

Infinite geodesics in hyperbolic random triangulations

Thomas BUDZINSKI *

April 29, 2019

Abstract

We study the structure of infinite geodesics in the Planar Stochastic Hyperbolic Triangulations \mathbb{T}_λ introduced in [16]. We prove that these geodesics form a supercritical Galton–Watson tree with geometric offspring distribution. The tree of infinite geodesics in \mathbb{T}_λ provides a new notion of boundary, which is a realization of the Poisson boundary. By scaling limit arguments, we also obtain a description of the tree of infinite geodesics in the hyperbolic Brownian plane. Finally, by combining our main result with a forthcoming paper [15], we obtain new hyperbolicity properties of \mathbb{T}_λ : they satisfy a weaker form of Gromov-hyperbolicity and admit bi-infinite geodesics.

Introduction

The construction and study of random infinite triangulations has been a very active field of research in recent years. The first such triangulation that was built is the UIPT [9, 3]. A key feature in the study of this object is its spatial Markov property, which motivated the introduction of a one-parameter family $(\mathbb{T}_\lambda)_{0 < \lambda \leq \lambda_c}$ of type-I triangulations¹ with $\lambda_c = \frac{1}{12\sqrt{3}}$, satisfying a similar property [16, 14] (see also [7] for similar constructions in the halfplanar case). The case $\lambda = \lambda_c$ corresponds to the UIPT, whereas for $\lambda < \lambda_c$ the triangulation \mathbb{T}_λ has hyperbolic behaviour. For example, it was proved that \mathbb{T}_λ has a.s. exponential volume growth and that the simple random walk on it has positive speed [16]. The goal of this work is to establish hyperbolicity properties of these maps related to their geodesics.

Leftmost geodesic rays. Our first goal in this work is to describe precisely the structure of infinite geodesics in the triangulations \mathbb{T}_λ . More precisely, all the triangulations considered here are *rooted*, that is, equipped with a distinguished oriented edge called the *root edge*. The *root vertex*, that we write ρ , is the starting point of the root edge. For any two vertices x and y in \mathbb{T}_λ , we call a geodesic γ from x to y *leftmost* if for any geodesic γ' from x to y , the union of γ and γ' cuts \mathbb{T}_λ in two parts, and the part on the left of γ is infinite. It is easy to see that for any vertices x and y there is a unique leftmost geodesic from x to y . A *leftmost geodesic ray* is a sequence of vertices $(\gamma(n))_{n \geq 0}$ such that $\gamma(0) = \rho$ and for any $n \geq 0$, the path $(\gamma(i))_{0 \leq i \leq n}$ is a leftmost geodesic from ρ to $\gamma(n)$. We denote by \mathbf{T}_λ^g the union of all the leftmost geodesic rays of \mathbb{T}_λ . We can see this set of vertices as a graph by relating two vertices if they are consecutive on the same geodesic ray. By uniqueness of the leftmost geodesic between two points, the graph \mathbf{T}_λ^g is an infinite tree with no leaf. Moreover, the tree \mathbf{T}_λ^g divides \mathbb{T}_λ into infinite maps with geodesic boundaries, that we call *strips* (see the left part of Figure 1). Our first main result describes the distribution of \mathbf{T}_λ^g and of these strips. Let $0 < \lambda \leq \lambda_c$. Let $h \in (0, \frac{1}{4}]$ be such that

$$\lambda = \frac{h}{(1 + 8h)^{3/2}}, \tag{1}$$

*ENS Paris and Université Paris-Saclay, thomas.budzinski@ens.fr

¹To be exact, type-II triangulations (i.e. with no loops) were considered in [16], while the type-I triangulations (with loops) were built in [14]. In this work, unless specified otherwise, we only consider type-I triangulations.

and let

$$m_\lambda = \frac{1 - 2h - \sqrt{1 - 4h}}{2h} \leq 1. \quad (2)$$

Theorem 1. • The tree \mathbf{T}_λ^g is a Galton–Watson tree with offspring distribution μ_λ , where $\mu_\lambda(0) = 0$ and $\mu_\lambda(k) = m_\lambda(1 - m_\lambda)^{k-1}$ for $k \geq 1$.

- There are two random infinite strips S_λ^0 and S_λ^1 such that the following holds. Conditionally on \mathbf{T}_λ^g :
 1. the strips delimited by \mathbf{T}_λ^g are independent,
 2. the strip containing the face lying on the right of the root edge has the same distribution as S_λ^1 ,
 3. all the other strips have the same distribution as S_λ^0 .

Note that for $\lambda = \lambda_c$, we have $h = \frac{1}{4}$ so $m_{\lambda_c} = 1$ and $\mu_{\lambda_c}(1) = 1$, so the tree $\mathbf{T}_{\lambda_c}^g$ consists of a single ray. This is reminiscent of the geodesics confluence properties already observed in [20] for the UIPQ (the natural analog of \mathbb{T}_{λ_c} for quadrangulations): there are infinitely many points that lie on every geodesic ray. See also [21] for similar results in the UIPT. However, the results of [20] are about all the geodesic rays whereas we only study the leftmost ones, so our result for $\lambda = \lambda_c$ does not obviously imply those of [20]. On the other hand, for $\lambda < \lambda_c$, we have $m_\lambda < 1$, so the offspring distribution μ_λ is supercritical and there are infinitely many leftmost geodesic rays. Moreover, the rate of exponential growth μ_λ of \mathbf{T}_λ^g is the same as the rate of exponential volume growth of \mathbb{T}_λ .

We will also describe the distributions of S_λ^0 and S_λ^1 explicitly in terms of reverse Galton–Watson trees. For $\lambda < \lambda_c$, these strips should be thought of as "thin", in the sense that their width is of constant order as the distance from the root goes to $+\infty$.

We also state right now a consequence of Theorem 1 that will be useful later. Let γ_ℓ (resp. γ_r) be the path in \mathbf{T}_λ^g that bounds S_λ^1 on its right (resp. on its left). Then γ_ℓ (resp. γ_r) can be thought of as the leftmost (resp. rightmost) path in \mathbf{T}_λ^g , seen from the root (see Figure 1). We write \mathbb{S}_λ for the part of \mathbb{T}_λ lying between γ_ℓ and γ_r , including the initial segment that γ_ℓ and γ_r have in common (cf. Figure 1). Then \mathbb{S}_λ can be seen as a gluing of infinitely many independent copies of S_λ^0 along \mathbf{T}_λ^g . This implies that \mathbb{S}_λ has an interesting self-similarity property. Indeed, let $r > 0$. We condition on $B_r(\mathbf{T}_\lambda^g)$, the finite subtree of \mathbf{T}_λ^g formed by those vertices lying at distance at most r from ρ . Let x be a vertex of \mathbf{T}_λ^g such that $d(\rho, x) = r$. Let γ_ℓ^x (resp. γ_r^x) be the leftmost (resp. rightmost) infinite path in \mathbf{T}_λ^g started from ρ and passing through x . Then the part of \mathbb{T}_λ lying between γ_ℓ^x and γ_r^x above x has the same distribution as \mathbb{S}_λ (see the right part of Figure 1). Indeed, this part consists of a gluing of i.i.d. copies of S_λ^0 in the faces of the tree of descendants of x in \mathbf{T}_λ^g , which has the same distribution (conditionally on $B_r(\mathbf{T}_\lambda^g)$) as \mathbf{T}_λ^g . We will denote this part of \mathbb{T}_λ by $\mathbb{S}[x]$. In all the rest of this work, "thin" maps such as S_λ^0 and S_λ^1 will be referred to as *strips*, and "thick" maps like \mathbb{S}_λ as *slices*.

Hyperbolicity properties related to geodesics. By Theorem 1, for $\lambda < \lambda_c$, the map \mathbb{T}_λ contains a supercritical Galton–Watson tree. By combining this with the results of our forthcoming paper [15], we obtain that \mathbb{T}_λ satisfies two metric hyperbolicity properties: a weak form of Gromov-hyperbolicity, and the existence of bi-infinite geodesics.

More precisely, we recall that a graph G is hyperbolic in the sense of Gromov if there is a constant $k \geq 0$ such that all the triangles are k -thin in the following sense. Let x, y and z be vertices of G and $\gamma_{xy}, \gamma_{yz}, \gamma_{zx}$ be geodesics from x to y , from y to z and from z to x . Then for any vertex v on γ_{xy} , the graph distance between v and $\gamma_{yz} \cup \gamma_{zx}$ is at most k . As usual, such a strong, uniform statement cannot hold for \mathbb{T}_λ since any finite triangulation appears somewhere in \mathbb{T}_λ . Therefore, we need to study an "anchored" version.




Figure 1: The strip decomposition of \mathbb{T}_λ . On the left, the tree \mathbf{T}_λ^g is in red. On the right $\mathbb{S}_\lambda = \mathbb{S}[\rho]$ is colored in yellow. The part $\mathbb{S}[x]$ of \mathbb{S}_λ that is shaded in green has the same distribution as \mathbb{S}_λ .

Definition 0.1. Let M be a planar map. We say that M is *weakly anchored hyperbolic* if there is $k \geq 0$ such that the following holds. Let x, y and z be three vertices of M , and let γ_{xy} (resp. γ_{yz}, γ_{zx}) be a geodesic from x to y (resp. y to z, z to x). Assume the triangle formed by γ_{xy}, γ_{yz} and γ_{zx} surrounds the root vertex ρ . Then

$$d(\rho, \gamma_{xy} \cup \gamma_{yz} \cup \gamma_{zx}) \leq k.$$

Another property studied in [15] is the existence of bi-infinite geodesics, i.e. paths $(\gamma(i))_{i \in \mathbb{Z}}$ such that for any i and j , the graph distance between $\gamma(i)$ and $\gamma(j)$ is exactly $|i - j|$. This is not strictly speaking a hyperbolicity property, since such geodesics exist e.g. in \mathbb{Z}^2 . However, they are expected to disappear after perturbations of the metric like first-passage percolation (see for example [25]). On the other hand, bi-infinite geodesics are much more stable in hyperbolic graphs [12]. In [15], we prove that any random planar map containing a supercritical Galton–Watson tree with no leaf is a.s. weakly anchored hyperbolic, and contains bi-infinite geodesics. In particular, the following result follows from Theorem 1.

Corollary 0.2. Let $0 < \lambda < \lambda_c$. Almost surely, the map \mathbb{T}_λ is weakly anchored hyperbolic and admits bi-infinite geodesics.

The existence of bi-infinite geodesics answers a question of Benjamini and Tessera [12]. Once again, there is a sharp contrast between the hyperbolic setting and "usual" random planar maps. For example, the results of [20] imply that such bi-infinite geodesics do not exist in the UIPQ.

Poisson boundary. Another goal of this work is to give a new description of the Poisson boundary of \mathbb{T}_λ for $0 < \lambda < \lambda_c$ in terms of the tree \mathbf{T}_λ^g . Let G be an infinite, locally finite graph, and let $G \cup \partial G$ be a compactification of G , i.e. a compact metric space in which G is dense. Let also (X_n) be the simple random walk on G . We say that ∂G is a *realization of the Poisson boundary* of G if the following two properties hold:

- (X_n) converges a.s. to a point $X_\infty \in \partial G$,
- every bounded harmonic function h on G can be written in the form

$$h(x) = \mathbb{E}_x [g(X_\infty)],$$

where g is a bounded measurable function from ∂G to \mathbb{R} .

A first realization of the Poisson boundary of \mathbb{T}_λ is given by a work of Angel, Hutchcroft, Nachmias and Ray [4]: let $\partial_{CP}\mathbb{T}_\lambda$ be the boundary of the circle packing of \mathbb{T}_λ in the unit disk.

We may equip $\mathbb{T}_\lambda \cup \partial_{CP}\mathbb{T}_\lambda$ with the topology induced by the usual topology on the unit disk. Then almost surely, $\partial_{CP}\mathbb{T}_\lambda$ is a realization of the Poisson boundary of \mathbb{T}_λ . Moreover, almost surely, the distribution of the limit point X_∞ has full support and no atoms in $\partial_{CP}\mathbb{T}_\lambda$.

We write $\partial\mathbf{T}_\lambda^g$ for the space of infinite rays of \mathbf{T}_λ^g . If $\gamma, \gamma' \in \partial\mathbf{T}_\lambda^g$, we write $\gamma \sim \gamma'$ if $\gamma = \gamma'$ or if γ and γ' are the left and right boundaries of the same strip. Then \sim is a.s. an equivalence relation in which countably many equivalence classes have cardinal 2, and all the others have cardinal 1. We write $\widehat{\partial\mathbf{T}}_\lambda^g = \partial\mathbf{T}_\lambda^g / \sim$. There is a natural way to equip $\mathbb{T}_\lambda \cup \widehat{\partial\mathbf{T}}_\lambda^g$ with a topology that makes it a compact space, see Section 3.1 for more details. Hence, $\mathbb{T}_\lambda \cup \widehat{\partial\mathbf{T}}_\lambda^g$ can be seen as a compactification of the infinite graph \mathbb{T}_λ . We show that $\widehat{\partial\mathbf{T}}_\lambda^g$ is also a realization of the Poisson boundary.

Theorem 2. Let $0 < \lambda < \lambda_c$. Then almost surely:

1. the limit $\lim X_n = X_\infty$ exists, and its distribution has full support and no atoms in $\widehat{\partial\mathbf{T}}_\lambda^g$,
2. $\widehat{\partial\mathbf{T}}_\lambda^g$ is a realization of the Poisson boundary of \mathbb{T}_λ .

Note that, by a result of Hutchcroft and Peres [24], the second point will follow from the first one. Note also that once we have Theorem 2, since the exit measure on $\widehat{\partial\mathbf{T}}_\lambda^g$ is nonatomic and only countably many pairs of points of $\partial\mathbf{T}_\lambda^g$ are identified, we also have almost sure convergence of (X_n) to a point of $\partial\mathbf{T}_\lambda^g$. Finally, we show in [15] that the existence of the limit X_∞ and the fact that it has full support are true in the more general setting of planar maps obtained by "filling the faces" of a supercritical Galton–Watson tree with i.i.d. strips. However, we did not manage to prove non-atomicity and to obtain a precise description of the Poisson boundary in this general setting. Our proof of non-atomicity here uses an argument specific to \mathbb{T}_λ , based on its peeling process.

Geodesic rays in the hyperbolic Brownian plane. Finally, the last goal of this work is to take the scaling limit of Theorem 1 to obtain results about geodesics in continuum objects. Indeed, another purpose of the theory of random planar maps is to build continuum random surfaces. The first such surface that was introduced is the Brownian map [31, 34], which is now known to be the scaling limit of a wide class of finite planar maps conditioned to be large [1, 2, 11, 13, 18, 33]. A noncompact version \mathcal{P} called the Brownian plane was introduced in [17] and is the scaling limit of the UIPQ and also of the UIPT [14]. Finally, it was shown in [14] that the hyperbolic random triangulations have a near-critical scaling limit called the hyperbolic Brownian plane and denoted \mathcal{P}^h . More precisely, let (λ_n) be a sequence of numbers in $(0, \lambda_c]$ satisfying

$$\lambda_n = \lambda_c \left(1 - \frac{2}{3n^4} + o\left(\frac{1}{n^4}\right) \right).$$

Then $\frac{1}{n}\mathbb{T}_{\lambda_n}$ converges for the local Gromov–Hausdorff distance to \mathcal{P}^h . By taking the scaling limit of Theorem 1 and checking that the tree of infinite leftmost geodesics behaves well in the scaling limit, we obtain a precise description of the geodesic rays in \mathcal{P}^h . Let \mathbf{B} be the infinite tree in which every vertex has exactly two children, except the root which has only one.

Theorem 3. The infinite geodesic rays of \mathcal{P}^h form a tree $\mathbf{T}^g(\mathcal{P}^h)$ that is distributed as a Yule tree with parameter $2\sqrt{2}$, i.e. the tree \mathbf{B} in which the lengths of the edges are i.i.d. exponential variables with parameter $2\sqrt{2}$.

Once again, this behaviour is very different from the non-hyperbolic setting: in the Brownian plane, there is only one geodesic ray (this is Proposition 15 of [17], and is an easy consequence of the local confluence of geodesics proved in [30] for the Brownian map). We also note that the rate $2\sqrt{2}$ of exponential growth of $\mathbf{T}^g(\mathcal{P}^h)$ is the same as the rate of exponential growth of the perimeters and volumes of the hulls of \mathcal{P}^h [14, Corollary 1].

