_{\mathbf{q}}(\mathcal{M}_{r,p}^0) = 0$, it is obviously enough to prove the convergence of Theorem 4.5 with M_n replaced by M_n^+ . The proof becomes much simpler because we do not need the shuffling operation. As previously, we introduce the labeled tree $(\theta^n, (\mathcal{L}_v^n)_{v \in \theta^n_{(1,2)}})$ associated with M_n^+ via the BDG bijection, but we now define $u_0^n = \emptyset, u_1^n, \dots, u_{p_n}^n = \emptyset$ as the modified contour sequence of θ_n (instead of $\tilde{\theta}_n$). We then define the contour process $C_i^n = |u_i^n|$ and the label process $V_i^n = \mathcal{L}_{u_i^n}^n$, for $0 \leq i \leq p_n$. Proposition 3.7 then holds in exactly the same form, as a consequence of the results of [17]. The reason why we do not need the shuffling operation is the fact that the label increments of $(\theta^n, (\mathcal{L}_v^n)_{v \in \theta^n_{(1,2)}})$ are centered in the bipartite case.

Once the convergence (8) is known to hold, it suffices to repeat all steps of the proof in Section 4.4, replacing $\tilde{\theta}^n$ by θ^n and v_i^n by u_i^n wherever this is needed. We leave the details to the reader.

5. The dimension of the Brownian cactus

In this section, we compute the Hausdorff dimension of the Brownian cactus KAC. We write $\mathbf{p}: \mathcal{T}_{\mathbf{e}} \rightarrow \text{KAC} = \mathcal{T}_{\mathbf{e}}/\asymp$ for the canonical projection. The uniform measure μ on KAC is the image of the mass measure Vol on the CRT (see

Section 3) under \mathfrak{p} . For every x in KAC and every $\delta \geq 0$, we denote the closed ball of center x and radius δ in KAC by $B_{\text{KAC}}(x, \delta)$. The following theorem gives information about the μ -measure of these balls around a typical point of KAC.

Proposition 5.1. (i) *We have*

$$\mathbb{E} \left[\int \mu(dx) \mu(B_{\text{KAC}}(x, \delta)) \right] = \frac{2^{5/4} \Gamma(1/4)}{3\sqrt{\pi}} \delta^3 + o(\delta^3),$$

as $\delta \rightarrow 0$.

(ii) *For every $\varepsilon > 0$,*

$$\limsup_{\delta \rightarrow 0} \frac{\mu(B_{\text{KAC}}(x, \delta))}{\delta^{4-\varepsilon}} = 0, \quad \mu(dx) \text{ a.e., a.s.}$$

Remark 5.2. *Let U be uniformly distributed over $[0, 1]$, so that $p_e(U)$ is distributed according to Vol and $X = \mathfrak{p} \circ p_e(U)$ is distributed according to μ . Assertion (i) of the theorem says that the mean volume of the ball $B_{\text{KAC}}(X, \delta)$ is of order δ^3 , whereas assertion (ii) shows that almost surely the volume of this ball will be bounded above by $\delta^{4-\varepsilon}$ when δ is small. This difference between the mean and the almost sure behavior is specific to the Brownian cactus. In the case of the Brownian map, results from Section 6 of [14] show that δ^4 is the correct order both for the mean and the almost sure behavior of the volume of a typical ball of radius δ .*

In relation with this, we see that in contrast with the CRT or the Brownian map, the Brownian cactus is not invariant under re-rooting according to the “uniform” measure μ . This means that KAC re-rooted at X does not have the same distribution as KAC. Indeed, since $d_{\text{KAC}}^{\mathbb{E}}(\rho, x) = d(\rho, x)$ for every pointed geodesic space $\mathbb{E} = (E, d, \rho)$, the previous considerations, and Proposition 3.2, entail that $\mu(B_{\text{KAC}}(\rho, \delta))$ is of order δ^4 both in the mean and in the a.s. sense.

Proof of Proposition 5.1. (i) Fix $\delta > 0$. Let U and U' be two independent random variables that are uniformly distributed over $[0, 1]$ and independent of (\mathbf{e}, Z) . By the very definition of μ , we have

$$\mathbb{E} \left[\int \mu(dx) \mu(B_{\text{KAC}}(x, \delta)) \right] = \mathbb{P}[d_{\text{KAC}}(p_e(U), p_e(U')) \leq \delta].$$

The value of $d_{\text{KAC}}(p_e(U), p_e(U'))$ is determined by the labels Z_a for $a \in \llbracket p_e(U), p_e(U') \rrbracket$. Write $(g_{U,U'}(t), 0 \leq t \leq d_e(U, U'))$ for the geodesic path from $p_e(U)$ to $p_e(U')$ in the tree $\mathcal{T}_{\mathbf{e}}$ (so that $\llbracket p_e(U), p_e(U') \rrbracket$ is the range of $g_{U,U'}$). Then, conditionally on the triplet (\mathbf{e}, U, U') the process

$$(Z_{g_{U,U'}(t)} - Z_{p_e(U)})_{0 \leq t \leq d_e(U, U')},$$

is a standard linear Brownian motion. Hence if $(B_t)_{t \geq 0}$ is a linear Brownian motion independent of (\mathbf{e}, U, U') , we have

$$\mathbb{P}[d_{\text{KAC}}(p_e(U), p_e(U')) \leq \delta] = \mathbb{P}\left[B_L - 2 \min_{0 \leq s \leq L} B_s \leq \delta\right],$$

where $L = d_e(U, U')$. Pitman’s theorem [23], Theorem VI.3.5, implies that, for every fixed $l \geq 0$, $B_l - 2 \min_{0 \leq s \leq l} B_s$ has the same distribution as $B_l^{(3)}$, where $(B_t^{(3)})_{t \geq 0}$ denotes a three-dimensional Bessel process started from 0. From the invariance under uniform re-rooting of the distribution of the CRT, the variable $d_e(U, U')$ has the same distribution as $d_e(0, U) = e_U$, which has density $4le^{-2l^2}$. Consequently, we can explicitly compute

$$\begin{aligned} \mathbb{P}[d_{\text{KAC}}(U, U') \leq \delta] &= 4 \int_0^\infty dl e^{-2l^2} \mathbb{P}[B_l^{(3)} \leq \delta] \\ &= 4 \int_0^\infty dl e^{-2l^2} \int_{\mathbb{R}^3} dz (2\pi l)^{-3/2} e^{-|z|^2/2l} \mathbf{1}_{\{|z| \leq \delta\}} \end{aligned}$$

$$\begin{aligned}
 &= 4\sqrt{\frac{2}{\pi}} \int_0^\infty dl l^{-1/2} e^{-2l^2} \int_0^\delta du u^2 e^{-u^2/2l} \\
 &= 4\sqrt{\frac{2}{\pi}} \int_0^\delta du u^2 \int_0^\infty dl l^{-1/2} \exp(-2l^2 - (u^2/2l)).
 \end{aligned}$$

The desired result follows since

$$\lim_{u \rightarrow 0} \int_0^\infty dl l^{-1/2} \exp(-2l^2 - (u^2/2l)) = \int_0^\infty dl l^{-1/2} \exp(-2l^2) = 2^{-5/4} \Gamma(1/4).$$

(ii) Let us fix $r \in]0, 1[$. For every $u \in [0, \mathbf{e}_r]$, set

$$G_{\mathbf{e}}(r, u) = \max\{s \in [0, r]: \mathbf{e}_s = \mathbf{e}_r - u\}, \quad D_{\mathbf{e}}(r, u) = \min\{s \in [r, 1]: \mathbf{e}_s = \mathbf{e}_r - u\}.$$

Then $p_{\mathbf{e}}(G_{\mathbf{e}}(r, u)) = p_{\mathbf{e}}(D_{\mathbf{e}}(r, u))$ is a point of $[[p_{\mathbf{e}}(0), p_{\mathbf{e}}(r)]]$, and more precisely the path $u \rightarrow p_{\mathbf{e}}(G_{\mathbf{e}}(r, u))$, $0 \leq u \leq \mathbf{e}_r$, is the geodesic from $p_{\mathbf{e}}(r)$ to $p_{\mathbf{e}}(0)$ in the tree $\mathcal{T}_{\mathbf{e}}$. As a consequence, conditionally on \mathbf{e} , the process

$$M_u^{(r)} := Z_r - \min\{Z_v: v \in [[p_{\mathbf{e}}(G_{\mathbf{e}}(r, u)), p_{\mathbf{e}}(r)]]\}, \quad 0 \leq u \leq \mathbf{e}_r,$$

has the same distribution as

$$- \min_{0 \leq v \leq u} B_v, \quad 0 \leq u \leq \mathbf{e}_r,$$

where B is as above. By classical results (see e.g. Theorem 6.2 in [10]), we have, for every $\varepsilon \in]0, 1/2[$,

$$\lim_{u \rightarrow 0} u^{-1/2-\varepsilon} M_u^{(r)} = \infty, \quad \text{a.s.} \tag{17}$$

On the other hand, if $t \in [0, 1] \setminus]G_{\mathbf{e}}(r, u), D_{\mathbf{e}}(r, u)[$, we have $\min_{t \wedge r \leq s \leq t \vee r} \mathbf{e}_s \leq \mathbf{e}_r - u$, which implies that the segment $[[p_{\mathbf{e}}(t), p_{\mathbf{e}}(r)]]$ contains $[[p_{\mathbf{e}}(G_{\mathbf{e}}(r, u)), p_{\mathbf{e}}(r)]]$, and therefore

$$d_{\text{KAC}}(p_{\mathbf{e}}(t), p_{\mathbf{e}}(r)) \geq M_u^{(r)}.$$

Using (17), it follows that, for every fixed $\varepsilon \in]0, 1/2[$, we have a.s. for all $u > 0$ small enough

$$B_{\text{KAC}}(p_{\mathbf{e}}(r), u^{1/2+\varepsilon}) \subset (\text{KAC} \setminus \mathfrak{p} \circ p_{\mathbf{e}}([0, G_{\mathbf{e}}(r, u)] \cup [D_{\mathbf{e}}(r, u), 1])),$$

and in particular

$$\mu(B_{\text{KAC}}(p_{\mathbf{e}}(r), u^{1/2+\varepsilon})) \leq D_{\mathbf{e}}(r, u) - G_{\mathbf{e}}(r, u).$$

However, the same standard results about Brownian motion that we already used to derive (17) imply that

$$\lim_{u \rightarrow 0} u^{-2+\varepsilon} (D_{\mathbf{e}}(r, u) - G_{\mathbf{e}}(r, u)) = 0, \quad \text{a.s.}$$

We conclude that, for every $\varepsilon \in]0, 1/2[$,

$$\lim_{u \rightarrow 0} u^{-2+\varepsilon} \mu(B_{\text{KAC}}(p_{\mathbf{e}}(r), u^{1/2+\varepsilon})) = 0, \quad \text{a.s.}$$

and property (ii) follows, in fact in a slightly stronger form than stated in the theorem. □

Corollary 5.3. *Almost surely, the Hausdorff dimension of KAC is 4.*

Proof. Classical density theorems for Hausdorff measures show that the existence of a non-trivial measure μ satisfying the property stated in part (ii) of Proposition 5.1 implies the lower bound $\dim(\text{KAC}) \geq 4$. To get the corresponding upper bound, we first note that the mapping $[0, 1] \ni t \rightarrow Z_t$ is a.s. Hölder continuous with exponent $1/4 - \varepsilon$, for any $\varepsilon \in]0, 1/4[$. Observing that $[[p_e(t), p_e(t')]] \subset p_e([t \wedge t', t \vee t'])$, for every $t, t' \in [0, 1]$, it readily follows that the composition $p \circ p_e$ defined on $[0, 1]$ and with values in KAC, is a.s. Hölder continuous with exponent $1/4 - \varepsilon$, for any $\varepsilon \in]0, 1/4[$. Hence, the Hausdorff dimension of KAC, which is the range of $p \circ p_e$, must be bounded above by 4. \square

6. Separating cycles

In this section, we study the existence and properties of a cycle with minimal length separating two points of the Brownian map, under the condition that this cycle contains a third point. This is really a problem about the Brownian map, but the cactus distance plays an important role in the statement. Our results in this section are related to the work of Bouttier and Guitter [4] for large random quadrangulations of the plane.