The skeleton decomposition. Our main tool for proving these results will be the skeleton decomposition of planar triangulations introduced by Krikun [28] for the type-II UIPT. See also [18] for the adaptation to the (slightly easier) type-I case. This decomposition encodes a triangulation by a reverse forest, where leftmost geodesics from the root pass between the trees and between their branches. The infinite forest describing the UIPT consists of a single tree, which can be seen as a reverse Galton–Watson tree with critical offspring distribution, started at time $-\infty$, and conditioned to have exactly 1 vertex at time 0 and to die at time 1. We obtain a similar description for the infinite forest encoding \mathbb{T}_λ for $0 < \lambda < \lambda_c$, but here the offspring distribution is subcritical. A key feature is that the forest now contains infinitely many infinite trees. The parts of \mathbb{T}_λ described by each of these trees are the strips delimited by the leftmost geodesic rays. Finally, let us highlight that the construction of infinite reverse Galton–Watson trees that we present in Section 2.6 holds for any subcritical or critical Galton–Watson process and might be of independent interest.

Structure of the paper. The structure of the paper is as follows. In Section 1, we recall some combinatorial results about planar triangulations, and we recall the definition and some basic properties of the maps \mathbb{T}_λ and their halfplanar analogs. In Section 2, we prove Theorem 1 by computing the skeleton decomposition of \mathbb{T}_λ . Section 3 is devoted to the proof of Theorem 2, and Section 4 to the proof of Theorem 3. In Appendix A, we prove a technical result needed in Section 4, which shows that a wide class of events related to geodesics are closed for the Gromov–Hausdorff distance.

Acknowledgments: I thank Nicolas Curien for carefully reading several earlier versions of this manuscript, and Itai Benjamini for his question about bi-infinite geodesics. I am also grateful to the anonymous reviewers for their useful comments. I acknowledge the support of ANR Liouville (ANR-15-CE40-0013) and ANR GRAAL (ANR-14-CE25-0014).

Contents

1	Combinatorics and preliminaries	6
1.1	Combinatorics	6
1.2	Planar and halfplanar hyperbolic type-I triangulations	6
2	The skeleton decomposition of hyperbolic triangulations	9
2.1	The skeleton decomposition of finite and infinite triangulations	9
2.2	Computation of the skeleton decomposition of the hulls of \mathbb{T}_λ	12
2.3	Slicing the skeleton	13
2.4	The distribution of S_λ^0 and S_λ^1	15
2.5	Proof of Theorem 1	17
2.6	Construction of reverse Galton–Watson trees and infinite strips	20
3	The Poisson boundary of \mathbb{T}_λ	24
3.1	Construction of the geodesic boundary	24
3.2	Proof of Theorem 2	26
4	The tree of infinite geodesics in the hyperbolic Brownian plane	30
4.1	The tree $\mathbf{T}^g(\mathcal{P}^h)$	30
4.2	Two lemmas about near-critical strips	31
4.3	Identification of the geodesic tree via Gromov–Hausdorff-closed events	34
A	A Gromov–Hausdorff closedness result	37

1 Combinatorics and preliminaries

1.1 Combinatorics

For $n \geq 0$ and $p \geq 1$, a *triangulation of the p -gon with n inner vertices* is a planar map with $n+p$ vertices and a distinguished face called the *outer face*, in which all faces except perhaps the outer face are triangles, and such that the boundary of the outer face is a simple cycle of length p . It is equipped with a root edge such that the outer face touches the root edge on its right. We consider type-I triangulations, which means we allow triangulations containing loops and multiple edges. We denote by $\mathcal{T}_{n,p}$ the set of triangulations of the p -gon with n inner vertices.

The number of triangulations with fixed volume and perimeter can be computed by a result of Krikun, as a special case of the main theorem of [29]:

$$\#\mathcal{T}_{n,p} = \frac{p(2p)!}{(p!)^2} \frac{4^{n-1}(2p+3n-5)!!}{n!(2p+n-1)!!} \underset{n \rightarrow +\infty}{\sim} c(p)\lambda_c^{-n}n^{-5/2}, \quad (3)$$

where $\lambda_c = \frac{1}{12\sqrt{3}}$ and

$$c(p) = \frac{3^{p-2}p(2p)!}{4\sqrt{2\pi}(p!)^2} \underset{p \rightarrow +\infty}{\sim} \frac{1}{36\pi\sqrt{2}} 12^p \sqrt{p}. \quad (4)$$

For $p \geq 1$ and $\lambda \geq 0$, we write $w_\lambda(p) = \sum_{n \geq 0} \#\mathcal{T}_{n,p} \lambda^n$. Note that by the asymptotics (3), we have $w_\lambda(p) < +\infty$ if and only if $\lambda \leq \lambda_c$. We finally write $W_\lambda(x) = \sum_{p \geq 1} w_\lambda(p)x^p$. Formula (4) of [29] computes W_λ after a simple change of variables:

$$W_\lambda(x) = \frac{\lambda}{2} \left(\left(1 - \frac{1+8h}{h}x \right) \sqrt{1 - 4(1+8h)x} - 1 + \frac{x}{\lambda} \right), \quad (5)$$

where $h \in (0, \frac{1}{4}]$ is given² by (1). From this formula, we easily get

$$w_\lambda(1) = \frac{1}{2} - \frac{1+2h}{2\sqrt{1+8h}} \quad (6)$$

and, for $p \geq 2$,

$$w_\lambda(p) = (2+16h)^p \frac{(2p-5)!!}{p!} \frac{(1-4h)p+6h}{4(1+8h)^{3/2}}. \quad (7)$$

We also define a *Boltzmann triangulation of the p -gon with parameter λ* as a random triangulation T such that $\mathbb{P}(T=t) = \frac{\lambda^n}{w_\lambda(p)}$ for every $n \geq 0$ and $t \in \mathcal{T}_{n,p}$.

1.2 Planar and halfplanar hyperbolic type-I triangulations

We recall from [14] the definition of the random triangulations \mathbb{T}_λ for $0 < \lambda \leq \lambda_c$. A *finite triangulation with a hole of perimeter p* is a rooted map with a distinguished face called the *hole* in which all the faces are triangles except perhaps the hole, and where the boundary of the hole is a simple cycle of length p . The difference with a triangulation of the p -gon is that we do not require the root to lie on the boundary. Let t be a finite triangulation with a hole of perimeter p and let T be an infinite, one-ended triangulation of the plane. We write $t \subset T$ if T can be obtained by filling the hole of t with an infinite triangulation of the p -gon. For $0 < \lambda \leq \lambda_c$, the distribution of \mathbb{T}_λ can be characterized as follows. For any finite triangulation t with a hole of perimeter p , we have

$$\mathbb{P}(t \subset \mathbb{T}_\lambda) = c_\lambda(p)\lambda^{|t|},$$

where $|t|$ is the total number of vertices of t and

$$c_\lambda(p) = \frac{1}{\lambda} \left(8 + \frac{1}{h} \right)^{p-1} \sum_{q=0}^{p-1} \binom{2q}{q} h^q, \quad (8)$$

²Note that our h corresponds to the h^3 of Krikun.

where h is as in (1). Equivalently, we can compute the generating function

$$C_\lambda(x) = \sum_{p \geq 1} c_\lambda(p) x^p = \frac{x}{\lambda \left(1 - \frac{1+8h}{h} x\right) \sqrt{1 - 4(1+8h)x}}. \quad (9)$$

Note that the numbers $c_{\lambda_c}(p)$ are equal to the $c(p)$ defined by (4) and \mathbb{T}_{λ_c} corresponds to the type-I UIPT [18, 36]. As in the type-II case [16], the triangulations \mathbb{T}_λ exhibit a spatial Markov property similar to that of the UIPT: if t is a finite triangulation with a hole of perimeter p , conditionally on $t \subset \mathbb{T}_\lambda$, the distribution of the infinite triangulation that fills the hole of t only depends on p . We denote by \mathbb{T}_λ^p a triangulation with this distribution. By a simple root transformation (more precisely duplicating the root edge, adding a loop in between and rooting the map at this new loop, see Figure 2 of [18]), triangulations of the plane are equivalent to infinite triangulations of the 1-gon. In particular, the image of \mathbb{T}_λ under this root transformation is \mathbb{T}_λ^1 , so studying one or the other are equivalent. In particular, the root transformation does not affect leftmost geodesics from the root, so all the results we will first obtain about geodesic rays in \mathbb{T}_λ^1 are immediate to transfer to \mathbb{T}_λ .

We also define the halfplanar analog of \mathbb{T}_λ , that we will denote by \mathbb{H}_λ . The goal of the next paragraphs is to adapt to the type-I setting results from [7] and [16] that are already well-known in the type-II setting. This will only be used in Section 3 to study the random walk on \mathbb{T}_λ , so all the rest of Section 1 can be skipped in the first read-through. A *triangulation of the halfplane* is an infinite planar map in which all the faces are triangles except one called the *outer face*, whose boundary is simple and infinite. Triangulations of the halfplane are rooted in such a way that the root edge touches the outer face on its right. We note that Angel and Ray build in [7] a family $(\mathbb{H}_\alpha^{\text{II}})_{2/3 < \alpha < 1}$ of hyperbolic triangulations of the halfplane in the type-II setting and explain in Section 3.4 how to "add loops" to obtain type-I triangulations. The triangulations \mathbb{H}_λ we define are a particular case of their construction. However, as in [14] and in order to limit the computations, we prefer to construct the maps \mathbb{H}_λ directly instead of relying on the type-II case.

The law of \mathbb{H}_λ is characterized by the following. Let t be a triangulation of the p -gon with a marked segment of edges $\partial_{\text{out}} t$ on its boundary, such that $\partial_{\text{out}} t$ contains the root edge. Such triangulations will be called *marked triangulations*. Let $\partial_{\text{in}} t = \partial t \setminus \partial_{\text{out}} t$. We write $|t|_{\text{in}}$ for the number of vertices of t that do not lie on $\partial_{\text{out}} t$. We also write $|\partial_{\text{in}} t|$ (resp. $|\partial_{\text{out}} t|$) for the number of edges on $\partial_{\text{in}} t$ (resp. $\partial_{\text{out}} t$). We write $t \subset \mathbb{H}_\lambda$ if \mathbb{H}_λ can be obtained by gluing a triangulation of the halfplane H to t in such a way that $\partial t \cap \partial H = \partial_{\text{in}} t$.

Proposition 1.1. For any $0 < \lambda \leq \lambda_c$, there is a unique (in distribution) random triangulation \mathbb{H}_λ of the halfplane such that for every marked triangulation t , we have

$$\mathbb{P}(t \subset \mathbb{H}_\lambda) = \left(8 + \frac{1}{h}\right)^{|\partial_{\text{in}} t| - |\partial_{\text{out}} t|} \lambda^{|t|_{\text{in}}}.$$

Proof. Uniqueness is standard. To prove existence, we construct this triangulation by peeling along the same lines as in [7]. If \mathbb{H}_λ exists, let f_0 be its triangular face that is adjacent to the root. Then f_0 has one of the three forms described by Figure 2. Moreover, we have $\mathbb{P}(\text{Case I occurs}) = \lambda \left(8 + \frac{1}{h}\right) = \frac{1}{\sqrt{1+8h}}$. By summing over all possible ways to fill the green zone, we also have

$$\mathbb{P}(\text{Case II}_i \text{ occurs}) = \mathbb{P}(\text{Case III}_i \text{ occurs}) = \left(8 + \frac{1}{h}\right)^{-i} w_\lambda(i+1)$$

for all $i \geq 0$. If we sum up these probabilities, we obtain

$$\frac{1}{\sqrt{1+8h}} + 2 \sum_{i \geq 0} \left(8 + \frac{1}{h}\right)^{-i} w_\lambda(i+1) = \frac{1}{\sqrt{1+8h}} + 2 \left(8 + \frac{1}{h}\right) W_\lambda\left(\frac{h}{1+8h}\right) = 1 \quad (10)$$




Figure 2: The three cases of peeling. In the two last cases, the parameter $i \geq 0$ corresponds to the number of edges of $\partial\mathbb{H}_\lambda$ that f_0 separates from infinity.

by (5). Since these probabilities sum up to 1, we can construct \mathbb{H}_λ by peeling with the transitions described above. Everytime case II_i or III_i occurs, we fill the green bounded region with a Boltzmann triangulation of the $(i+1)$ -gon with parameter λ . As in [7] (see also [16] in the planar case), we can check that we indeed obtain a triangulation of the halfplane, that its distribution does not depend on the choice of the peeling algorithm, and that the random triangulation we obtain has the right distribution. \square

We now state a coupling result between \mathbb{T}_λ and \mathbb{H}_λ similar to the one stated in [6] and (implicitly) [16] in the type-II case. We recall that a *peeling algorithm* is a way to assign to every triangulation with a hole an edge on its boundary (see e.g. [16, Section 1.3]). To any peeling algorithm is naturally associated a filled-in exploration of \mathbb{T}_λ . By *filled-in*, we mean that every time the face just explored separates a finite region from infinity, the interior of the finite region is entirely discovered.

Lemma 1.2. (i) For $0 < \lambda \leq \lambda_c$, the triangulation \mathbb{H}_λ is the local limit as $p \rightarrow +\infty$ of \mathbb{T}_λ^p .

(ii) For $0 < \lambda < \lambda_c$, consider a filled-in peeling algorithm \mathcal{A} with infinitely many peeling steps, and let T be the part of \mathbb{T}_λ that is discovered by \mathcal{A} . Let $(T_i)_{i \in I}$ be the infinite connected components of $\mathbb{T}_\lambda \setminus T$ and, for every $i \in I$, let e_i be an edge of ∂T that is glued to T_i , chosen in a way that only depends on T . Then conditionally on T , the maps T_i rooted at e_i are independent copies of \mathbb{H}_λ .

Proof. (i) If $p \geq 1$ and t is a marked triangulation with $|\partial_{\text{out}} t| \leq p$, let t_0 be a triangulation with a hole of perimeter p and let $t_0 + t$ be a triangulation obtained by gluing $\partial_{\text{out}} t$ to a segment of ∂t_0 . Then by the definition of \mathbb{T}_λ^p , we have

$$\begin{aligned} \mathbb{P}(t \subset \mathbb{T}_\lambda^p) &= \frac{\mathbb{P}(t_0 + t \subset \mathbb{T}_\lambda)}{\mathbb{P}(t_0 \subset \mathbb{T}_\lambda)} = \frac{c_\lambda(p + |\partial_{\text{in}} t| - |\partial_{\text{out}} t|)}{c_\lambda(p)} \lambda^{|\partial_{\text{in}} t|} \\ &\xrightarrow{p \rightarrow +\infty} \left(8 + \frac{1}{h}\right)^{|\partial_{\text{in}} t| - |\partial_{\text{out}} t|} \lambda^{|\partial_{\text{in}} t|} \\ &= \mathbb{P}(t \subset \mathbb{H}_\lambda), \end{aligned}$$

which is enough to conclude.

(ii) We first note that there are infinitely many peeling steps and all the finite holes are filled-in, so every connected component of $\mathbb{T}_\lambda \setminus T$ is halfplanar. The second point then follows from the first one since the perimeter of the region discovered after n peeling steps a.s. goes to $+\infty$ as $n \rightarrow +\infty$. See the proof of Lemma 2.16 in [6] for the same result in the type-II case. Note that for $\lambda = \lambda_c$, we have $T = \mathbb{T}_\lambda$ a.s. by Corollary 7 of [19], so the statement (ii) is irrelevant. \square

2 The skeleton decomposition of hyperbolic triangulations

The aim of this section is to prove Theorem 1. It is organized as follows. In Section 2.1, we describe the finite skeleton decomposition, which associates to every finite triangulation a finite forest, and its infinite counterpart. In Section 2.2, we compute the distribution of the skeletons of the hulls of \mathbb{T}_λ . This characterizes the skeleton of \mathbb{T}_λ entirely, but in a form that is not convenient for the proof of Theorem 1. In Section 2.3, we explain why the infinite skeleton of \mathbb{T}_λ is related to infinite leftmost geodesics in the triangulation. Section 2.4 contains a description of the strips S_λ^0 and S_λ^1 by the distribution of their hulls, without a proof of their existence. In Section 2.5, we use all that precedes to prove Theorem 1. Finally, Section 2.6 is devoted to the construction of S_λ^0 and S_λ^1 .