We consider the Brownian map $m_\infty = (m_\infty, D, \rho_*)$ (see Section 3). We set $\mathbf{p} = \Pi \circ p_e$, which corresponds to the canonical projection from $[0, 1]$ onto m_∞ . If U is uniformly distributed over $[0, 1]$, the point $\mathbf{p}(U)$ is distributed according to the volume measure λ on m_∞ .

A loop in m_∞ is a continuous path $\gamma : [0, T] \rightarrow m_\infty$, where $T > 0$, such that $\gamma(0) = \gamma(T)$. If x and y are two distinct points of m_∞ , we say that the loop γ separates the points x and y if x and y lie in distinct connected components of $m_\infty \setminus \{\gamma(t) : 0 \leq t \leq T\}$. It is known [16] that (m_∞, D) is homeomorphic to the 2-sphere, so that separating loops do exist. We denote by $S(x, y, \rho_*)$ the set of all loops γ such that $\gamma(0) = \rho_*$ and γ separates x and y . Recall from Section 2.2 the definition of the length of a curve in a metric space.

Theorem 6.1. *Let U_1 and U_2 be independent and uniformly distributed over $[0, 1]$. Then almost surely there exists a unique loop $\gamma_* \in S(\mathbf{p}(U_1), \mathbf{p}(U_2), \rho_*)$ with minimal length, up to reparametrization and time-reversal. This loop is obtained as the concatenation of the two distinct geodesic paths from $\Pi(\beta)$ to ρ_* , where β is the a.s. unique point of $[[p_e(U_1), p_e(U_2)]]$ such that*

$$Z_\beta = \min_{a \in [[p_e(U_1), p_e(U_2)]]} Z_a.$$

In particular, the length of γ_ is*

$$L(\gamma_*) = 2D(\rho_*, \Pi(\beta)) = D(\rho_*, \mathbf{p}(U_1)) + D(\rho_*, \mathbf{p}(U_2)) - 2d_{\text{KAC}}(p_e(U_1), p_e(U_2)).$$

The complement in m_∞ of the range of γ_ has exactly two components C_1 and C_2 , such that $\mathbf{p}(U_1) \in C_1$ and $\mathbf{p}(U_2) \in C_2$, and the pair $(\lambda(C_1), \lambda(C_2))$ is distributed according to the beta distribution with parameters $(\frac{1}{4}, \frac{1}{4})$:*

$$\mathbb{E}[f(\lambda(C_1), \lambda(C_2))] = \frac{\Gamma(1/2)}{\Gamma(1/4)^2} \int_0^1 dt (t(1-t))^{-3/4} f(t, 1-t)$$

for any non-negative Borel function f on \mathbb{R}_+^2 .

Proof. We first explain how the loop γ_* is constructed. As in the previous section, write $(g_{U_1, U_2}(r))_{0 \leq r \leq d_e(U_1, U_2)}$ for the geodesic path from $p_e(U_1)$ to $p_e(U_2)$ in the tree \mathcal{T}_e , whose range is the segment $[[p_e(U_1), p_e(U_2)]]$. We already noticed that, conditionally on the triplet (e, U_1, U_2) the process

$$(Z_{g_{U_1, U_2}(r)} - Z_{p_e(U_1)})_{0 \leq r \leq d_e(U_1, U_2)},$$

is a standard linear Brownian motion. Hence this process a.s. attains its minimal value at a unique time $r_0 \in]0, d_e(U_1, U_2)[$, and we set $\beta = g_{U_1, U_2}(r_0)$. Since there are only countably many values of $r \in]0, d_e(U_1, U_2)[$ such that $g_{U_1, U_2}(r)$ has multiplicity 3 in \mathcal{T}_e , it is also clear that β has multiplicity 2 in \mathcal{T}_e , a.s. Write C_1° and C_2° for the two connected components of $\mathcal{T}_e \setminus \{\beta\}$, ordered in such a way that $p_e(U_1) \in C_1^\circ$ and $p_e(U_2) \in C_2^\circ$, and set $C_1 = C_1^\circ \cup \{\beta\}$, $C_2 = C_2^\circ \cup \{\beta\}$. Then $\Pi(C_1)$ and $\Pi(C_2)$ are closed subsets of m_∞ whose union is m_∞ . Furthermore, the discussion at

the beginning of Section 3 of [14] shows that the boundary of $\Pi(\mathcal{C}_1)$, or equivalently the boundary of $\Pi(\mathcal{C}_2)$, coincides with the set $\Pi(\mathcal{C}_1) \cap \Pi(\mathcal{C}_2)$ of all points $x \in m_\infty$ that can be written as $x = \Pi(a_1) = \Pi(a_2)$ for some $a_1 \in \mathcal{C}_1$ and $a_2 \in \mathcal{C}_2$. In particular, the interiors of $\Pi(\mathcal{C}_1)$ and of $\Pi(\mathcal{C}_2)$ are disjoint. Notice that $\mathbf{p}(U_1)$ belongs to the interior of $\Pi(\mathcal{C}_1)$, and $\mathbf{p}(U_2)$ belongs to the interior of $\Pi(\mathcal{C}_2)$, almost surely: To see this, observe that for almost every (in the sense of the volume measure Vol) point a of \mathcal{T}_e , the equivalence class of a for \approx is a singleton, and thus $\Pi^{-1}(\mathbf{p}(U_1))$ and $\Pi^{-1}(\mathbf{p}(U_2))$ must be singletons almost surely.

Since β has multiplicity 2 in \mathcal{T}_e , Theorem 7.6 in [14] implies that there are exactly two distinct geodesic paths from ρ_* to $\Pi(\beta)$, and that these paths are simple geodesics in the sense of [14], Section 4. We denote these geodesic paths by ϕ_1 and ϕ_2 . From the definition of simple geodesics, one easily gets that $\phi_1(s) = \phi_2(s)$ for every $0 \leq s \leq s_0$, where

$$s_0 := \max\left(\min_{a \in \mathcal{C}_1} Z_a, \min_{a \in \mathcal{C}_2} Z_a\right) - \underline{Z}.$$

Note that $\{\phi_1(s): 0 \leq s < s_0\}$ is contained in the interior of $\Pi(\mathcal{C}_i)$, where $i \in \{1, 2\}$ is determined by the condition $a_* \in \mathcal{C}_i$. Furthermore, the definition of simple geodesics shows that

$$\Pi(\mathcal{C}_1) \cap \Pi(\mathcal{C}_2) = \{\phi_1(s): s_0 \leq s \leq D(\rho_*, \Pi(\beta))\} \cup \{\phi_2(s): s_0 \leq s \leq D(\rho_*, \Pi(\beta))\}.$$

We define γ_* by setting

$$\gamma_*(t) = \begin{cases} \phi_1(t) & \text{if } 0 \leq t \leq D(\rho_*, \Pi(\beta)), \\ \phi_2(2D(\rho_*, \Pi(\beta)) - t) & \text{if } D(\rho_*, \Pi(\beta)) \leq t \leq 2D(\rho_*, \Pi(\beta)). \end{cases}$$

Then γ_* is a loop starting and ending at ρ_* . Furthermore γ_* separates $\mathbf{p}(U_1)$ and $\mathbf{p}(U_2)$, since any continuous path in m_∞ starting from $\mathbf{p}(U_1)$ will have to hit the boundary of $\Pi(\mathcal{C}_1)$ before reaching $\mathbf{p}(U_2)$. Finally the length of γ_* is

$$L(\gamma_*) = 2D(\rho_*, \Pi(\beta)) = 2(Z_\beta - \underline{Z}) = D(\rho_*, \mathbf{p}(U_1)) + D(\rho_*, \mathbf{p}(U_2)) - 2d_{\text{KAC}}(p_e(U_1), p_e(U_2)).$$

We next verify that γ_* is the unique loop in $S(\mathbf{p}(U_1), \mathbf{p}(U_2), \rho_*)$ with minimal length. Let γ be a path in $S(\mathbf{p}(U_1), \mathbf{p}(U_2), \rho_*)$ indexed by the interval $[0, T]$. The image under Π of the path g_{U_1, U_2} is a continuous path from $\mathbf{p}(U_1)$ to $\mathbf{p}(U_2)$, which must intersect the range of γ . Hence the range of γ contains at least one point y such that $y = \Pi(a)$ for some $a \in \llbracket p_e(U_1), p_e(U_2) \rrbracket$. Since $\gamma(0) = \gamma(T) = \rho_*$, we have

$$L(\gamma) \geq 2D(\rho_*, y) = 2(Z_a - \underline{Z})$$

using (5). Since $Z_a \geq Z_\beta$, we thus obtain that $L(\gamma) \geq L(\gamma_*)$.

Let $\tau \in [0, T]$ be such that $y = \gamma(\tau)$. The preceding considerations show that the equality $L(\gamma) = L(\gamma_*)$ can hold only if $a = \beta$ and if furthermore the paths $(\gamma(\tau - t), 0 \leq t \leq \tau)$ and $(\gamma(\tau + t), 0 \leq t \leq T - \tau)$ have length $D(\rho_*, \Pi(\beta))$, so that these paths must coincide (up to reparametrization) with geodesics from $\Pi(\beta)$ to ρ_* . We conclude that any minimizing path γ coincides with γ_* , up to reparametrization and time-reversal.

In order to complete the proof of the theorem, we first need to identify the connected components of the complement of the range of γ_* in m_∞ . Consider the case when a_* belongs to \mathcal{C}_1 , and set

$$\mathcal{R} := \{\phi_1(s): 0 \leq s < s_0\} \subset \Pi(\mathcal{C}_1).$$

Write $\text{Int}(\Pi(\mathcal{C}_i))$ for the interior of $\Pi(\mathcal{C}_i)$, for $i = 1, 2$. Then the connected components of the complement of the range of γ_* in m_∞ are

$$\mathcal{C}_1 = \text{Int}(\Pi(\mathcal{C}_1)) \setminus \mathcal{R}, \quad \mathcal{C}_2 = \text{Int}(\Pi(\mathcal{C}_2)).$$

This easily follows from the preceding considerations: Note for instance that $\text{Int}(\Pi(\mathcal{C}_2))$ is the image under Π of a connected subset of \mathcal{C}_2 , and is therefore connected. From this identification, we get

$$\lambda(\mathcal{C}_1) = \text{Vol}(\mathcal{C}_1), \quad \lambda(\mathcal{C}_2) = \text{Vol}(\mathcal{C}_2) = 1 - \text{Vol}(\mathcal{C}_1), \tag{18}$$

using the fact that the range of γ_* has zero λ -measure (this can be seen from the uniform estimates on the measure of balls found in Section 6 of [14]). Clearly the same identities (18) remain valid in the case when a_* belongs to \mathcal{C}_2 .

To complete the proof, we need to compute the distribution of $\text{Vol}(\mathcal{C}_1)$. To this end it will be convenient to use the invariance of the law of \mathcal{T}_e under uniform re-rooting. Let U be a random variable uniformly distributed over $[0, 1]$, and let α be the (almost surely unique) vertex of $[[p_e(0), p_e(U)]]$ such that $Z_\alpha = \min_{a \in [[p_e(0), p_e(U)]]} Z_a$. Then, if \mathcal{C}° is the connected component of $\mathcal{T}_e \setminus \{\alpha\}$ containing $p_e(U)$, the invariance of the CRT under uniform re-rooting implies that

$$\text{Vol}(\mathcal{C}_1) \stackrel{(d)}{=} \text{Vol}(\mathcal{C}^\circ).$$

Now notice that conditionally on the pair (e, U) , the random variable $H = d_e(p_e(0), \alpha)$ is distributed according to the arc-sine law on $[0, e_U]$, with density

$$\frac{1}{\pi\sqrt{s(e_U - s)}}.$$

Moreover,

$$\text{Vol}(\mathcal{C}^\circ) = D_e(U, e_U - H) - G_e(U, e_U - H),$$

where we use the same notation as in the preceding section, for $r \in]0, 1[$ and $u \in [0, e_r]$,

$$G_e(r, u) = \max\{s \leq r : e_s = e_r - u\}, \quad D_e(r, u) = \min\{s \geq r : e_s = e_r - u\}. \tag{19}$$

From the previous remarks, we have, for any non-negative measurable function g on $[0, 1]$,

$$\mathbb{E}[g(\text{Vol}(\mathcal{C}_1))] = \mathbb{E}[g(\text{Vol}(\mathcal{C}^\circ))] = \mathbb{E}\left[\int_0^1 ds \int_0^{e_s} \frac{dh}{\pi\sqrt{h(e_s - h)}} g(D_e(s, h) - G_e(s, h))\right]. \tag{20}$$