2.1 The skeleton decomposition of finite and infinite triangulations

The finite setting: skeleton decomposition of triangulations of the cylinder. We first recall the skeleton decomposition of triangulations introduced by Krikun [28, 27] for type-II triangulations and quadrangulations, and described in [18] for type-I triangulations (see also [8, 35]). This decomposition applies to so-called triangulations of the cylinder. Most of the presentation here is adapted from [18].

Definition 2.1. Let $r \geq 1$. A *triangulation of the cylinder of height r* is a rooted planar map in which all faces are triangles except two distinguished faces called the top and the bottom faces, such that the following properties hold. The boundaries of the top and bottom faces are simple cycles. The bottom face lies on the right of the root edge. Finally, every vertex incident to the top face is at distance r from the boundary of the bottom face, and every edge adjacent to the top face is also adjacent to a face whose third vertex is at distance $r - 1$ from the boundary of the bottom face.

If Δ is a triangulation of the cylinder of height r , we write $\partial\Delta$ and $\partial_*\Delta$ for the boundaries of the bottom and top faces. Let $p = |\partial\Delta|$ and $q = |\partial_*\Delta|$. The skeleton decomposition encodes Δ by a forest of q plane trees and a family of triangulations of polygons indexed by the vertices of this forest. For $1 \leq j \leq r - 1$, we define the ball $B_j(\Delta)$ to be the map formed by all the faces of Δ having at least one vertex at distance at most $j - 1$ from $\partial\Delta$, along with their vertices and edges. We also define the hull $B_j^\bullet(\Delta)$ as the union of $B_j(\Delta)$ and all the connected components of its complement, except the one that contains $\partial_*\Delta$. It is easy to see that $B_j^\bullet(\Delta)$ is a triangulation of the cylinder of height j . We denote by $\partial_j\Delta$ the top boundary of $B_j^\bullet(\Delta)$, with the conventions $\partial_0\Delta = \partial\Delta$ and $\partial_r\Delta = \partial_*\Delta$.

If $1 \leq j \leq r$, every edge of $\partial_j\Delta$ is incident to exactly one triangle whose third vertex belongs to $\partial_{j-1}\Delta$. Such triangles are called *downward triangles at height j* . We can define a genealogy on $\bigcup_{j=0}^r \partial_j\Delta$ by saying that $e \in \partial_j\Delta$ for $j \geq 1$ is the parent of $e' \in \partial_{j-1}\Delta$ if the downward triangle adjacent to e is the first one that one encounters when moving along $\partial_{j-1}\Delta$ in clockwise order starting from the middle of the edge e' (see Figure 3). We obtain q trees rooted on $\partial_*\Delta$. Let (t_1, \dots, t_q) be the forest obtained by listing these trees in clockwise order in such a way that the root edge of Δ lies in t_1 . This forest is called the *skeleton of Δ* and we denote it by $\text{Skel}(\Delta)$. Note that t_1 has a distinguished vertex at height r . The set of possible values of $\text{Skel}(\Delta)$ is called the set of (p, q, r) -*admissible forests* and is described by the next definition.

Definition 2.2. Let $p, q, r \geq 1$: a (p, q, r) -pre-admissible forest is a sequence $f = (t_1, \dots, t_q)$ of plane trees equipped with a distinguished vertex ρ such that:

- the maximal height of the trees t_i is r ,
- the total number of vertices at height r in the trees t_i is p ,
- ρ lies at height r .




Figure 3: A triangulation of the cylinder Δ and its skeleton $\text{Skel}(\Delta) \in \mathcal{F}'_{4,8,3}$ in purple. The cycles $\partial_j \Delta$ are in blue, and the root edge is in red. The leftmost geodesics from the vertices of $\partial_* \Delta$ to $\partial \Delta$ are in orange. The green holes must be filled by triangulations of polygons.

We write $\mathcal{F}_{p,q,r}$ for the set of (p, q, r) -pre-admissible forests. If furthermore $\rho \in t_1$, we say that f is (p, q, r) -admissible, and we write $\mathcal{F}'_{p,q,r}$ for the set of (p, q, r) -admissible forests.

Let $f = (t_1, \dots, t_p) \in \mathcal{F}_{p,q,r}$. Most of the time, we will represent f with the roots of the trees t_i on the top. Hence, if $x \in t_i$ is a vertex of f , we define the *reverse height of x in f* as r minus the height of x in t_i , and we write it $h_f^{\text{rev}}(v)$. In particular, the roots of the trees t_i have reverse height r and ρ has reverse height 0. Although quite unusual, this convention is natural because the reverse heights in $\text{Skel}(\Delta)$ match the distances to the root in the triangulation Δ . It will also be more convenient when we will deal with reverse forests with infinite height.

The forest $\text{Skel}(\Delta)$ is not enough to completely describe Δ : if we consider all the downward triangles of Δ , there is a family of holes, each of which is naturally associated to an edge of $\bigcup_{j=0}^r \partial_j \Delta$. If $e \in \partial_j \Delta$ with $1 \leq j \leq r$, the associated hole is bounded by the edges of $\partial_{j-1} \Delta$ that are children of e and by two vertical edges connecting the initial vertex of e to two vertices of $\partial_{j-1} \Delta$. This hole has perimeter $c_e + 2$, where c_e is the number of children of e , so it must be filled by a triangulation of a $(c_e + 2)$ -gon. If $c_e = 0$, it is possible that the hole of perimeter 2 is filled by the triangulation of the 2-gon consisting of a single edge, which means that the two vertical edges are simply glued together.

If f is a (p, q, r) -admissible forest, let f^* be the set of those vertices v of f such that $h_f^{\text{rev}}(v) > 0$. The decomposition we just described is a bijection between triangulations of the cylinder Δ with height r such that $\partial \Delta = p$ and $\partial_* \Delta = q$, and pairs consisting of a (p, q, r) -admissible forest f and a family $(M_v)_{v \in f^*}$ of maps such that M_v is a finite triangulation of a $(c_v + 2)$ -gon for every v .

Moreover, this decomposition encodes informations about leftmost geodesics from the vertices of $\partial_* \Delta$ to $\partial \Delta$: as can be seen on Figure 3, these geodesics (in orange on Figure 3) are the paths going between the trees of $\text{Skel}(\Delta)$. These geodesics cut Δ into q slices, each of which contains an edge of $\partial_* \Delta$. Moreover, the slice containing the i -th edge of $\partial_* \Delta$ can be completely described by the i -th tree t_i of $\text{Skel}(\Delta)$ and the maps M_v for $v \in t_i$.

The infinite setting: infinite reverse forests. If $f \in \mathcal{F}_{p,q,r}$ and $0 \leq j \leq r$, let x_1^j, \dots, x_s^j be the vertices of f lying at reverse height j in f , from left to right. For every $1 \leq i \leq s$, let t_i^j be the tree of descendants of x_i^j . Let also i_0 be the index such that the distinguished vertex




Figure 4: A forest f of height 4 and the forests $B_2(f)$ and $B'_2(f)$. The ancestors lie on the top and the distinguished vertex is in red.

belongs to $t_{i_0}^j$. We define $B_j(f)$ as the forest (t_1^j, \dots, t_s^j) , with the same distinguished vertex as f (see Figure 4), and $B'_j(f)$ as the forest $(t_{i_0}^j, t_{i_0+1}^j, \dots, t_s^j, t_1^j, \dots, t_{i_0-1}^j)$. Note that $B_j(f) \in \mathcal{F}_{p,s,j}$ and $B'_j(f) \in \mathcal{F}'_{p,s,j}$, but it is not always the case that $B_j(f) \in \mathcal{F}'_{p,s,j}$. Note also that if v is a vertex of $B_j(f)$, then $h_{B_j(f)}^{\text{rev}}(x) = h_{B'_j(f)}^{\text{rev}}(x) = h_f^{\text{rev}}(x)$.

Definition 2.3. Let $p \geq 1$. A p -pre-admissible infinite forest is a sequence $\mathbf{f} = (f_r)_{r \geq 1}$ of forests of plane trees such that

- (i) for every $r \geq 1$, there is $q \geq 1$ such that the forest f_r is (p, q, r) -pre-admissible,
- (ii) for every $s \geq r \geq 1$, we have $B_r(f_s) = f_r$.

We write $\mathcal{F}_{p,\infty,\infty}$ for the set of p -pre-admissible infinite forests. We will also write $B_r(\mathbf{f}) = f_r$ for $r \geq 1$.

Note that $\mathbf{f} \in \mathcal{F}_{p,\infty,\infty}$ can also be seen as an increasing sequence of finite graphs, and therefore as an infinite graph. We call *infinite reverse trees* the connected components of \mathbf{f} . If v is a vertex of \mathbf{f} , we have $v \in f_r$ for r large enough. Note that $h_{f_r}^{\text{rev}}(v)$ does not depend on r as long as $v \in f_r$. We call it the *reverse height of v in \mathbf{f}* and denote it by $h_{\mathbf{f}}^{\text{rev}}(v)$. We also write \mathbf{f}^* for the set of vertices of \mathbf{f} that do not lie at reverse height 0.

Definition 2.4. A p -pre-admissible forest \mathbf{f} is called p -admissible if the distinguished vertex lies in the leftmost infinite tree of \mathbf{f} , i.e. if for every $r \geq 0$, the leftmost vertex of \mathbf{f} at reverse height r is in the same infinite tree as the distinguished vertex. We write $\mathcal{F}'_{p,\infty,\infty}$ for the set of p -admissible infinite forests.

Skeleton decomposition of infinite triangulations of the p -gon. We now introduce the skeleton decomposition of infinite triangulations of the p -gon. Let T be an infinite, one-ended triangulation of the p -gon. For every $r \geq 1$, the hull $B_r^\bullet(T)$ is a triangulation of the cylinder of height r , so we can define its skeleton $f'_r = \text{Skel}(B_r^\bullet(T)) \in \mathcal{F}'_{p,q,r}$ for some q . It is also easy to see that the forests f'_r are consistent in the sense that $B'_r(f'_s) = f'_r$ for every $s \geq r \geq 1$. We claim that such a family (f'_r) always defines an infinite p -admissible forest. More precisely, if $\mathbf{f} \in \mathcal{F}'_{p,\infty,\infty}$ and $r \geq 1$, let $B'_r(\mathbf{f})$ be the reordered ball of radius r in \mathbf{f} (that is, the ball $B_r(\mathbf{f})$ in which the trees have been cyclically permuted so that the distinguished vertex lies in the first tree). Then there is a unique $\mathbf{f} \in \mathcal{F}'_{p,\infty,\infty}$ such that $B'_r(\mathbf{f}) = f'_r$ for every $r \geq 1$. We do not prove this formally, but explain how to build \mathbf{f} from (f'_r) : for any $s \geq r$, the forest $B_r(f'_s)$ is a cyclic permutation of f'_r , and this cyclic permutation does not depend on s for s large enough. Hence, we can set $f_r = B_r(f'_s)$ for s large enough and $\mathbf{f} = (f_r)_{r \geq 1}$. Therefore, there is a unique p -admissible forest, that we denote by $\text{Skel}(T)$ and call the *skeleton* of T , such that $B'_r(\text{Skel}(T)) = \text{Skel}(B_r^\bullet(T))$ for every $r \geq 1$. As in the finite case, the skeleton decomposition establishes a bijection between one-ended infinite triangulations of the p -gon and pairs consisting of a p -admissible infinite forest \mathbf{f} and a family $(M_v)_{v \in \mathbf{f}^*}$ of maps such that M_v is a finite triangulation of a $(c_v + 2)$ -gon for every v .

Remark 2.5. In order to define the skeleton decomposition, it would have been more convenient to define an infinite reverse forest \mathbf{f} as the sequence $(B'_r(\mathbf{f}))$ instead of $(B_r(\mathbf{f}))$. The reason why we chose this definition is that it will later make the decomposition of an infinite forest in infinite reverse trees much more convenient to define.

2.2 Computation of the skeleton decomposition of the hulls of \mathbb{T}_λ

We now compute the law of the skeletons of the hulls of \mathbb{T}_λ^p . The map \mathbb{T}_λ^p can be seen as a triangulation of the cylinder with infinite height. For every $r \geq 1$, the map $B_r^\bullet(\mathbb{T}_\lambda^p)$ is a triangulation of the cylinder of height r with bottom boundary length equal to p .

Lemma 2.6. Let $0 < \lambda \leq \lambda_c$. Let Δ be a triangulation of the cylinder of height r . We write p (resp. q) for the length of its bottom (resp. top) boundary. Let $f = \text{Skel}(\Delta) \in \mathcal{F}'_{p,q,r}$. Then

$$\mathbb{P}(B_r^\bullet(\mathbb{T}_\lambda^p) = \Delta) = \frac{q h_\lambda(q)}{p h_\lambda(p)} \prod_{v \in f^*} \theta_\lambda(c_v) \prod_{v \in f^*} \frac{\lambda^{|M_v|}}{w_\lambda(c_v + 2)},$$

where:

- c_v is the number of children of v in f ,
- $h \in (0, \frac{1}{4}]$ is given by (1),
- $h_\lambda(p) = \frac{1}{p} (8 + \frac{1}{h})^{-p} c_\lambda(p)$,
- θ_λ is the offspring distribution whose generating function g_λ is given by

$$g_\lambda(x) = \sum_{i \geq 0} \theta_\lambda(i) x^i = \frac{1}{x} - \frac{(1-x)(1 - \sqrt{1 - 4hx})}{2hx^2}. \quad (11)$$

Proof. The proof is exactly the same as in the case $\lambda = \lambda_c$ (Lemma 2 in [18]), up to changes of notation (the ρ of [18] corresponds to our λ and the α corresponds to $8 + \frac{1}{h}$). The same computations yield

$$\theta_\lambda(i) = \frac{1}{\sqrt{1 + 8h}} \left(\frac{h}{1 + 8h} \right)^i w_\lambda(i + 2), \quad (12)$$

and the computation of g_λ follows from (5). \square

Let $p, q, r \geq 1$ and $f \in \mathcal{F}'_{p,q,r}$. We sum the formula of Lemma 2.6 over all families $(M_v)_{v \in f^*}$ such that M_v is a triangulation of a $(c_v + 2)$ -gon for every v . By the definition of $w_\lambda(i + 2)$, we have $\sum_{n \geq 0} \# \mathcal{T}_{i+2,n} \frac{\lambda^n}{w_\lambda(i+2)} = 1$ for every $i \geq 0$, so we get

$$\mathbb{P}(\text{Skel}(B_r^\bullet(\mathbb{T}_\lambda^p)) = f) = \frac{q h_\lambda(q)}{p h_\lambda(p)} \prod_{v \in f^*} \theta_\lambda(c_v). \quad (13)$$

Note that (13) describes explicitly the distribution of $B'_r(\text{Skel}(\mathbb{T}_\lambda^p))$, so we completely know the law of $\text{Skel}(\mathbb{T}_\lambda^p)$. As we will see in Section 2.3, this is in theory enough to prove Theorem 1. However, infinite leftmost geodesics are not very tractable in this characterization. Hence, we will need to find another construction of the $\text{Skel}(\mathbb{T}_\lambda^p)$ and prove it is equivalent to (13). This will be the main goal of the rest of Section 2. Before moving on to the proof of Theorem 1, we end this subsection with a few remarks about the perimeter process of \mathbb{T}_λ .

We notice that (13) can be used to study the perimeter process of \mathbb{T}_λ in the same way as the perimeter process of \mathbb{T}_{λ_c} is studied in [28] and [35]. More precisely, by the same computation as in the proof of Lemma 3 in [18], by summing (13) over all (p, q, r) -admissible forests, we obtain

$$\mathbb{P}(|\partial B_r^\bullet(\mathbb{T}_\lambda^p)| = q) = \frac{h_\lambda(q)}{h_\lambda(p)} \mathbb{P}_q(X_\lambda(r) = p),$$

where X_λ is the Galton–Watson process with offspring distribution θ_λ . Since we know that $(|\partial B_r^\bullet(\mathbb{T}_\lambda)|)_{r \geq 0}$ is a Markov chain (by the spatial Markov property) and has the same transitions as the perimeter process of \mathbb{T}_λ^1 , we even get, for every $p, q \geq 1$ and $s \geq r \geq 0$,

$$\mathbb{P}(|\partial B_s^\bullet(\mathbb{T}_\lambda)| = q | |\partial B_r^\bullet(\mathbb{T}_\lambda)| = p) = \frac{h_\lambda(q)}{h_\lambda(p)} \mathbb{P}_q(X_\lambda(s-r) = p). \quad (14)$$

Let $m_\lambda = \sum_{i \geq 0} i \theta_\lambda(i)$ be the mean number of children. By (11), we can compute $m_\lambda = g'_\lambda(1)$ and obtain (2). In particular, we have $m_\lambda \leq 1$, with equality if and only if $\lambda = \lambda_c$. Hence, the Galton–Watson process X_λ is subcritical for $\lambda < \lambda_c$. We can therefore see the perimeter process of \mathbb{T}_λ for $\lambda < \lambda_c$ as a time-reversed subcritical branching process. Note that the perimeters and volumes of the hulls in \mathbb{T}_λ are already quite well-known. Sharp exponential growth for a fixed $\lambda < \lambda_c$ is proved in Section 2 of [16], whereas the near-critical scaling limit as $\lambda \rightarrow \lambda_c$ is studied in Section 3 of [14]. Equation (14) together with Lemma 2.8 can give explicit formulas for the generating function of the perimeters of the hull, so it should be possible to recover these results by using the same techniques as in [35], but we do not do this in this work.

2.3 Slicing the skeleton

The goal of this subsection is twofold. First, we explain the link between the skeleton decomposition and the infinite leftmost geodesics of an infinite triangulation. Second, we introduce some formalism, that will later allow us to obtain a construction of $\text{Skel}(\mathbb{T}_\lambda)$ that is more suitable for our purpose than (13).

Decomposition of a 1-admissible forest in reverse trees. An infinite 1-admissible forest \mathbf{f} may contain several infinite reversed trees, and we will need to study the way these trees are placed with respect to each other. This can be encoded by a genealogical structure, which is described by the red tree on Figure 5. If \mathbf{t} is one of the infinite trees of \mathbf{f} , let $h_{\min}^{\text{rev}}(\mathbf{t})$ be the reverse height of the lowest vertex of \mathbf{t} . We consider the set of pairs (\mathbf{t}, i) where \mathbf{t} is an infinite tree of \mathbf{f} and $i \geq h_{\min}^{\text{rev}}(\mathbf{t})$. If $i > h_{\min}^{\text{rev}}(\mathbf{t})$, the parent of (\mathbf{t}, i) is $(\mathbf{t}, i-1)$. If $i = h_{\min}^{\text{rev}}(\mathbf{t}) > 0$, let \mathbf{t}' be the first infinite tree on the left of \mathbf{t} such that $h_{\min}^{\text{rev}}(\mathbf{t}') \leq i-1$ (note that \mathbf{t}' always exists because \mathbf{f} is admissible). Then the parent of (\mathbf{t}, i) is $(\mathbf{t}', i-1)$. Finally, if $i = 0$, then (\mathbf{t}, i) has no parent. This genealogy is encoded in an infinite plane tree with no leaf that we denote by $\mathbf{U}(\mathbf{f})$ (see Figure 5). More intuitively, the tree $\mathbf{U}(\mathbf{f})$ is the tree whose branches pass between the infinite trees of \mathbf{f} . Note that the genealogy in $\mathbf{U}(\mathbf{f})$ is "reversed" compared to the genealogy in \mathbf{f} : the parent of a vertex x of $\mathbf{U}(\mathbf{f})$ lies below x , whereas in the forest \mathbf{f} , the parent of a vertex lies above it. Therefore, the heights in $\mathbf{U}(\mathbf{f})$ match the reverse heights in \mathbf{f} . We also write $B_r(\mathbf{U}(\mathbf{f}))$ for the subtree of $\mathbf{U}(\mathbf{f})$ whose vertices are the (\mathbf{t}, i) with $i \leq r$. Note that the tree $B_r(\mathbf{U}(\mathbf{f}))$ is not a function of $B_r(\mathbf{f})$ (it is impossible by looking at $B_r(\mathbf{f})$ to know if two vertices belong to the same infinite tree). Finally, it is easy to see that a 1-admissible forest is completely described by the tree $\mathbf{U}(\mathbf{f})$ and the infinite trees it contains.

Leftmost infinite geodesics and decomposition of the skeleton in reverse trees. Let T be an infinite triangulation of a 1-gon, let ρ be its root vertex (i.e. the unique point on its boundary), and let $\mathbf{f} = \text{Skel}(T)$. We have seen that for every $r \geq 0$, the paths going between the trees of $B_r(\mathbf{f})$ correspond to the leftmost geodesics from ρ to the vertices of $\partial B_r^\bullet(T)$ (cf. Figure 3). Therefore, infinite paths started from ρ in $\mathbf{U}(\mathbf{f})$ correspond to leftmost geodesic rays in T , so the tree of leftmost infinite geodesics in T is isomorphic to $\mathbf{U}(\mathbf{f})$.

The skeleton decomposition of infinite strips. We will also need to describe the skeleton decomposition of *strips*, which are infinite triangulations with two infinite geodesic boundaries. They correspond to the S^0 and S^1 appearing in Theorem 1.




Figure 5: An infinite 1-admissible forest \mathbf{f} in which 4 trees t_1, t_2, t_3 and t_4 reach reverse height 4. In red, the tree $\mathbf{U}(\mathbf{f})$ and the names of some of its vertices. In the proof of Proposition 2.13 for $r = 4$, the p_j^r are equal to 2, 1, 2 and 1 and the $h(f_j)$ are equal to 4, 0, 3 and 2.

Definition 2.7. An *infinite strip* is a one-ended planar triangulation bounded by two infinite geodesics γ_ℓ (on its left) and γ_r (on its right), and equipped with a root vertex ρ , such that:

- (i) ρ is the only common point of γ_ℓ and γ_r ,
- (ii) for every $i, j \geq 0$, the path γ_r is the only geodesic from $\gamma_r(i)$ to $\gamma_r(j)$,
- (iii) γ_r and γ_ℓ are the only leftmost geodesic rays in S .

Exactly as for infinite triangulations of the p -gon, if S is an infinite strip, we define its ball $B_r(S)$ of radius r as the union of all its faces containing a vertex at distance at most $r - 1$ from ρ . We also define its hull $B_r^\bullet(S)$ of radius r as the union of $B_r(S)$ and all the finite connected components of its complement. To define the skeleton of an infinite strip S , we note that there is a simple transformation that associates to S an infinite triangulation of the p -gon, where p is the number of edges on $\partial B_1^\bullet(S)$. More precisely, we write \tilde{S} for the map obtained by rooting $S \setminus B_1^\bullet(S)$ at the leftmost edge of $\partial B_1^\bullet(S)$ and gluing γ_ℓ and γ_r together. We define the skeleton of S as the skeleton of \tilde{S} (see Figure 6). Like for triangulations of the plane, the skeleton decomposition is a bijection between infinite strips and pairs consisting:

- on the one hand of an infinite reverse tree \mathbf{t} that is rooted at its leftmost vertex of reverse height 0,
- on the other hand of a family of maps $(M_v)_{v \in \mathbf{t}}$ such that M_v is a triangulation of a $(c_v + 2)$ -gon for every v .

The fact that \mathbf{t} must be connected follows from the uniqueness of the leftmost geodesic rays γ_ℓ and γ_r in a strip (recall that the trees of the skeleton are separated by infinite leftmost geodesics). Note that this time, the triangulations filling the holes are indexed by \mathbf{t} and not \mathbf{t}^* since the M_v for v at reverse height 0 are used to encode $B_1^\bullet(S)$ (see Figure 6).

Finally, let T be an infinite triangulation of a 1-gon, and let \mathbf{f} be its skeleton. As we have seen above, the tree $\mathbf{U}(\mathbf{f})$ can be seen as the tree of leftmost infinite geodesics of T . Moreover, $\mathbf{U}(\mathbf{f})$ cuts T into strips, whose skeletons are the infinite reverse trees of \mathbf{f} . This reduces the proof of Theorem 1 to the study of $\text{Skel}(\mathbb{T}_\lambda^1)$.




Figure 6: An infinite strip S and its skeleton.

2.4 The distribution of S_λ^0 and S_λ^1

The goal of this subsection is to describe (without to build them explicitly) the strips S_λ^0 and S_λ^1 in a way that will allow us to prove Theorem 1. This description will involve the quasi-stationary distribution of the branching process X_λ . Hence, we start with two explicit computations about X_λ . First, we give an explicit formula for the iterates of the generating function g_λ . We define $g_\lambda^{\circ r}$ by $g_\lambda^{\circ 0} = \text{Id}$ and $g_\lambda^{\circ(r+1)} = g_\lambda \circ g_\lambda^{\circ r}$ for every $r \geq 0$. Note that $g_{\lambda_c}^{\circ r}$ is explicitly computed in [18]: we have

$$g_{\lambda_c}^{\circ r}(x) = 1 - \left(r + \frac{1}{\sqrt{1-x}} \right)^{-2} \quad (15)$$

for every $x \in [0, 1]$. The iterates of g_λ in the subcritical case are also computed in [35], with different notations. Note that when $\lambda \rightarrow \lambda_c$ in the formula below, we recover (15).

Lemma 2.8. Let $0 < \lambda < \lambda_c$. For every $r \geq 0$, we have

$$g_\lambda^{\circ r}(x) = 1 - \frac{1 - 4h}{4h \sinh^2 \left(\text{argsh} \sqrt{\frac{1-4h}{4h(1-x)}} + rb_\lambda \right)},$$

where $b_\lambda = \text{argch} \frac{1}{\sqrt{4h}} = -\frac{1}{2} \ln m_\lambda$, and h is as in (1).

Proof. This computation is already done in Section 3.1 of [35], with different notation. In [35], the function $\varphi_t(u)$ for $0 \leq t \leq 1$ is defined by

$$\varphi_t(u) = \frac{\rho^s}{\alpha t} \sum_{i \geq 0} (\alpha t)^i w_{\rho s}(i+2) u^i,$$

where $\alpha = \frac{1}{12}$, $\rho = \frac{1}{12\sqrt{3}} = \lambda_c$ and $s = t\sqrt{3-2t}$. For $t = \frac{12h}{1+8h}$, we have $s = \frac{\lambda}{\lambda_c}$, so by (12), we have $\frac{\rho^s}{\alpha t} (\alpha t)^i w_{\rho s}(i+2) = \theta_\lambda(i)$ for every $i \geq 0$, and our function g_λ corresponds to the φ_t of [35] for $t = \frac{12h}{1+8h}$. Our formula follows then immediately from Lemma 3 of [35]. Finally, using (2), it is easy to check that $e^{-2b_\lambda} = m_\lambda$. \square

Remark 2.9. It is also possible to prove the last lemma directly by using (14). Indeed, the probability must add up to 1 when we sum them over q , which shows that $(h_\lambda(p))$ is an invariant measure for X_λ . If we write $H_\lambda(x) = \sum_{p \geq 1} h_\lambda(p) x^p$, this easily gives

$$H_\lambda(g_\lambda(x)) = H_\lambda(x) + H_\lambda(g_\lambda(0)),$$

so $g_\lambda^{\circ r}(x) = H_\lambda^{-1}(H_\lambda(x) + rH_\lambda(g_\lambda(0)))$. Since H_λ can be explicitly computed, this also gives the result.

Our second computation deals with the quasi-stationary measure of X_λ . We first recall a few facts about quasi-stationary measures of Galton–Watson processes (see for example Chapter 7 of [10]). For every $p \geq 1$, the ratio $\frac{\mathbb{P}_1(X_\lambda(n)=p)}{\mathbb{P}_1(X_\lambda(n)=1)}$, where \mathbb{P}_1 is the distribution of X_λ started from 1, is nondecreasing in n and converges to $\pi_\lambda(p) < +\infty$. Moreover, let Π_λ be the generating function of $(\pi_\lambda(p))_{p \geq 1}$. In our case, it is explicit:

Lemma 2.10. For $0 < \lambda < \lambda_c$, we have

$$\Pi_\lambda(x) = \frac{1}{\sqrt{1-4h}} \left(1 - (1-x) \left(\frac{\sqrt{1-4h} + 1}{\sqrt{1-4h} + \sqrt{1-4hx}} \right)^2 \right).$$

We also have

$$\Pi_{\lambda_c}(x) = 2 \left(\frac{1}{\sqrt{1-x}} - 1 \right).$$

In particular, we have $\Pi_\lambda(\theta_\lambda(0)) = \Pi_\lambda(1-h) = \frac{1-\sqrt{1-4h}}{2h}$.

Proof. By the definition of $\pi_\lambda(p)$, we have

$$\Pi_\lambda(x) = \lim_{n \rightarrow +\infty} \frac{g_\lambda^{on}(x) - g_\lambda^{on}(0)}{(g_\lambda^{on})'(0)}.$$

Hence, by using Lemma 2.8, the computation of $\Pi_\lambda(x)$ is straightforward. Note that the case $\lambda = \lambda_c$ already appears in [18] (in the proof of Lemma 3). \square

We can now describe the distribution of S_λ^0 and S_λ^1 . We will first admit the existence of reverse trees with a certain distribution (described by the next lemma), and later (Section 2.6) build these trees by a spine decomposition approach.

Lemma 2.11. There are two infinite reverse trees τ_λ^0 and τ_λ^1 whose distributions are characterized as follows. For every $r \geq 0$ and every forest (t_1, \dots, t_p) of height exactly r , we have

$$\mathbb{P}(B_r(\tau_\lambda^0) = (t_1, \dots, t_p)) = \frac{\pi_\lambda(p) m_\lambda^{-r}}{\Pi_\lambda(\theta_\lambda(0))} \prod_{i=1}^p \prod_{v \in t_i} \theta_\lambda(c_v) \quad (16)$$

and, if f has only one vertex at height r ,

$$\mathbb{P}(B_r(\tau_\lambda^1) = (t_1, \dots, t_p)) = \frac{\pi_\lambda(p) m_\lambda^{-r}}{\theta_\lambda(0)} \prod_{i=1}^p \prod_{v \in t_i} \theta_\lambda(c_v), \quad (17)$$

where c_v is the number of children of v for every vertex v .

Moreover, let $\tau_\lambda^{1,*}$ be the tree τ_λ^1 in which we have cut the only vertex at reverse height 0. The reverse heights in $\tau_\lambda^{1,*}$ are shifted by 1, so that the minimal reverse height in $\tau_\lambda^{1,*}$ is 0. This allows us to define the two strips that appear in Theorem 1.

Definition 2.12. We denote by S_λ^0 (resp. S_λ^1) the random infinite strip whose skeleton is τ_λ^0 (resp. $\tau_\lambda^{1,*}$) and where conditionally on the skeleton, all the holes are filled with independent Boltzmann triangulations with parameter λ .

The reason why we need to replace τ_λ^1 by $\tau_\lambda^{1,*}$ in this last definition is linked to the root transformation between \mathbb{T}_λ and \mathbb{T}_λ^1 , and will be explained in details in the end of Subsection 2.5 (see Figure 7).

2.5 Proof of Theorem 1

The goal of this subsection is to prove Theorem 1. For this, we build a random infinite 1-admissible forest \mathbf{F}_λ directly in terms of its decomposition in reverse infinite trees, and we show that it has the same distribution as $\text{Skel}(\mathbb{T}_\lambda^1)$.

We recall that $\mu_\lambda(0) = 0$ and $\mu_\lambda(k) = m_\lambda(1 - m_\lambda)^{k-1}$ for $k \geq 1$. We have seen that an infinite 1-admissible forest \mathbf{f} is completely described by the tree $\mathbf{U}(\mathbf{f})$ and the infinite trees that \mathbf{f} contains. Therefore, there is a unique (in distribution) random 1-admissible forest \mathbf{F}_λ such that $\mathbf{U}(\mathbf{F}_\lambda)$ is a Galton–Watson tree with offspring distribution μ_λ and, conditionally on $\mathbf{U}(\mathbf{F}_\lambda)$:

- (i) the trees of \mathbf{F}_λ are independent,
- (ii) the unique tree that reaches reverse height 0 has the same distribution as τ_λ^1 ,
- (iii) all the other trees have the same distribution as τ_λ^0 described above.

A more rigorous (but heavier) way to define \mathbf{F} would be to build explicitly $B_r(\mathbf{F})$ by concatenating independent forests of the form $B_j(\tau^0)$ and $B_j(\tau^1)$.