In order to compute the right-hand side, it is convenient to argue first under the Itô measure $n(de)$ of positive excursions of linear Brownian motion (see e.g. Chapter XII of [23], where the notation $n_+(de)$ is used). Let $\sigma(e)$ denote the duration of excursion e , and define $D_e(r, u)$ and $G_e(r, u)$, for $r \in]0, \sigma(e)[$ and $0 \leq u \leq e(r)$, in a way analogous to (19). Also write

$$q_h(t) = \frac{h}{\sqrt{2\pi t^3}} \exp\left(-\frac{h^2}{2t}\right)$$

for the density of the hitting time of $h > 0$ by a standard linear Brownian motion. Then, an application of Bismut's decomposition of the Itô measure (see e.g. Theorem XII.4.7 in [23]) gives for every non-negative measurable function f on \mathbb{R}_+^2 ,

$$\begin{aligned} & \int n(de) \int_0^{\sigma(e)} ds \int_0^{e(s)} \frac{dh}{\pi\sqrt{h(e(s) - h)}} f(\sigma(e), D_e(s, h) - G_e(s, h)) \\ &= \int_0^\infty du \int_0^u \frac{dh}{\pi\sqrt{h(u - h)}} \int_0^\infty dt q_{2h}(t) \int_0^\infty dt' q_{2(u-h)}(t') f(t + t', t) \\ &= \frac{1}{\pi} \int_0^\infty \frac{dh}{\sqrt{h}} \int_0^\infty \frac{dh'}{\sqrt{h'}} \int_0^\infty dt q_{2h}(t) \int_0^\infty dt' q_{2h'}(t') f(t + t', t) \\ &= \frac{1}{\pi} \int_0^\infty dt \int_0^\infty dt' f(t + t', t) \left(\int_0^\infty \frac{dh}{\sqrt{h}} q_{2h}(t)\right) \left(\int_0^\infty \frac{dh'}{\sqrt{h'}} q_{2h'}(t')\right). \end{aligned} \tag{21}$$

We easily compute

$$\int_0^\infty \frac{dh}{\sqrt{h}} q_{2h}(t) = 2^{-3/4} (2\pi)^{-1/2} \Gamma(3/4) t^{-3/4}.$$

Hence, using also the identity $\Gamma(1/4)\Gamma(3/4) = \pi\sqrt{2}$, we see that the right-hand side of (21) is equal to

$$\frac{2^{-3/2}}{\Gamma(1/4)^2} \int_0^\infty d\ell \int_0^\ell dt f(\ell, t) (t(\ell - t))^{-3/4}.$$

We can condition the resulting formula on $\{\sigma = 1\}$, using the fact that the density of $\sigma(e)$ under $n(de)$ is equal to $\frac{1}{2}(2\pi\ell^3)^{-1/2}$, and we conclude that

$$\begin{aligned} & \mathbb{E} \left[\int_0^1 ds \int_0^{e_s} \frac{dh}{\pi\sqrt{h(e_s - h)}} g(D_{e_s}(s, h) - G_{e_s}(s, h)) \right] \\ &= n \left(\int_0^{\sigma(e)} ds \int_0^{e(s)} \frac{dh}{\pi\sqrt{h(e(s) - h)}} g(D_e(s, h) - G_e(s, h)) \middle| \sigma = 1 \right) \\ &= \frac{\sqrt{\pi}}{\Gamma(1/4)^2} \int_0^1 dt (t(1 - t))^{-3/4} g(t). \end{aligned}$$

We now see that the last assertion of the theorem follows from (20). □

Appendix

This section is devoted to the proof of the fact, mentioned in Section 4.5, that if $\mathbf{q} = (q_1, q_2, \dots)$ is a sequence with finite support, such that $q_k > 0$ for some $k \geq 3$, then there exists a constant $a > 0$ such that $a\mathbf{q} = (aq_1, aq_2, \dots)$ is regular critical in the sense of [17,19]. We briefly discuss case (A2), which is easier. Following [17], we define

$$f_{\mathbf{q}}(x) = \sum_{k \geq 0} x^k \binom{2k+1}{k} q_{2k+2}, \quad x \geq 0.$$

By [17], Proposition 1, the Boltzmann measure $W_{\mathbf{q}}$ defined in Section 4.5 is a finite measure if and only if the equation

$$f_{\mathbf{q}}(x) = 1 - \frac{1}{x}, \quad x > 1, \tag{22}$$

has a solution. Since $q_k > 0$ for some $k \geq 3$, the function $f_{\mathbf{q}}$ is a strictly convex polynomial, so there can be either one or two solutions to this equation. In the first situation, the graphs of $f_{\mathbf{q}}$ and $x \mapsto 1 - 1/x$ are tangent at the unique solution, in which case \mathbf{q} is said to be critical in the sense of [17], Definition 1 (it will even be *regular* critical in our case since $f_{\mathbf{q}}(x)$ is finite for every $x > 0$). It is then trivial that there exists a unique $a = a_c > 0$ such that the graphs of $f_{a\mathbf{q}}$ and $x \mapsto 1 - 1/x$ intersect at a tangency point, and then $a_c\mathbf{q}$ is regular critical.

Let us turn to case (A1), which is more delicate. For every $x, y \geq 0$, we set

$$\begin{aligned} f_{\mathbf{q}}^\bullet(x, y) &= \sum_{k, k' \geq 0} x^k y^{k'} \binom{2k+k'+1}{k+1} \binom{k+k'}{k} q_{2+2k+k'}, \\ f_{\mathbf{q}}^\diamond(x, y) &= \sum_{k, k' \geq 0} x^k y^{k'} \binom{2k+k'}{k} \binom{k+k'}{k} q_{1+2k+k'}, \end{aligned}$$

defining two polynomials, which are convex in each separate variable x, y . Proposition 1 of [19] asserts that the Boltzmann measure $W_{\mathbf{q}}$ is finite (then \mathbf{q} is said to be *admissible*) if and only if the equations

$$\begin{cases} f_{\mathbf{q}}^\bullet(x, y) = 1 - \frac{1}{x}, & x > 1, \\ f_{\mathbf{q}}^\diamond(x, y) = y, & y > 0 \end{cases} \tag{23}$$

have a solution (x, y) , such that the spectral radius of the matrix

$$M(x, y) = \begin{pmatrix} 0 & 0 & x - 1 \\ \frac{x}{y} \partial_x f_{\mathbf{q}}^{\diamond}(x, y) & \partial_y f_{\mathbf{q}}^{\diamond}(x, y) & 0 \\ \frac{x^2}{x-1} \partial_x f_{\mathbf{q}}^{\bullet}(x, y) & \frac{xy}{x-1} \partial_y f_{\mathbf{q}}^{\bullet}(x, y) & 0 \end{pmatrix}$$

is at most 1. Moreover, a solution (x, y) with these properties is then unique.