In order to prove that \mathbf{F}_λ has indeed the same distribution as $\text{Skel}(\mathbb{T}_\lambda^1)$, we need to introduce one last notation. Let \mathbf{f} be a 1-admissible forest, and let $r \geq 1$. Let ℓ be the number of infinite reverse trees of \mathbf{f} that intersect $B_r(\mathbf{f})$, and let $\mathbf{t}_1^r, \dots, \mathbf{t}_\ell^r$ be these trees, from left to right. For every $1 \leq j \leq \ell$, we denote by $p_j^r(\mathbf{f})$ the number of vertices of \mathbf{t}_j^r whose reverse height in \mathbf{f} is exactly r (see Figure 5 for an example). Note that for each $1 \leq j \leq \ell$, the trees of $B_r(\mathbf{f})$ that belong to \mathbf{t}_j^r are $p_j^r(\mathbf{f})$ consecutive trees of $B_r(\mathbf{f})$. Therefore, if we already know $B_r(\mathbf{f})$, knowing the values $p_j^r(\mathbf{f})$ is equivalent to knowing which of the trees of $B_r(\mathbf{f})$ lie in the same infinite reverse tree of \mathbf{f} . We write $\tilde{B}_r(\mathbf{f})$ for the pair $(B_r(\mathbf{f}), (p_1^r(\mathbf{f}), \dots, p_\ell^r(\mathbf{f})))$. The reason why we introduce this object is that its distribution for $\mathbf{f} = \mathbf{F}_\lambda$ will be easier to compute.

The key result is the next proposition. As explained in Section 2.1, a 1-admissible forest \mathbf{f} is characterized by its reordered balls $B'_r(\mathbf{f})$, so the distribution of $\text{Skel}(\mathbb{T}_\lambda^1)$ is characterized by the distribution of its reordered balls. Therefore, the next proposition will imply that \mathbf{F}_λ and $\text{Skel}(\mathbb{T}_\lambda^1)$ have the same distribution, which implies Theorem 1.

Proposition 2.13. For every $r \geq 0$, the forests $B'_r(\mathbf{F}_\lambda)$ and $\text{Skel}(B_r^\bullet(\mathbb{T}_\lambda^1))$ have the same distribution.

Proof. In all this proof, we will fix $0 < \lambda \leq \lambda_c$ and omit the parameter λ in the notation. The idea of the proof is the following: we will first compute the law of $\tilde{B}_r(\mathbf{F})$, then that of $B_r(\mathbf{F})$ and finally that of $B'_r(\mathbf{F})$.

First, we know that $B_r(\mathbf{U}(\mathbf{F}))$ is a tree with height r in which all the leaves lie at height r . Moreover, if t is such a tree, we have

$$\mathbb{P}(B_r(\mathbf{U}(\mathbf{F})) = t) = \prod_{v \in t_{<r}} m(1 - m)^{c_v - 1} = m^{|t_{<r}|} (1 - m)^{|t| - 1 - |t_{<r}|}, \quad (18)$$

where $t_{<r}$ is the set of vertices of t at height strictly less than r , and c_v is the number of children of a vertex v .

Now let $(f, (p_1, \dots, p_\ell))$ be a possible value of $\tilde{B}_r(\mathbf{F})$, i.e. $f = (t_1, \dots, t_p)$ is a forest of height r , the positive integers p_1, \dots, p_ℓ satisfy $p_1 + \dots + p_\ell = p$, and f has a unique vertex at reverse height 0 which lies in one of the trees t_1, \dots, t_{p_1} . For every $1 \leq j \leq \ell$, we write $f_j = (t_{p_1 + \dots + p_{j-1} + 1}, \dots, t_{p_1 + \dots + p_j})$ and $h(f_j)$ for the height of the forest f_j . Each of these forests corresponds to one of the infinite trees that reach reverse height r .

We now check that the tree $B_r(\mathbf{U}(\mathbf{F}))$ is a measurable function of $\tilde{B}_r(\mathbf{F})$ (although not of $B_r(\mathbf{F})$). The reader may find helpful to look at Figure 5 while reading what follows. We consider the tree u whose vertices are the pairs (j, i) with $1 \leq j \leq \ell$ and $r - h(f_j) \leq i \leq r$, and in which:

- (i) the pair $(1, 0)$ is the root vertex,
- (ii) if $i > r - h(f_j)$, then the parent of (j, i) is $(j, i - 1)$,
- (iii) if $i = r - h(f_j)$, then the parent of (j, i) is $(k, i - 1)$, where k is the greatest integer such that $k < j$ and $h(f_k) \geq r - i + 1$.

This is the natural analog of $\mathbf{U}(\mathbf{f})$ for the finite forest f (cf. Figure 5). Note that for every $1 \leq j \leq \ell$, there are exactly $h(f_j) + 1$ vertices of u of the form (j, i) , exactly one of which lies at height r . Hence, we have

$$\sum_{j=1}^{\ell} h(f_j) = |u_{<r}|. \quad (19)$$

It is easy to see that if $\tilde{B}_r(\mathbf{F}) = (f, (p_1, \dots, p_\ell))$, then $B_r(\mathbf{U}(\mathbf{F})) = u$. It is also possible to read the heights $h(f_j)$ on the tree u : first, we have $h(f_1) = r$. Moreover, if $2 \leq j \leq \ell$, let x_j be the j -th leaf in u starting from the left. Then $h(f_j)$ is the smallest h such that the ancestor of x_j at height $r - h$ is not the leftmost child of its parent (cf. Figure 5). Hence, conditionally on $B_r(\mathbf{U}(\mathbf{F})) = u$, the forest $\tilde{B}_r(\mathbf{F})$ is a concatenation of ℓ independent forests of heights $h(f_1), \dots, h(f_\ell)$. The first forest has the same distribution as $B_r(\tau^1)$ and, for $j \geq 2$, the j -th forest has the same distribution as $B_{h(f_j)}(\tau^0)$. By combining this observation with (18) and Lemma 2.11, we obtain

$$\begin{aligned} \mathbb{P}\left(\tilde{B}_r(\mathbf{F}) = (f, (p_i)_{1 \leq i \leq \ell})\right) &= m^{|u_{<r}|} (1 - m)^{|u| - |u_{<r}| - 1} \times \frac{\pi(p_1) m^{-r}}{\theta(0)} \prod_{v \in f_1} \theta(c_v) \\ &\quad \times \prod_{j=2}^{\ell} \frac{\pi(p_j) m^{-h(f_j)}}{\Pi(\theta(0))} \prod_{v \in f_j} \theta(c_v) \\ &= \frac{m^{|u_{<r}| - \sum_{j=1}^{\ell} h(f_j)} (1 - m)^{|u| - |u_{<r}| - 1}}{\theta(0) \Pi(\theta(0))^{\ell-1}} \prod_{j=1}^{\ell} \pi(p_j) \times \prod_{v \in f} \theta(c_v) \\ &= \frac{1}{\theta(0)} \left(\frac{1 - m}{\Pi(\theta(0))} \right)^{\ell-1} \prod_{j=1}^{\ell} \pi(p_j) \times \prod_{v \in f} \theta(c_v), \end{aligned} \quad (20)$$

where in the end, we used (19). Moreover, by Lemma 2.10, we can compute $\frac{1-m}{\Pi(g(0))} = \sqrt{1 - 4h}$.

We now compute the distribution of $B_r(\mathbf{F})$: let $f = (t_1, \dots, t_p)$ be a possible value of $B_r(\mathbf{F})$, and let i_0 be the index such that the only vertex of reverse height 0 belongs to t_{i_0} . We need to sum (20) over all the possible values of $\ell \geq 1$ and $p_1, \dots, p_\ell \geq 1$ such that $\sum_{j=1}^{\ell} p_j = p$ and $p_1 \geq i_0$ (by construction of \mathbf{F} , the lowest vertex always belongs to the leftmost infinite tree). We obtain

$$\mathbb{P}(B_r(\mathbf{F}) = f) = \frac{1}{\theta(0)} \left(\sum_{\ell \geq 1} (1 - 4h)^{\frac{\ell-1}{2}} \sum_{\substack{p_1 + \dots + p_\ell = p \\ p_1 \geq i_0}} \prod_{j=1}^{\ell} \pi(p_j) \right) \prod_{v \in f} \theta(c_v). \quad (21)$$

Now let $f' = (t_1, \dots, t_p)$ be a possible value of $B'_r(\mathbf{F})$, i.e. a forest of height r in which the only vertex of reverse height 0 lies in t_1 . To obtain $\mathbb{P}(B'_r(\mathbf{F}) = f')$, we need to sum Equation (21) over all the forests one may get by applying a cyclic permutation to the trees of f' . The values of p and $\prod_{v \in f} \theta(c_v)$ are the same for all these forests, but the value of i_0 ranges from 1

to p , so we have

$$\begin{aligned} \mathbb{P}(B'_r(\mathbf{F}) = f') &= \frac{1}{\theta(0)} \left(\sum_{i_0=1}^p \sum_{\ell \geq 1} (1-4h)^{\frac{\ell-1}{2}} \sum_{\substack{p_1+\dots+p_\ell=p \\ p_1 \geq i_0}} \prod_{j=1}^{\ell} \pi(p_j) \right) \prod_{v \in f} \theta(c_v) \\ &= \frac{1}{\theta(0)} \left(\sum_{\ell \geq 1} (1-4h)^{\frac{\ell-1}{2}} \sum_{p_1+\dots+p_\ell=p} p_1 \prod_{j=1}^{\ell} \pi(p_j) \right) \prod_{v \in f} \theta(c_v). \end{aligned}$$

By comparing this to (13), we only need to prove that for every $p \geq 1$,

$$\sum_{\ell \geq 1} (1-4h)^{\frac{\ell-1}{2}} \sum_{p_1+\dots+p_\ell=p} p_1 \prod_{j=1}^{\ell} \pi(p_j) = \frac{ph_\lambda(p)}{h_\lambda(1)}. \quad (22)$$

It is enough to show that the generating functions of both sides coincide. But the generating function of the left-hand side is

$$\begin{aligned} \sum_{p \geq 1} \left(\sum_{\ell \geq 1} (1-4h)^{\frac{\ell-1}{2}} \sum_{p_1+\dots+p_\ell=p} p_1 \prod_{j=1}^{\ell} \pi(p_j) \right) y^p &= \sum_{\ell \geq 1} (1-4h)^{\frac{\ell-1}{2}} y \Pi'(y) \Pi(y)^{\ell-1} \\ &= \frac{y \Pi'(y)}{1 - \sqrt{1-4h} \Pi(y)}, \end{aligned}$$

which is explicitly known by Lemma 2.10. On the other hand, we recall from Lemma 2.6 that $h_\lambda(p) = \frac{1}{p} \left(8 + \frac{1}{h}\right)^{-p} c_\lambda(p)$. Hence, the generating function of the right-hand side of (22) is

$$\frac{1}{c_\lambda(1)} \left(8 + \frac{1}{h}\right) \sum_{p \geq 1} \left(8 + \frac{1}{h}\right)^{-p} c_\lambda(p) y^p = \frac{1}{y c_\lambda(1)} \left(8 + \frac{1}{h}\right) C_\lambda \left(\frac{h}{1+8h} y\right),$$

where C_λ is given by (9). Hence, it is easy to check that the generating functions of both sides of (22) coincide, which concludes the proof. \square

End of the proof of Theorem 1. By Proposition 2.13, the skeleton of \mathbb{T}_λ^1 has the same distribution as \mathbf{F}_λ . In particular, the infinite leftmost geodesics are the paths separating infinite trees in \mathbf{F}_λ , so their union is isomorphic to $\mathbf{U}(\mathbf{F}_\lambda)$. Moreover, consider the strips delimited by $\mathbf{U}(\mathbf{F}_\lambda)$ in \mathbb{T}_λ^1 . Let e_1 be the root edge of \mathbb{T}_λ^1 (i.e. the loop on its boundary). The skeletons of the strips that are not adjacent to e_1 are independent copies of τ^0 and all the holes are filled with independent Boltzmann triangulations with parameter λ , so these strips are independent copies of S^0 .

The case of the strip containing e_1 needs to be handled more carefully because of the presence of the bottom boundary (actually, this is not exactly a strip in the sense of Definition 2.7). More precisely, let $S^{1,+}$ be the random strip obtained from τ^1 as on the left part of Figure 7, where all the green holes are filled with independent Boltzmann triangulations with parameter λ . Then the strip of \mathbb{T}_λ^1 adjacent to e_1 has skeleton τ^1 and its holes are filled with Boltzmann triangulations, but it has a bottom boundary of length one, so it has the same distribution as $S^{1,+}$.

Finally, we recall that \mathbb{T}_λ is the image of \mathbb{T}_λ^1 by the following root transformation. Let f_1 be the face of \mathbb{T}_λ^1 adjacent to its boundary loop e_1 . We obtain \mathbb{T}_λ by contracting e_1 and gluing together the two other sides of f_1 . This transformation does not affect the tree of infinite leftmost geodesics and the strips that are not adjacent to e_1 . Its effect on the strip $S^{1,+}$ adjacent to e_1 is described on Figure 7, and the strip we obtain is then S^1 . Note also that the root edge on the right of Figure 7 may be on the boundary of the strip, but in this case it is necessarily on the left boundary, so S^1 is always the strip containing the face on the *right* of the root edge, as in the statement of the Theorem. This proves Theorem 1. \square



Figure 7: On the left, the strip $S^{1,+}$ and its skeleton τ^1 . On the right, the strip S^1 and its skeleton $\tau^{1,*}$. The root transformation contracting the face f_1 (whose sides are in red on the left) sends the first strip to the second. The skeletons are in violet.

2.6 Construction of reverse Galton–Watson trees and infinite strips

Our goal in this subsection is to construct the reverse subcritical trees τ^0 and τ^1 and to prove Lemma 2.11 and a few useful estimates. We will also show that τ^0 is the local limit as $n \rightarrow +\infty$ of a subcritical Galton–Watson tree conditioned to die at generation n , seen from its last generation. This is different from the more usual Galton–Watson tree conditioned to survive [26], where the tree is seen from its root. Moreover, τ^1 is just τ^0 , conditioned to have only one vertex at reverse height 0. Our trees will be built by a spine decomposition approach, which will be useful to obtain geometric estimates on these trees in Section 4. Since our proofs hold in a more general setting, we present the results of this subsection for any critical or subcritical Galton–Watson process, and denote by g the generating function of the offspring distribution. The trees τ_λ^0 and τ_λ^1 that we need are obtained by taking $g = g_\lambda$.

We start with a vertical half-line which is infinite on the top side, that we call the *spine*. The root of the tree will be the lowest point on the spine. For every $r \geq 1$, we sample a random pair (L_r, R_r) of integers such that all these pairs are independent and, for every $i, j \geq 0$, we have

$$\mathbb{P}(L_r = i, R_r = j) = \frac{g^{or}(0) - g^{\circ(r-1)}(0)}{g^{\circ(r+1)}(0) - g^{\circ r}(0)} \theta(i + j + 1) g^{\circ(r-1)}(0)^i g^{\circ r}(0)^j. \quad (23)$$

To check that (23) defines a probability measure, we just need to sum the right-hand side over pairs (i, j) with a fixed value of $k = i + j$ and then sum over k . Similarly, by a straightforward computation, we obtain

$$\mathbb{E}[L_r + R_r] = \frac{g^{or}(0)g'(g^{or}(0)) - g^{\circ(r-1)}(0)g'(g^{\circ(r-1)}(0))}{g^{\circ(r+1)}(0) - g^{\circ r}(0)} - 1. \quad (24)$$

We call s_r the vertex at height r in the spine. For all $r \geq 2$, let $t_1^r, \dots, t_{L_r}^r$ be L_r independent Galton–Watson trees with offspring distribution θ conditioned on having height at most $r - 2$. We graft L_r edges to the left of the vertex s_r , and the trees $t_1^r, \dots, t_{L_r}^r$ to the other ends of these edges. Similarly, for all $r \geq 1$, let $u_1^r, \dots, u_{R_r}^r$ be R_r independent Galton–Watson trees with offspring distribution θ conditioned on having height at most $r - 1$. We graft R_r edges to the right of s_r and the trees $u_1^r, \dots, u_{R_r}^r$ to the other ends of these edges (see Figure 8).

We denote by τ^0 the infinite tree we obtain, and we define a genealogy on it: for every $r \geq 1$, the children of s_r are s_{r-1} and the roots of the trees t_i^r and u_j^r . Inside the trees t_i^r and u_j^r , the genealogy is the usual one in a Galton–Watson tree. We fix the reverse height of the root at 0 and declare that the parent of a vertex of reverse height r has reverse height $r + 1$ (it corresponds to the height of the vertices on Figure 8). By the conditioning we have chosen for the trees t_i^r and u_j^r , every vertex of the tree has a nonnegative reverse height, and the root is the leftmost vertex with reverse height 0.

Lemma 2.14. For every $r \geq 0$, the tree τ^0 has a finite number of vertices at height r .




Figure 8: Construction of the tree τ^0 . Every vertex lies below its parent.

Proof. We fix $r \geq 0$ and take $r' > r$, and we define $A_r^{r'}$ as the event

{one of the trees grafted at $s_{r'}$ reaches reverse height r in τ^0 }.

We want to show that $\mathbb{P}(A_r^{r'})$ decreases fast enough in r' . Consider one of the trees $t_i^{r'}$ grafted on the left of the spine at $s_{r'}$. This is a Galton–Watson tree conditioned to have height at most $r' - 2$, so the probability that $t_i^{r'}$ has height at least $r' - r - 1$ is $\frac{g^{\circ(r'-1)}(0) - g^{\circ(r'-r-1)}(0)}{g^{\circ(r'-1)}(0)}$. The denominator is larger than $\frac{1}{2}$ for r' large enough. Hence, the probability that one of the $L_{r'}$ trees grafted on the left of $s_{r'}$ reaches reverse height r is at most $2 \left(g^{\circ(r'-1)}(0) - g^{\circ(r'-r-1)}(0) \right) \mathbb{E}[L_{r'}]$. Similarly, the probability of the analog event on the right is at most $2 \left(g^{\circ r'}(0) - g^{\circ(r'-r-1)}(0) \right) \mathbb{E}[R_{r'}]$. Therefore, for r' large enough, we have

$$\mathbb{P}\left(A_r^{r'}\right) \leq 2 \left(g^{\circ r'}(0) - g^{\circ(r'-r-1)}(0) \right) \mathbb{E}[L_{r'} + R_{r'}].$$

Moreover, the right-hand side of (24) can be rewritten $\frac{\tilde{g}(x_r) - \tilde{g}(x_{r-1})}{x_{r+1} - x_r}$, where we write $x_r = g^{\circ r}(0)$ and $\tilde{g}(x) = xg'(x) - g(x)$ for $x \in [0, 1]$. We want to prove $\sum_{r' > r} \mathbb{P}\left(A_r^{r'}\right) < +\infty$, that is

$$2 \sum_{r' > r} \frac{x_{r'} - x_{r'-r-1}}{x_{r'+1} - x_{r'}} (\tilde{g}(x_{r'}) - \tilde{g}(x_{r'-1})) < +\infty. \quad (25)$$

Since g is 1-Lipschitz, the sequence $(x_{r'+1} - x_{r'})_{r' \geq 0}$ is nonincreasing, so we have

$$\frac{x_{r'} - x_{r'-r-1}}{x_{r'+1} - x_{r'}} \leq (r+1) \frac{x_{r'-r} - x_{r'-r-1}}{x_{r'+1} - x_{r'}} = \frac{(r+1)(x_{r'-r} - x_{r'-r-1})}{g^{\circ(r+1)}(x_{r'-r}) - g^{\circ(r+1)}(x_{r'-r-1})} \rightarrow \frac{r+1}{(g^{\circ(r+1)})'(1)}$$

as $r' \rightarrow +\infty$. In particular, in (25), the first factor is bounded while the second is nonnegative and telescopic, so the sum converges. Therefore, a.s. $A_r^{r'}$ does not occur for r' large enough, which proves the lemma. \square

For $r \geq 0$, let $d_r(\tau^0)$ be the tree of descendants of the r -th vertex s_r of the spine.

Lemma 2.15. The tree $d_r(\tau^0)$ has the same distribution as a Galton–Watson tree with offspring distribution θ , conditioned to have height exactly r .

Proof. Let t be a finite plane tree of height r . Let $s_0(t)$ be its leftmost vertex of reverse height 0 and let $s_r(t)$ be its root. Let $(s_0(t), s_1(t), \dots, s_r(t))$ be the unique geodesic path from $s_0(t)$ to

$s_r(t)$. For every $1 \leq i \leq r$, let $\ell_i(t)$ (resp. $r_i(t)$) be the number of children of $s_i(t)$ on the left (resp. on the right) of $s_{i-1}(t)$, and let $(v_j^i(t))_{1 \leq j \leq \ell_i(t)}$ (resp. $(w_j^i(t))_{1 \leq j \leq r_i(t)}$) be those children. Finally, for every i and j , let $t_j^i(t)$ (resp. $u_j^i(t)$) be the tree of descendants of $v_j^i(t)$ (resp. $w_j^i(t)$). Then we have

$$\mathbb{P}(d_r(\boldsymbol{\tau}^0) = t) = \prod_{i=1}^r \mathbb{P}(L_i = \ell_i(t), R_i = r_i(t)) \times \prod_{i=1}^r \left(\prod_{j=1}^{\ell_i(t)} \mathbb{P}(t_j^i = t_j^i(t)) \times \prod_{j=1}^{r_i(t)} \mathbb{P}(u_j^i = u_j^i(t)) \right). \quad (26)$$

Moreover, by definition of t_j^i and u_j^i , we have

$$\mathbb{P}(t_j^i = t_j^i(t)) = \frac{1}{g^{\circ(r-1)}(0)} \prod_{v \in t_j^i(t)} \theta(c_v) \quad \text{and} \quad \mathbb{P}(u_j^i = u_j^i(t)) = \frac{1}{g^{\circ r}(0)} \prod_{v \in u_j^i(t)} \theta(c_v).$$

By combining this with (23) and (26), we obtain

$$\mathbb{P}(d_r(\boldsymbol{\tau}^0) = t) = \frac{1}{g^{\circ(r+1)}(0) - g^{\circ r}(0)} \prod_{v \in t} \theta(c_v),$$

which concludes. \square

Now let $r \geq 0$. By Lemma 2.14, there is $r' \geq r$ such that all the vertices of reverse height at most r lie in $d_{r'}(\boldsymbol{\tau}^0)$. Hence, $\boldsymbol{\tau}^0$ is the a.s. local limit as $r' \rightarrow +\infty$ of $d_{r'}(\boldsymbol{\tau}^0)$ rooted at its leftmost vertex of height r' . By Lemma 2.15, this proves that $\boldsymbol{\tau}^0$ is the local limit (in distribution) as $r' \rightarrow +\infty$ of Galton–Watson trees conditioned on extinction at time r' , seen from the last generation.

In particular, for every $r \geq 0$, let $Y(r)$ be the number of vertices of $\boldsymbol{\tau}^0$ at reverse height r . Then $(Y(-r))_{r \leq 0}$ is the limit in distribution as $n \rightarrow +\infty$ of a Galton–Watson process with offspring distribution θ , started from 1 at time $-n$ and conditioned on extinction at time exactly 1. Hence, Y has the same distribution as the reverse Galton–Watson process described by Esty in [22], started from 0 at time -1 . In particular, by Equation 3 of [22], we obtain explicitly the distribution of $Y(r)$.

Lemma 2.16. For all $r \geq 0$ and $p \geq 1$, we have

$$\mathbb{P}(Y(r) = p) = \frac{\pi(p)m^{-r}}{\Pi(\theta(0))} \left(g^{\circ(r+1)}(0)^p - g^{\circ r}(0)^p \right).$$

We also define $\boldsymbol{\tau}^1$ as the tree $\boldsymbol{\tau}^0$, conditioned on $Y(0) = 1$. We can now prove Lemma 2.11.

Proof of Lemma 2.11. We start with the first part of the lemma (i.e. the part related to $\boldsymbol{\tau}^0$).

Let R be the smallest r' such that $s_{r'}$ is an ancestor of all the vertices at reverse height r in $\boldsymbol{\tau}^0$. For every $r' > r$, the event $\{R \leq r'\}$ depends only on the trees t_j^i and u_j^i for $i > r'$, so it is independent of the tree $d_{r'}(\boldsymbol{\tau}^0)$. We write $Y(r, r')$ for the number of descendants of $s_{r'}$ at reverse height r , and $B_r^{r'}(\boldsymbol{\tau}^0)$ for the forest consisting of the trees of descendants of these vertices. We have

$$\begin{aligned} \mathbb{P}(B_r(\boldsymbol{\tau}^0) = (t_1, \dots, t_p)) &= \lim_{r' \rightarrow +\infty} \mathbb{P}(R \leq r', B_r(\boldsymbol{\tau}^0) = (t_1, \dots, t_p)) \\ &= \lim_{r' \rightarrow +\infty} \mathbb{P}(R \leq r', B_r^{r'}(\boldsymbol{\tau}^0) = (t_1, \dots, t_p)) \\ &= \lim_{r' \rightarrow +\infty} \mathbb{P}(R \leq r') \mathbb{P}(Y(r, r') = p) \\ &\quad \times \mathbb{P}(B_r^{r'}(\boldsymbol{\tau}^0) = (t_1, \dots, t_p) | Y(r, r') = p). \end{aligned}$$

But by Lemma 2.15, the tree $d_{r'}(\tau^0)$ has the same distribution as a Galton–Watson tree conditioned to have height r' . From here, it is easy to show that the distribution of $B_r^{r'}(\tau^0)$ conditioned on $Y(r, r') = p$ is that of a Galton–Watson forest of p trees conditioned to have height r , i.e.

$$\mathbb{P}\left(B_r^{r'}(\tau^0) = (t_1, \dots, t_p) \mid Y(r, r') = p\right) = \frac{1}{g^{\circ(r+1)}(0)^p - g^{\circ r}(0)^p} \prod_{i=1}^p \prod_{v \in t_i} \theta(c_v).$$

On the other hand, by using again the independence of $\{R \leq r'\}$ and $d_{r'}(\tau^0)$, we have

$$\begin{aligned} \lim_{r' \rightarrow +\infty} \mathbb{P}(R \leq r') \mathbb{P}(Y(r, r') = p) &= \lim_{r' \rightarrow +\infty} \mathbb{P}(R \leq r', Y(r, r') = p) \\ &= \lim_{r' \rightarrow +\infty} \mathbb{P}(R \leq r', Y(r) = p) \\ &= \mathbb{P}(Y(r) = p) \\ &= \frac{\pi(p)m^{-r}}{\Pi(\theta(0))} \left(g^{\circ(r+1)}(0)^p - g^{\circ r}(0)^p \right) \end{aligned}$$

by Lemma 2.16, which proves the first part of the lemma.

Given the definition of τ^1 , the proof of the second part is now easy. For every forest (t_1, \dots, t_p) of height r with exactly one vertex at reverse height 0, we have

$$\begin{aligned} \mathbb{P}(B_r(\tau^1) = (t_1, \dots, t_p)) &= \mathbb{P}(B_r(\tau^0) = (t_1, \dots, t_p) \mid Y(0) = 1) \\ &= \frac{\mathbb{P}(B_r(\tau^0) = (t_1, \dots, t_p))}{\mathbb{P}(Y(0) = 1)}. \end{aligned}$$

By Lemma 2.16 we have $\mathbb{P}(Y(0) = 1) = \frac{\theta(0)}{\Pi(\theta(0))}$, so the second part of the lemma follows from the first one. \square

We end up this subsection by showing that the strips S_λ^0 and S_λ^1 constructed from τ_λ^0 and τ_λ^1 are in some sense very close to each other. This will allow us in Section 4 to conclude a strip verifies a property if the other does, and therefore to avoid annoying case distinctions.

Lemma 2.17. Let $\lambda_n \rightarrow \lambda_c$ and let (A_n) be measurable events. Then the following two assertions are equivalent:

- (i) $\mathbb{P}(S_{\lambda_n}^0 \in A_n) \xrightarrow{n \rightarrow +\infty} 0$,
- (ii) $\mathbb{P}(S_{\lambda_n}^1 \in A_n) \xrightarrow{n \rightarrow +\infty} 0$.

Proof. We will show that the strips $S_{\lambda_n}^0$ and $S_{\lambda_n}^1$ are absolutely continuous with respect to each other, uniformly in n . For this, we recall that $\tau^{1,*}$ is the tree τ^1 in which we have cut the only vertex at reverse height 0, and we have shifted the reverse heights by 1. We will first prove that $\tau^{1,*}$ has the same distribution as τ^0 biased by $Y(0)$, and then extend this absolute continuity relation to the strips.

Indeed, for any forest $f \in \mathcal{F}_{p,q,r}$, we have

$$\mathbb{P}(B_r(\tau^{1,*}) = f) = \sum_{v \in f \setminus f^*} \mathbb{P}(B_{r+1}(\tau^1) = f_v^+),$$

where $f \setminus f^*$ is the set of vertices of f at reverse height 0, and f_v^+ is the forest of $\mathcal{F}_{1,q,r+1}$ obtained by adding a unique child to v in f . By the second part of Lemma 2.11, we obtain

$$\mathbb{P}(B_r(\tau^{1,*}) = f) = \frac{\theta(1)}{\theta(0)} p \pi(q) m^{-r-1} \prod_{v \in f} \theta(c_v).$$

Combined with the first part of Lemma 2.11, this yields

$$\frac{\mathbb{P}(B_r(\boldsymbol{\tau}^{1,*}) = f)}{\mathbb{P}(B_r(\boldsymbol{\tau}^0) = f)} = \frac{\theta(1)}{\theta(0)} \frac{\Pi(\theta(0))}{m} p. \quad (27)$$

In other words, the forest $B_r(\boldsymbol{\tau}^{1,*})$ has the distribution of $B_r(\boldsymbol{\tau}^0)$, biased by $Y(0)$. Since it is true for all r , we can conclude that $\boldsymbol{\tau}^{1,*}$ has the distribution of $\boldsymbol{\tau}^0$ biased by $Y(0)$. Since S^1 is constructed from $\boldsymbol{\tau}^{1,*}$ in the exact same way as S^0 is constructed from $\boldsymbol{\tau}^0$, we deduce that $S_{\lambda_n}^1$ has the same distribution as $S_{\lambda_n}^0$ biased by $Y_{\lambda_n}(0)$. It remains to prove that this absolute continuity is "uniform in n ".

More precisely, if (A_n) is a sequence of measurable events, for any $n \geq 0$, we have

$$\mathbb{P}(S_{\lambda_n}^1 \in A_n) = \frac{\mathbb{E}\left[Y_{\lambda_n}(0) \mathbb{1}_{S_{\lambda_n}^0 \in A_n}\right]}{\mathbb{E}[Y_{\lambda_n}(0)]}. \quad (28)$$

Moreover, (27) shows that

$$\mathbb{E}[Y_\lambda(0)] = \frac{\theta_\lambda(0)}{\theta_\lambda(1)} \frac{m_\lambda}{\Pi_\lambda(\theta_\lambda(0))} = \frac{1-h}{2h(1-h)} \frac{1-2h-\sqrt{1-4h}}{1-\sqrt{1-4h}},$$

which is continuous at $\lambda = \lambda_c$. Therefore, the denominator in (28) converges to a positive limit, so it is enough to prove that $\mathbb{P}(S_{\lambda_n}^0 \in A_n) \rightarrow 0$ if and only if $\mathbb{E}\left[Y_{\lambda_n}(0) \mathbb{1}_{S_{\lambda_n}^0 \in A_n}\right] \rightarrow 0$. The indirect implication is immediate since $Y_{\lambda_n}(0) \geq 1$. To prove the direct one, it is enough to check the variables $Y_{\lambda_n}(0)$ are uniformly integrable. We know that they converge to $Y_{\lambda_c}(0)$ in distribution. By the Skorokhod representation theorem, we may assume the convergence is almost sure. Since we also have convergence of the expectations by (28), by Scheffé's Lemma we have $Y_{\lambda_n}(0) \rightarrow Y_{\lambda_c}(0)$ in L^1 . In particular, the family $(Y_{\lambda_n}(0))$ is uniformly integrable, which proves the direct implication and finally the lemma. \square

3 The Poisson boundary of \mathbb{T}_λ

The goal of this section is to prove Theorem 2. We fix $0 < \lambda < \lambda_c$ until the end of the section and omit the parameter λ in most of the notation.