If the spectral radius of $M(x, y)$ (for this unique solution (x, y)) equals 1, then we say that \mathbf{q} is critical. It is here even *regular* critical in the terminology of [19], since the functions $f_{\mathbf{q}}^{\bullet}, f_{\mathbf{q}}^{\diamond}$ are everywhere finite in our case. Note that the matrix $M(x, y)$ has non-negative coefficients, and the Perron–Frobenius theorem ensures that the spectral radius of $M(x, y)$ is also the largest real eigenvalue of $M(x, y)$. Thus, assuming that \mathbf{q} is admissible, and letting (x, y) be the unique solution of (23) such that $M(x, y)$ has spectral radius bounded by 1, we see that \mathbf{q} is regular critical if and only if 1 is an eigenvalue of $M(x, y)$, which holds if and only if the determinant of $\text{Id} - M(x, y)$ vanishes.

For every $x, y > 0$, set

$$G(x, y) = f_{\mathbf{q}}^{\bullet}(x, y) - 1 + 1/x \quad \text{and} \quad H(x, y) = f_{\mathbf{q}}^{\diamond}(x, y) - y.$$

Then G and H are convex functions in each variable $x, y \in (0, \infty)$. A pair $(x, y) \in (0, \infty)^2$ satisfies (23) if and only if $G(x, y) = H(x, y) = 0$ (notice that the condition $G(x, y) = 0$ forces $x > 1$). The set $\{G = 0\}$, resp. $\{H = 0\}$ is the boundary of the closed set $C_G = \{G \leq 0\}$, resp. of $C_H = \{H \leq 0\}$, in $(0, \infty)^2$.

Lemma A.1. (i) *The set C_G is contained in $(1, \infty) \times (0, A)$, for some $A > 0$.*

(ii) *The set C_H is bounded.*

(iii) *If $(x, y) \in C_G$ then $(x, y') \in C_G$ for every $y' \in (0, y)$. If $(x, y) \in C_H$ then $(x', y) \in C_H$ for every $x' \in (0, x)$. There exists $\varepsilon > 0$ such that C_H does not intersect $[1, \infty) \times (0, \varepsilon)$.*

(iv) *For every $a > 0$, let G_a , resp. H_a , be the function analogous to G , resp. to H , when \mathbf{q} is replaced by $a\mathbf{q}$. Then $C_{H_a} \subset (0, 1] \times (0, \infty)$ for every large enough $a > 0$. Consequently $C_{H_a} \cap C_{G_a} = \emptyset$ for every large enough $a > 0$.*

Proof. (i) This is obvious since $f_{\mathbf{q}}^{\bullet}(x, y) \geq Cy^{\ell}$ for every $x, y > 0$, for some constant $C > 0$ and some integer $\ell \geq 3$.

(ii) Suppose first that there exists an odd integer $\ell \geq 3$ such that $q_{\ell} > 0$. Then, the definition of $f_{\mathbf{q}}^{\diamond}$ shows that there is a positive constant c such that

$$f_{\mathbf{q}}^{\diamond}(x, y) \geq c(x^{(\ell-1)/2} + y^{\ell-1}),$$

and it readily follows that C_H is bounded. Consider then the case when there is an even integer $\ell \geq 4$ such that $q_{\ell} > 0$. Then there is a positive constant c such that

$$f_{\mathbf{q}}^{\diamond}(x, y) \geq c(x^{(\ell-2)/2}y + y^{\ell-1}),$$

and again this implies that C_H is bounded.

(iii) The first property is clear since $y \mapsto G(x, y)$ is non-decreasing, for every $x > 0$. Similarly, the second property in (iii) follows from the fact that $x \mapsto H(x, y)$ is non-decreasing, for every $y > 0$. The last property is also clear since we can find $\varepsilon > 0$ such that $f_{\mathbf{q}}^{\diamond}(x, y) > \varepsilon$ for every $x \geq 1$ and $y > 0$ (we use the fact that \mathbf{q} is not supported on even integers).

(iv) Suppose first that there exists an odd integer $\ell \geq 3$ such that $q_{\ell} > 0$. Using the same bound as in the proof of (ii), and noting that $f_{a\mathbf{q}}^{\diamond} = af_{\mathbf{q}}^{\diamond}$, we see that $H_a(x, y) \leq 0$ can only hold if

$$x^{(\ell-1)/2} + y^{\ell-1} \leq \frac{y}{ca}.$$

It is elementary to check that this implies $x \leq 1$ as soon as a is large enough. The case when there is an even integer $\ell \geq 4$ such that $q_{\ell} > 0$ is treated similarly using the bound stated in the proof of (ii). Finally the last assertion in (iv) follows by using (i). □

Recall that $f_{\mathbf{q}}^\bullet$ and $f_{\mathbf{q}}^\diamond$ are polynomials. It follows that the set $\{G = 0\}$ is either empty or a smooth curve depending on whether the set $\{G \leq 0\}$ is empty or not (a priori it could happen that $\{G = 0\} = \{G \leq 0\}$ is a singleton, but assertion (iii) in the previous lemma shows that this case does not occur). Similar properties hold for the set $\{H = 0\}$. A simple calculation also shows that

$$\det(\text{Id} - M(x, y)) = x^2 \det(\nabla G(x, y), \nabla H(x, y)). \tag{24}$$

Consequently, if we assume that (x, y) satisfies (23), the condition $\det(\text{Id} - M(x, y)) = 0$ will hold if and only if the curves $\{G = 0\}$ and $\{H = 0\}$ are tangent at (x, y) .