3.1 Construction of the geodesic boundary

We start by defining precisely the compactification of \mathbb{T} that we will afterwards prove to be a realization of its Poisson boundary. We recall that $\partial\mathbf{T}^g$ is the set of ends of the tree \mathbf{T}^g , i.e. the set of infinite self-avoiding paths from the root in \mathbf{T}^g . If $\gamma, \gamma' \in \partial\mathbf{T}^g$, we write $\gamma \sim \gamma'$ if $\gamma = \gamma'$ or if γ and γ' are the left and right boundaries of one of the copies of S^0 or S^1 (cf. left part of Figure 1). Note that a.s., every ray of \mathbf{T}^g branches infinitely many times, so no ray is equivalent to two distinct other rays. It follows that \sim is a.s. an equivalence relation, for which countably many equivalence classes have cardinal 2, and all the others have cardinal 1. We write $\widehat{\partial\mathbf{T}^g} = \partial\mathbf{T}^g / \sim$, i.e. we identify two geodesic rays if they are the left and right boundaries of the same strip. We also write $\gamma \rightarrow \widehat{\gamma}$ for the canonical projection from $\partial\mathbf{T}^g$ to $\widehat{\partial\mathbf{T}^g}$. If S is one of the copies of S^0 or S^1 appearing in the strip decomposition, then the two geodesics bounding S correspond to the same point of $\widehat{\partial\mathbf{T}^g}$, that we denote by $\widehat{\gamma}_S$.

Our goal is now to define a topology on $\mathbb{T} \cup \widehat{\partial\mathbf{T}^g}$. It should be possible to define it by an explicit distance, but such a distance would be tedious to write down, so we prefer to give an "abstract" construction. Let S and S' be two distinct strips appearing in the strip decomposition of \mathbb{T} (cf. left part of Figure 1). Consider the smallest r such that S and S' both intersect $B_r^\bullet(\mathbb{T})$. Then $\mathbb{T} \setminus (B_r^\bullet(\mathbb{T}) \cup S \cup S')$ has two connected components, that we denote by (S, S') and (S', S)

(the vertices on the geodesics bounding S and S' do not belong to (S, S') and (S', S)). We also write

$$\partial_g(S, S') = \{\widehat{\gamma} \mid \gamma \text{ is a ray of } \mathbf{T}^g \text{ such that } \gamma(i) \in (S, S') \text{ for } i \text{ large enough}\}$$

We define $\partial_g(S', S)$ similarly. Note that $\partial_g(S, S')$ and $\partial_g(S', S)$ are disjoint, and their union is $\widehat{\partial\mathbf{T}}^g \setminus \{\widehat{\gamma}_S, \widehat{\gamma}_{S'}\}$.

Definition 3.1. The *geodesic compactification* of \mathbb{T} is the set $\mathbb{T} \cup \widehat{\partial\mathbf{T}}^g$, equipped with the topology generated by

- the singletons $\{x\}$, where x is a vertex of \mathbb{T} ,
- the sets $(S, S') \cup \partial_g(S, S')$, where S and S' are two distinct strips appearing in the strip decomposition of \mathbb{T} .

This topology is separated (if $\widehat{\gamma}_1 \neq \widehat{\gamma}_2$, then there are two strips separating γ_1 and γ_2) and has a countable basis, so it is induced by a distance. Moreover, any open set of our basis intersects \mathbb{T} , so \mathbb{T} is dense in $\mathbb{T} \cup \widehat{\partial\mathbf{T}}^g$. The end of this subsection is devoted to the proof of two very intuitive topological properties of the geodesic compactification. The second one states that the boundary $\widehat{\partial\mathbf{T}}^g$ is homeomorphic to the circle, which is natural since this is also the standard topology on the space of ends quotiented by our equivalence relation.

Lemma 3.2. The space $\mathbb{T} \cup \widehat{\partial\mathbf{T}}^g$ is compact.

Proof. Let (x_n) be a sequence with values in $\mathbb{T} \cup \widehat{\partial\mathbf{T}}^g$. We first assume $x_n \in \mathbb{T}$ for every n . We may assume $d(\rho, x_n) \rightarrow +\infty$. If there is a strip S that contains infinitely many x_n , then $\widehat{\gamma}_S$ is a subsequential limit of (x_n) . We now assume it is not the case. We recall that for every vertex v of \mathbf{T}^g , the slice $\mathbb{S}[v]$ is the part of \mathbb{T} lying between the leftmost and the rightmost rays passing through v , above v (see Figure 1). We will construct a ray γ of \mathbf{T}^g step by step, in such a way that for every $k \geq 0$, there are infinitely many points x_n in $\mathbb{S}[\gamma(k)]$.

Assume we have already built $\gamma(0), \dots, \gamma(k)$. If $\gamma(k)$ has only one child in \mathbf{T}^g , then $\gamma(k+1)$ must be this child. If $\gamma(k)$ has $d \geq 2$ children, we call them y_1, \dots, y_d . Then $\mathbb{S}[\gamma(k)]$ is the union of the slices $\mathbb{S}[y_i]$ for $1 \leq i \leq d$ and of the $d-1$ strips whose lowest vertex is $\gamma(k)$. We know that $\mathbb{S}[\gamma(k)]$ contains infinitely many of the vertices x_n , but the $d-1$ strips contain finitely many of them. Therefore, there is an index $1 \leq i_0 \leq d$ such that $\mathbb{S}[y_{i_0}]$ contains infinitely many of them. We choose $\gamma(k+1) = y_{i_0}$. We can check that the class $\widehat{\gamma}$ of the geodesic we built is a subsequential limit of (x_n) , which concludes the case where $x_n \in \mathbb{T}$ for every n .

Finally, let δ be a distance on $\mathbb{T} \cup \widehat{\partial\mathbf{T}}^g$ that generates its topology, and let (x_n) be any sequence in $\mathbb{T} \cup \widehat{\partial\mathbf{T}}^g$. If $x_n \in \widehat{\partial\mathbf{T}}^g$, let $y_n \in \mathbb{T}$ be such that $\delta(x_n, y_n) \leq \frac{1}{n}$ (it exists by density). If $x_n \in \mathbb{T}$, we take $y_n = x_n$. By the first case (y_n) has a subsequential limit, so (x_n) also has one. \square

Lemma 3.3. The boundary $\widehat{\partial\mathbf{T}}^g$ is homeomorphic to the circle.

Proof. We build an explicit homeomorphism. Consider a ray γ of \mathbf{T}^g . For every $i \geq 0$, let $c_\gamma(i)$ be the number of children of $\gamma(i)$ in \mathbf{T}^g . We denote these children by $x_0, \dots, x_{c_\gamma(i)-1}$ and we denote by $j_\gamma(i)$ the index $j \in \{0, 1, \dots, c_\gamma(i)-1\}$ such that $\gamma(i+1) = x_j$. We also define

$$\begin{aligned} \Psi : \partial\mathbf{T}^g &\longrightarrow \mathbb{R}/\mathbb{Z} \\ \gamma &\longrightarrow \sum_{k \geq 0} \frac{j_\gamma(k)}{\prod_{i=0}^k c_\gamma(i)} \pmod{1}. \end{aligned}$$

If γ_ℓ and γ_r are the left and right boundaries of a copy of S^0 , then there is i_0 such that:

- for $i < i_0$, we have $c_{\gamma_r}(i) = c_{\gamma_\ell}(i)$ and $j_{\gamma_r}(i) = j_{\gamma_\ell}(i)$,
- we have $c_{\gamma_r}(i_0) = c_{\gamma_\ell}(i_0)$ and $j_{\gamma_r}(i_0) = j_{\gamma_\ell}(i_0) + 1$,

- for $i > i_0$, we have $j_{\gamma_\ell}(i) = c_{\gamma_\ell}(i) - 1$ and $j_{\gamma_r}(i) = 0$,

which implies $\Psi(\gamma_\ell) = \Psi(\gamma_r)$. Moreover, if γ_ℓ and γ_r are respectively the leftmost and rightmost rays of \mathbf{T}^g , then $\Psi(\gamma_\ell) = \Psi(\gamma_r)$ (the sum is equal to 0 for γ_ℓ and to 1 for γ_r). Hence, we have $\Psi(\gamma) = \Psi(\gamma')$ as soon as $\gamma \sim \gamma'$, so Ψ defines an application from $\widehat{\partial\mathbf{T}^g}$ to \mathbb{R}/\mathbb{Z} . The verification that this application is a homeomorphism is easy and left to the reader. \square

3.2 Proof of Theorem 2

We first argue why the second point of Theorem 2 is an easy consequence of the first one. Assume the first point is proved. Since $\widehat{\partial\mathbf{T}^g}$ is homeomorphic to the circle, we can embed $\mathbb{T} \cup \widehat{\partial\mathbf{T}^g}$ in the unit disk \mathbb{D} in such a way that $\widehat{\partial\mathbf{T}^g}$ is sent to $\partial\mathbb{D}$ (we do not describe the embedding explicitly). In this embedding, the simple random walk converges to a point of $\partial\mathbb{D}$ and the law of the limit point is a.s. non-atomic. Therefore, by Theorem 1.3 of [24], $\widehat{\partial\mathbf{T}^g}$ is a realization of the Poisson boundary of \mathbb{T} .

The idea of the proof of the first point is to first show that two independent random walks are quite well separated in terms of $\widehat{\partial\mathbf{T}^g}$. Let X^1 and X^2 be two independent simple random walks started from ρ . By Proposition 11 of [16] we know that \mathbb{T} is transient and does not have the intersection property (although [16] deals with type-II triangulations, all the arguments of the proof still hold in our case). Hence, the complement of $\{X_n^1 | n \in \mathbb{N}\} \cup \{X_n^2 | n \in \mathbb{N}\}$ has two infinite connected components with infinite boundaries. We denote them by $[X^1, X^2]$ and $[X^2, X^1]$. By point (ii) of Lemma 1.2 applied to the peeling along X_1 and X_2 (see also Section 3.2 of [16] for a similar argument), the halfplanar triangulations $[X^1, X^2]$ and $[X^2, X^1]$ are independent copies of $\mathbb{H} = \mathbb{H}_\lambda$.

Therefore, we will need to study geodesics in halfplanar triangulations.

Definition 3.4. Let H be a halfplanar triangulation and ∂H its boundary. An *infinite geodesic away from the boundary* is a sequence $(\gamma(n))_{n \geq 0}$ of vertices of H such that for any $n \geq 0$:

- (i) the vertices $\gamma(n)$ and $\gamma(n+1)$ are neighbours,
- (ii) we have $d(\gamma(n), \partial H) = n$.

Remark 3.5. Contrary to geodesics away from a point, the existence of such geodesics is not obvious. For example, they do not exist in the UIHPT \mathbb{H}_{λ_c} .

Lemma 3.6. Almost surely, there is a geodesic away from the boundary in \mathbb{H}_λ .

Proof. We write $P_r = |\partial B_r^\bullet(\mathbb{T})|$ and $L_r = |\partial B_r^\bullet(\mathbb{T}) \cap \mathbf{T}^g|$. We recall that \mathbf{T}^g is a Galton–Watson tree with offspring distribution μ given by Theorem 1. In particular, we have $\sum_{i \geq 0} i\mu(i) = m_\lambda^{-1}$ and $\sum_{i \geq 0} (i \log i)\mu(i) < +\infty$, so by the Kesten–Stigum Theorem $m_\lambda^r L_r$ converges a.s. to a positive random variable. On the other hand, as in Section 2 of [16], it holds that $m_\lambda^r P_r$ converges a.s. to a positive random variable ([16] only deals with type-II triangulations but all the arguments of the proof still work in the type-I case). Hence, there is a constant $c > 0$ such that, for r large enough,

$$\mathbb{P}(L_r \geq cP_r) \geq \frac{1}{2}.$$

Therefore, if z_r is a random vertex chosen uniformly on $\partial B_r^\bullet(\mathbb{T})$, for r large enough we have $\mathbb{P}(z_r \in \mathbf{T}^g) \geq \frac{c}{2}$. Hence, for any $s > 0$, with probability at least $\frac{c}{2}$, there is a point $x \in \partial B_{r+s}^\bullet(\mathbb{T})$ at distance exactly s from z_r .

But \mathbb{H} is the local limit as $r \rightarrow +\infty$ of $\mathbb{T} \setminus B_r^\bullet(\mathbb{T})$ (rooted at a uniform edge on its boundary), so if ρ denotes the root vertex of \mathbb{H} , for any $s \geq 0$, we have

$$\mathbb{P}(\text{there is } x \in \mathbb{H} \text{ such that } d(x, \rho) = d(x, \partial\mathbb{H}) = s) \geq \frac{c}{2}.$$




Figure 9: From an infinite geodesic away from the boundary in H , we can build an infinite leftmost geodesic in \mathbb{T} , that lies in H eventually.

This event is nonincreasing in s , so with probability at least $\frac{\epsilon}{2}$ it occurs for every s . By a compactness argument, with probability at least $\frac{\epsilon}{2}$, there is an infinite geodesic away from the boundary γ with $\gamma(0) = \rho$. Finally, we claim that \mathbb{H} is invariant under root translation, and that the root translation is ergodic, which is enough to conclude. Indeed, by local limit, the invariance is a consequence of the invariance of \mathbb{T}^p for every p under re-rooting along the boundary. The ergodicity is proved in the type-II case in [7] (this is Proposition 1.3), and the proof adapts well here. \square

We can now show that X^1 and X^2 are a.s. separated by an infinite leftmost geodesic.

Lemma 3.7. Almost surely, there is a ray $\hat{\gamma}$ of \mathbf{T}^g such that for n large enough, we have $\hat{\gamma}(n) \in [X^1, X^2]$, and the same is true for $[X^2, X^1]$.

Proof. In this proof, we will write $H = [X^1, X^2]$ to avoid too heavy notations. Let γ be an infinite geodesic away from the boundary in H , which exists by Lemma 3.6. For $n \geq 0$, let $(\gamma_n(i))_{0 \leq i \leq d(\rho, \gamma(n))}$ be the leftmost geodesic from ρ to $\gamma(n)$ (cf. Figure 9). For $n \geq i \geq d(\rho, \gamma(0))$, we have $i \leq n = d(\gamma(n), \partial H) \leq d(\gamma(n), \rho)$ so $\gamma_n(i)$ is well defined, and

$$\begin{aligned} d(\gamma(n), \gamma_n(i)) &= d(\rho, \gamma(n)) - i && \text{(since } \gamma \text{ is a geodesic)} \\ &\leq n + d(\rho, \gamma(0)) - i && \text{(by the triangular inequality)} \\ &\leq n \\ &= d(\gamma(n), \partial H), \end{aligned}$$

so $\gamma_n(i) \in H$. But by a compactness argument, there is an infinite path $\tilde{\gamma}$ in \mathbb{T} such that for any i , there are infinitely many n such that $\tilde{\gamma}(i) = \gamma_n(i)$. It is easy to check that $\tilde{\gamma}$ is an infinite leftmost geodesic in \mathbb{T} . Moreover, since $\gamma_n(i) \in H$ for $n \geq i \geq d(\rho, \gamma(0))$, we have $\tilde{\gamma}(i) \in H$ for i large enough, so there is an infinite leftmost geodesic of \mathbb{T} that lies in H eventually. \square

Proof of Theorem 2. We can now prove the convergence of the simple random walk to a point of $\hat{\partial}\mathbf{T}^g$. If X is a simple random walk, we write $\text{Acc}(X)$ for the set of accumulation points of X on $\hat{\partial}\mathbf{T}^g$. By Lemma 3.2, it is enough to prove that $\text{Acc}(X)$ is reduced to a point. We first claim that $\text{Acc}(X)$ is a circle arc of $\hat{\partial}\mathbf{T}^g$. Indeed, assume $\hat{\gamma}_1 \neq \hat{\gamma}_2$ are two points of $\text{Acc}(X)$. Then $\hat{\partial}\mathbf{T}^g \setminus \{\hat{\gamma}_1, \hat{\gamma}_2\}$ has two connected components, that we denote by $(\hat{\gamma}_1, \hat{\gamma}_2)$ and $(\hat{\gamma}_2, \hat{\gamma}_1)$. To oscillate between $\hat{\gamma}_1$ and $\hat{\gamma}_2$, the walk X must intersect infinitely many times either all the γ such that $\hat{\gamma} \in (\hat{\gamma}_1, \hat{\gamma}_2)$ or all the γ such that $\hat{\gamma} \in (\hat{\gamma}_2, \hat{\gamma}_1)$ (see Figure 10). In both cases, $\text{Acc}(X)$




Figure 10: Proof that $\text{Acc}(X)$ is an arc circle: assume X oscillates between γ_1 and γ_2 and there is $\gamma' \in (\widehat{\gamma}_2, \widehat{\gamma}_1)$ that X intersects only finitely many times. Then X intersects infinitely many times every γ with $\widehat{\gamma} \in (\widehat{\gamma}_1, \widehat{\gamma}_2)$.

contains one of the two arcs from $\widehat{\gamma}_1$ to $\widehat{\gamma}_2$. Hence, $\text{Acc}(X)$ is closed and connected, so it is a circle arc.