Lemma A.1(iii) implies that whenever non-empty, the curves $\{G = 0\}, \{H = 0\}$ are the graphs of two smooth functions $g : I \rightarrow (0, \infty)$ and $h : J \rightarrow (0, \infty)$ in the sense that

$$\{G = 0\} = \{(x, g(x)) : x \in I\}, \quad \{H = 0\} = \{(h(y), y) : y \in J\},$$

and moreover, $\{G \leq 0\}, \{H \leq 0\}$ are exactly the points below the graphs of these functions g, h . By (i) and (ii) in Lemma A.1, I is a subset of $(1, \infty)$ and J is a bounded subset of $(0, \infty)$. Moreover, the fact that $x \mapsto G(x, y)$ and $y \mapsto H(x, y)$ are convex functions implies that I and J are intervals, and that there exists two subintervals I' and J' of respectively I and J , such that $\inf I' = \inf I$, $\inf J' = \inf J$, and the functions g and h are non-decreasing on I' and J' respectively, and non-increasing on $I \setminus I'$ and $J \setminus J'$ respectively (we may have $J = J'$).

Since $\nabla G, \nabla H$ are orthogonal to level lines of G, H and pointing towards larger values of these functions, we deduce that for every $x \in I', y \in J'$, it holds that

$$\partial_x G(x, g(x)) \leq 0, \quad \partial_y G(x, g(x)) \geq 0, \quad \partial_x H(h(y), y) \geq 0, \quad \partial_y H(h(y), y) \leq 0.$$

For $x \in I \setminus I', y \in J \setminus J'$, the first and last inequality should be reversed, the other two remaining as such. We now claim that the restrictions of g, h to the intervals I', J' are concave. To check this, we differentiate twice the equation $G(x, g(x)) = 0$ to obtain

$$g''(x) = -\frac{1}{\partial_y G} (\partial_{xx} G + 2g'(x) \partial_{xy} G + g'(x)^2 \partial_{yy} G),$$

where in the right-hand side all derivatives of G should be evaluated at the point $(x, g(x))$. Since $\partial_{xx} G, \partial_{xy} G, \partial_{yy} G$ are respectively equal to $\partial_{xx} f_{\mathbf{q}}^\bullet + 2/x^3, \partial_{xy} f_{\mathbf{q}}^\bullet, \partial_{yy} f_{\mathbf{q}}^\bullet$, which are all non-negative, we obtain that $g''(x) \leq 0$ for $x \in I'$, as claimed. The proof of the concavity of h on J' is similar.

Let us now consider possible intersections of $\{G = 0\}$ and $\{H = 0\}$. First of all, the monotonicity properties of g, h imply that if the graphs $\{(x, g(x)) : x \in I'\}$ and $\{(h(y), y) : y \in J'\}$ do not intersect, then $\{G = 0\} \cap \{H = 0\}$ is empty as well, so that \mathbf{q} cannot be admissible. Next, consider a point $(x, y) \in \{G = 0\} \cap \{H = 0\}$ and such that $y \in J \setminus J'$. Then we have seen that $\partial_y H(x, y) \geq 0$. Since $\partial_y H = \partial_y f_{\mathbf{q}}^\diamond - 1$, this means that $\partial_y f_{\mathbf{q}}^\diamond(x, y) \geq 1$. But $\partial_y f_{\mathbf{q}}^\diamond(x, y)$ is a diagonal coefficient of the irreducible matrix $M(x, y)$, so that it is a strict lower bound for its spectral radius. The last property is easy to show from the well-known fact that a Perron eigenvector (one corresponding to the largest, real eigenvalue) has only positive entries. We deduce that such an intersection between $\{G = 0\}$ and $\{H = 0\}$ cannot correspond to a ‘‘good’’ solution such that $M(x, y)$ has spectral radius at most 1. Finally, we see that such a good solution (x, y) can correspond only to a solution with $y \in J'$, and from the concavity of h on this interval and the monotonicity properties of g , we see that there are at most two such solutions. Of these, the smallest one will also have $x \in I'$, and the largest one, if different from the smallest one, corresponds to a point where g crosses the curve $\{(h(y), y) : y \in J'\}$ from above. In such a situation, the property $\det(\nabla G, \nabla H) < 0$ must hold, so 24 implies that the characteristic polynomial of $M(x, y)$ has a root strictly larger than 1.

Finally, we see that there can be only one solution (x, y) to the equations $G = H = 0$ for which $M(x, y)$ has spectral radius at most 1. If it exists, this solution then satisfies $(x, y) \in I' \times J'$, and is the solution to $G = H = 0$ which has minimal coordinates, in the sense that if (x', y') is another solution then $x < x'$ and $y < y'$. Moreover, if $M(x, y)$ has spectral radius equal to 1, then the curves are tangent at this solution, which is then also the unique solution to $G = H = 0$. With this at hand, we can finally prove the following result.

Proposition A.2. *Under Assumption (A1), there exists a unique positive real a_c such that $a_c \mathbf{q}$ is regular critical.*




Fig. A.1. Illustration of the sets C_{G_a} and C_{H_a} for $0 < a < a_c$ and for $a = a_c$.

Proof. For every $a > 0$, write $M_a(x, y)$ for the analogue of the matrix $M(x, y)$ when \mathbf{q} is replaced by $a\mathbf{q}$. Simple counting arguments (using for instance the BDG bijections and the fact that the sequence \mathbf{q} has finite support, so that the degrees of faces in maps m such that $W_{\mathbf{q}}(m) > 0$ are bounded) show that the Boltzmann measure $W_{a\mathbf{q}}$ is finite for $a > 0$ small enough. Consequently we can fix $a_0 > 0$ small enough so that $a_0\mathbf{q}$ is admissible. By previous observations, there exists a pair (x_{a_0}, y_{a_0}) belonging to the intersection of the curves $\{G_{a_0} = 0\}$ and $\{H_{a_0} = 0\}$ and such that the spectral radius of the matrix $M_{a_0}(x_{a_0}, y_{a_0})$ is bounded above by 1. This pair is also the intersection point with minimal coordinates. If the curves $\{G_{a_0} = 0\}$ and $\{H_{a_0} = 0\}$ are tangent at (x_{a_0}, y_{a_0}) , then (24) shows that this spectral radius is equal to 1, and thus $a_0\mathbf{q}$ is regular critical.

Suppose that the curves $\{G_{a_0} = 0\}$ and $\{H_{a_0} = 0\}$ are not tangent at (x_{a_0}, y_{a_0}) . Note that both sets C_{G_a} and C_{H_a} are decreasing functions of a , and vary continuously with a (as long as they are non-empty). Geometric arguments, together with property (iv) of Lemma A.1, show that there exists a critical value $a_c > a_0$ such that for $a_0 \leq a < a_c$ the curves $\{G_a = 0\}$ and $\{H_a = 0\}$ intersect in at least two points, the one with minimal coordinates being denoted by (x_a, y_a) , and furthermore the curves $\{G_{a_c} = 0\}$ and $\{H_{a_c} = 0\}$ are tangent at a point denoted by (x_{a_c}, y_{a_c}) . See Figure A.1 for an illustration. Moreover the mapping $a \mapsto (x_a, y_a)$ is continuous on $[a_0, a_c]$. It follows that the spectral radius of $M_a(x_a, y_a)$ remains bounded above by 1 for $a \in [a_0, a_c]$: If this were not the case, this spectral radius would take the value 1 at some $a_1 \in (a_0, a_c)$ but then by (24) the curves $\{G_{a_1} = 0\}$ and $\{H_{a_1} = 0\}$ would be tangent at (x_{a_1}, y_{a_1}) , which is a contradiction. Finally by letting $a \uparrow a_c$ we get that the spectral radius of $M_{a_c}(x_{a_c}, y_{a_c})$ is bounded above by 1, hence equal to 1 by (24) and the fact that $\{G_{a_c} = 0\}$ and $\{H_{a_c} = 0\}$ are tangent at (x_{a_c}, y_{a_c}) . We conclude that $a_c\mathbf{q}$ is regular critical.

The uniqueness of a_c is clear since we can start the previous argument from an arbitrarily small value of a_0 and since the curves $\{G_a = 0\}$ and $\{H_a = 0\}$ will not intersect when $a > a_c$. \square

Acknowledgement

We thank Itai Benjamini for the name cactus as well as for suggesting the study of this mathematical object.

References

- [1] D. Aldous. The continuum random tree I. *Ann. Probab.* **19** (1991) 1–28. [MR1085326](#)
- [2] J. Ambjørn, B. Durhuus and T. Jonsson. *Quantum Geometry. A Statistical Field Theory Approach*. *Cambridge Monographs on Mathematical Physics*. Cambridge Univ. Press, Cambridge, 1997. [MR1465433](#)

- [3] J. Bouttier, P. Di Francesco and E. Guitter. Planar maps as labeled mobiles. *Electron. J. Combin.* **11** (2004) R69. [MR2097335](#)
- [4] J. Bouttier and E. Guitter. Confluence of geodesic paths and separating loops in large planar quadrangulations. *J. Stat. Mech. Theory Exp.* (2009) P03001. [MR2495866](#)
- [5] D. Burago, Y. Burago and S. Ivanov. *A Course in Metric Geometry. Graduate Studies in Mathematics* **33**. Amer. Math. Soc., Boston, 2001. [MR1835418](#)
- [6] S. N. Evans. *Probability and Real Trees. Lecture Notes in Math.* **1920**. Springer, Berlin, 2008. [MR2351587](#)
- [7] S. N. Evans, J. W. Pitman and A. Winter. Rayleigh processes, real trees, and root growth with re-grafting. *Probab. Theory Related Fields* **134** (2006) 81–126. [MR2221786](#)
- [8] C. Favre and M. Jonsson. *The Valuative Tree. Lecture Notes in Math.* **1853**. Springer, Berlin, 2004. [MR2097722](#)
- [9] M. Gromov. *Metric Structures for Riemannian and Non-Riemannian Spaces*. Birkhäuser, Boston, 2001. [MR1699320](#)
- [10] N. C. Jain and S. J. Taylor. Local asymptotic laws for Brownian motion. *Ann. Probab.* **1** (1973) 527–549. [MR0365732](#)
- [11] S. K. Lando and A. K. Zvonkin. *Graphs on Surfaces and Their Applications. Encyclopedia of Mathematical Sciences* **141**. Springer, Berlin, 2004. [MR2036721](#)
- [12] J. F. Le Gall. Random trees and applications. *Probab. Sur.* **2** (2005) 245–311. [MR2203728](#)
- [13] J. F. Le Gall. The topological structure of scaling limits of large planar maps. *Invent. Math.* **169** (2007) 621–670. [MR2336042](#)
- [14] J. F. Le Gall. Geodesics in large planar maps and in the Brownian map. *Acta Math.* **205** (2010) 287–360. [MR2746349](#)
- [15] J. F. Le Gall. Uniqueness and universality of the Brownian map. Preprint. Available at arXiv:[1105.4842](#).
- [16] J. F. Le Gall and F. Paulin. Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere. *Geomet. Funct. Anal.* **18** (2008) 893–918. [MR2438999](#)
- [17] J. F. Marckert and G. Miermont. Invariance principles for random bipartite planar maps. *Ann. Probab.* **35** (2007) 1642–1705. [MR2349571](#)
- [18] J. F. Marckert and A. Mokkadem. Limit of normalized quadrangulations. The Brownian map. *Ann. Probab.* **34** (2006) 2144–2202. [MR2294979](#)
- [19] G. Miermont. An invariance principle for random planar maps. In *Fourth Colloquium on Mathematics and Computer Science, Algorithms, Trees, Combinatorics and Probabilities* 39–57 (electronic). Discrete Math. Theor. Comput. Sci. Proc., AG. Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2006. [MR2509622](#)
- [20] G. Miermont. Invariance principles for spatial multitype Galton–Watson trees. *Ann. Inst. Henri Poincaré Probab. Stat.* **44** (2008) 1128–1161. [MR2469338](#)
- [21] G. Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. Preprint. Available at arXiv:[1104.1606](#).
- [22] G. Miermont and M. Weill. Radius and profile of random planar maps with faces of arbitrary degrees. *Electron. J. Probab.* **13** (2008) 79–106. [MR2375600](#)
- [23] D. Revuz and M. Yor. *Continuous Martingales and Brownian Motion*. Springer, Berlin, 1991. [MR1083357](#)
- [24] O. Schramm. Conformally invariant scaling limits: An overview and a collection of problems. In *Proceedings of the International Congress of Mathematicians (Madrid, 2006)*, Vol. I 513–543. European Math. Soc., Zürich, 2007. [MR2334202](#)
- [25] W. T. Tutte. A census of planar maps. *Canad. J. Math.* **15** (1963) 249–271. [MR0146823](#)