Now let ν be a probability measure with no atom and full support on $\widehat{\partial}\mathbf{T}^g$ (for example, one can consider the exit measure of the nonbacktracking random walk on \mathbf{T}^g). Either $\text{Acc}(X)$ is a singleton, or it has positive measure, so it is enough to show that $\mathbb{P}(\nu(\text{Acc}(X)) > 0) = 0$. By Lemma 3.7, we know that if X^1 and X^2 are two independent simple random walks started from ρ , then there are two rays of \mathbf{T}^g lying respectively in $[X_1, X_2]$ and $[X_2, X_1]$ eventually, so they separate X^1 and X^2 . Therefore, $\text{Acc}(X^1) \cap \text{Acc}(X^2)$ contains at most two points, so $\nu(\text{Acc}(X^1) \cap \text{Acc}(X^2)) = 0$ a.s.. Let X^i for $i \in \mathbb{N}$ be infinitely many independent simple random walks started from ρ . We have

$$\sum_{i \geq 0} \nu(\text{Acc}(X^i)) = \nu\left(\bigcup_{i \geq 0} \text{Acc}(X^i)\right) \leq 1.$$

But the $\nu(\text{Acc}(X^i))$ are i.i.d. (conditionally on \mathbb{T}), so they must be 0 a.s.. Therefore, $\text{Acc}(X)$ cannot have positive measure so it is a.s. a point.

Hence, the simple random walk a.s. converges to a point of $\widehat{\partial}\mathbf{T}^g$, so it defines an exit measure on $\widehat{\partial}\mathbf{T}^g$, that we denote by ν_∂ . We now prove that ν_∂ is nonatomic. Once again, our main tool is Lemma 3.7. Assume ν_∂ has an atom with positive probability, and let $(X^i)_{1 \leq i \leq 4}$ be four independent SRW started from ρ . For every $1 \leq i \leq 4$, let X_∞^i be the limit of X^i in $\widehat{\partial}\mathbf{T}^g$. Then $\mathbb{P}(X_\infty^1 = X_\infty^2 = X_\infty^3 = X_\infty^4) > 0$. If this happens, we can assume (up to a factor $\frac{1}{24}$) that they lie in clockwise order, and that $\widehat{\partial}\mathbf{T}^g \setminus \{X_\infty^1\}$ lies in the part between X^4 and X^1 . By Lemma 3.7, there are at least one ray of \mathbf{T}^g between X^1 and X^2 , one between X^2 and X^3 and one between X^3 and X^4 (cf. Figure 11). Hence, there are at least three rays of \mathbf{T}^g in the part between X^1 and X^4 that does not contain $\widehat{\partial}\mathbf{T}^g \setminus \{X_\infty^1\}$. In particular, two of them are not equivalent for \sim (we recall that the equivalence classes have cardinal at most 2). Hence, the part lying between X^1 and X^4 containing X^2 and X^3 (the left part on Figure 11) also contains a slice of the form $\mathbb{S}[y]$ (see Figure 1), so $X_\infty^1 \neq X_\infty^4$. We get a contradiction.

We finally show that ν_∂ has full support. Since ν_∂ is nonatomic, we have $\nu_\partial(\{\widehat{\gamma}_{S^1}\}) = 0$ a.s..




Figure 11: Sketch of the proof that ν_∂ is nonatomic: if X^1, \dots, X^4 are four independent simple random walks, then there are three infinite leftmost geodesics (in red) between X^1 and X^4 , so there is a slice $\mathbb{S}[y]$ there. The hulls of the four random walks are in green.

Hence, a.s., we have $X_n \in \mathbb{S}[\rho] = \mathbb{S}$ for n large enough. Therefore, we also have

$$\mathbb{P}(\forall k \geq 0, X_k \in \mathbb{S}[\rho] | X_0 = \rho) > 0 \quad \text{a.s.}$$

Now fix $r > 0$, condition on $B_r^\bullet(\mathbb{T})$ and take $x \in \mathbf{T}^g$ such that $d(\rho, x) = r$. Since $\mathbb{S}[x]$ has the same distribution as $\mathbb{S}[\rho]$ (see Figure 1), we have

$$\mathbb{P}(\forall k \geq 0, X_k \in \mathbb{S}[x] | X_0 = x) > 0 \quad \text{a.s.}$$

Hence, we have $\mathbb{P}(X_k \in \mathbb{S}[x] \text{ for } k \text{ large enough} | X_0 = \rho) > 0$ a.s. If this occurs, then X_∞ is of the form $\hat{\gamma}$, where $\gamma(k) \in \mathbb{S}[x]$ for k large enough. Therefore, we have

$$\nu_\partial(\{\hat{\gamma} | \gamma \text{ lie in } \mathbb{S}[x] \text{ eventually}\}) > 0 \quad \text{a.s.}$$

Almost surely, this holds for every $x \in \mathbf{T}^g$, which is enough to ensure that ν_∂ has full support. This ends the proof of Theorem 2. \square

Remark 3.8. We end this section with a remark about the Gromov boundary [23], which is another natural notion of boundary for an infinite graph $G = (V, E)$. Let $\mathcal{C}(G)$ be the space of functions $f : V \rightarrow \mathbb{R}$ equipped with the product topology. We say that two functions of $\mathcal{C}(G)$ are equivalent if they are equal up to an additive constant. We quotient $\mathcal{C}(G)$ by this equivalence relation to obtain the quotient space $\mathcal{C}(G)/\mathbb{R}$. If $x \in V$, we define $f_x \in \mathcal{C}(G)/\mathbb{R}$ by $f_x(y) = d_G(x, y)$ for any $y \in V$. The *Gromov compactification* \hat{G} of G is the closure of $\{f_x | x \in V\}$ in $\mathcal{C}(G)/\mathbb{R}$ and the *Gromov boundary* $\partial_{G_r} G$ of G is the set $\hat{G} \setminus \{f_x | x \in V\}$.

It is easy to show that for any geodesic ray γ , the sequence $f_{\gamma(n)}$ converges in $\mathcal{C}(G)/\mathbb{R}$, so it defines a point $f_\gamma \in \partial_{G_r} \mathbb{T}$. A natural question is to ask whether there is a natural correspondence between $\hat{\partial} \mathbf{T}^g$ and $\partial_{G_r} \mathbb{T}$. The answer is no.

To prove it, we show that if γ_1 and γ_2 are respectively the left and right boundary of the same strip, then $f_{\gamma_1} \neq f_{\gamma_2}$. Indeed, let n be such that $\gamma_1(n) \notin \gamma_2$. We have $f_{\gamma_1}(\gamma_1(n)) - f_{\gamma_1}(\rho) = -n$ by definition of f_{γ_1} . But if we had $f_{\gamma_2}(\gamma_1(n)) - f_{\gamma_2}(\rho) = -n$, this would mean that there is $m > n$ such that $d(\gamma_2(m), \gamma_1(n)) - d(\gamma_2(m), \rho) = -n$, i.e. $d(\gamma_2(m), \gamma_1(n)) = m - n$. Take such an m minimal. Then by concatenating γ_1 from ρ to $\gamma_1(n)$ and a geodesic from $\gamma_1(n)$ to $\gamma_2(m)$, we obtain a geodesic from ρ to $\gamma_2(m)$ that lies strictly to the left of γ_2 . This contradicts the fact that γ_2 is a leftmost geodesic in \mathbb{T} . This suggests that $\partial_{G_r} \mathbb{T}$ should not be homeomorphic to the circle, but rather to a Cantor set.




Figure 12: Proposition 4.2 : almost surely, for any x , none of these two cases occurs.

4 The tree of infinite geodesics in the hyperbolic Brownian plane

4.1 The tree $\mathbf{T}^g(\mathcal{P}^h)$

The goal of this section is to take the scaling limit of Theorem 1 and to prove Theorem 3. For all this section, we fix a sequence (λ_n) of numbers in $(0, \lambda_c]$ such that $\lambda_n = \lambda_c \left(1 - \frac{2}{3n^4}\right) + o\left(\frac{1}{n^4}\right)$. We know, by the main result of [14], that $\frac{1}{n}\mathbf{T}_{\lambda_n}$ converges for the local Gromov–Hausdorff distance to \mathcal{P}^h . Therefore, it seems reasonable that $\mathbf{T}^g(\mathcal{P}^h)$ should be the scaling limit of the trees $\mathbf{T}_{\lambda_n}^g$. This scaling limit is easy to describe. We recall that \mathbf{B} is the infinite tree in which every vertex has exactly two children, except the root which has one. For $\alpha > 0$, we denote by \mathbf{Y}_α the Yule tree of parameter α , i.e. the tree \mathbf{B} in which the lengths of the edges are i.i.d. exponential variables with parameter α .

Lemma 4.1. The trees $\frac{1}{n}\mathbf{T}_{\lambda_n}^g$ converge for the local Gromov–Hausdorff distance to $\mathbf{Y}_{2\sqrt{2}}$.

Sketch of proof. We recall that $\mathbf{T}_{\lambda_n}^g$ is a Galton–Watson tree with offspring distribution μ_{λ_n} , where $\mu_{\lambda_n}(k) = m_{\lambda_n} (1 - m_{\lambda_n})^{k-1}$ for every $k \geq 1$, and m_λ is explicitly given by (2). We can compute $m_{\lambda_n} = 1 - \frac{2\sqrt{2}}{n} + O\left(\frac{1}{n^2}\right)$. Let $\tilde{\mu}_\lambda$ be the distribution defined by $\tilde{\mu}_\lambda(1) = m_\lambda$ and $\tilde{\mu}_\lambda(2) = 1 - m_\lambda$, and let $\tilde{\mathbf{T}}_\lambda^g$ be a Galton–Watson tree with offspring distribution $\tilde{\mu}_\lambda$. Then $\tilde{\mathbf{T}}_\lambda^g$ is a copy of \mathbf{B} where the length of each edge is a geometric variable of parameter $1 - m_\lambda$, so $\frac{1}{n}\tilde{\mathbf{T}}_{\lambda_n}^g$ converges to $\mathbf{Y}_{2\sqrt{2}}$. Moreover, \mathbf{T}_λ^g can be obtained by adding children to some of the vertices of $\tilde{\mathbf{T}}_\lambda^g$ with two children. Since $\mu_{\lambda_n}([3, +\infty[) = O\left(\frac{1}{n^2}\right)$, the probability to affect a vertex is of order $\frac{1}{n^2}$. The number of vertices at height of order n in $\mathbf{T}_{\lambda_n}^g$ is of order n , so the difference between $\mathbf{T}_{\lambda_n}^g$ and $\tilde{\mathbf{T}}_{\lambda_n}^g$ does not affect the scaling limit. \square

However, taking the scaling limit of $\mathbf{T}_{\lambda_n}^g$ is not enough to obtain a description of infinite geodesics in \mathcal{P}^h . Three different kinds of problems could prevent this:

- (i) it is not completely clear that the infinite geodesics in \mathcal{P}^h form a tree,
- (ii) two different discrete leftmost geodesics might be too close and collide in the scaling limit,
- (iii) discrete paths that are not infinite leftmost geodesics might become infinite geodesics in the scaling limit.

We take care of item (i) right now, while the goal of Lemmas 4.3 and 4.5 will be to rule out items (ii) and (iii).

The fact that the infinite geodesics of \mathcal{P}^h indeed form a tree is a quite strong result, that follows from the confluence of geodesics properties in the Brownian map. More precisely, it will be a consequence of [5, Proposition 28], which we recall here. We write m_∞ for the Brownian map, and denote its root by ρ .

Proposition 4.2. [5] Almost surely, for any $x \in m_\infty$, if γ and γ' are two geodesics from ρ to x that coincide on a neighbourhood of x , then $\gamma = \gamma'$.

We write $\mathbf{T}^g(\mathcal{P}^h)$ for the set of those points of \mathcal{P}^h that lie on an infinite geodesic of \mathcal{P}^h started from ρ . By local isometry of the Brownian plane and the Brownian map and scale

invariance of the Brownian plane [17], Proposition 4.2 also holds for the Brownian plane. By the absolute continuity relations of [14], it also holds for \mathcal{P}^h . We claim that it implies that for every $x \in \mathbf{T}^g(\mathcal{P}^h)$, there is a unique geodesic from ρ to x in \mathcal{P}^h . Indeed, let $x \in \mathbf{T}^g(\mathcal{P}^h)$ and let γ be an infinite geodesic of \mathcal{P}^h passing through x . Let γ' be a geodesic from ρ to x . Finally, let y be a point on the ray γ such that $d(\rho, y) > d(\rho, x)$. The concatenation γ'' of γ' and the part of γ between x and y is a geodesic from ρ to y that coincides with γ in a neighbourhood of y . By Proposition 4.2, the path γ'' must coincide with γ , so γ' must coincide with γ , which proves our claim.

We equip $\mathbf{T}^g(\mathcal{P}^h)$ with its natural tree metric: if $x, y \in \mathbf{T}^g(\mathcal{P}^h)$, the intersection of the geodesics from ρ to x and y is compact so there is a unique point z on it that maximizes $d(\rho, z)$. We write $d_{\mathbf{T}^g(\mathcal{P}^h)}(x, y) = d_{\mathcal{P}^h}(x, z) + d_{\mathcal{P}^h}(z, y)$. Equipped with this distance $\mathbf{T}^g(\mathcal{P}^h)$ is a real tree. However, it is not obvious that it is locally compact (for example it is not if \mathcal{P}^h is replaced by \mathbb{R}^2 equipped with the Euclidean norm).

4.2 Two lemmas about near-critical strips

We first show that two disjoint geodesics in \mathbf{T}_λ^g are quite well-separated, which rules out problem (ii). We denote by γ_ℓ and γ_r the left and right boundaries of the infinite strip $S_{\lambda_n}^1$.

Lemma 4.3. Let $b > a > 0$ and $\varepsilon > 0$. Then there is $\delta > 0$ such that, for n large enough, the following holds:

$$\mathbb{P}\left(\forall i, j \in [an, bn], d_{S_{\lambda_n}^1}(\gamma_\ell(i), \gamma_r(j)) \geq \delta n\right) \geq 1 - \varepsilon.$$

Note that this lemma is very similar to Lemma 3.4 of [31]. However, the slices considered here are not exactly the same, and the lemma of [31] only gives the result with positive probability.

Proof. The idea of the proof is as follows: assume that two points on γ_r and γ_ℓ at height in $[a, b]$ are too close from each other. Then with positive probability S_λ^1 is the only strip of \mathbb{T}_{λ_n} that reaches height b . In this case, we have a small separating cycle in \mathbb{T}_{λ_n} , which becomes a pinch point in the scaling limit. This contradicts the homeomorphicity of \mathcal{P}^h to the plane (this is Proposition 18 of [14], and a consequence of the homeomorphicity of the Brownian map to the sphere [32]).

More precisely, assume the lemma is not true. Then up to extraction, for all $\delta > 0$ and for n large enough, we have

$$\mathbb{P}\left(\exists i, j \in [an, bn], d_{S_{\lambda_n}^1}(\gamma_\ell(i), \gamma_r(j)) < \delta n\right) \geq \varepsilon.$$

Note that if this happens, then by the triangle inequality we must have $|i - j| < \delta n$.

On the other hand, by Theorem 1, the probability that the tree $\mathbf{T}_{\lambda_n}^g$ has only one vertex at height $(b+1)n$ is $\mu_{\lambda_n}(1)^{(b+1)n} = m_{\lambda_n}^{(b+1)n} \rightarrow e^{-2\sqrt{2}(b+1)}$. Since the tree $\mathbf{T}_{\lambda_n}^g$ is independent of the strip $S_{\lambda_n}^1$, for n large enough, the following event occurs with probability at least $\frac{1}{2}e^{-2\sqrt{2}(b+1)}\varepsilon$:

$$\begin{aligned} &\{\text{there are } i, j \in [an, bn] \text{ with } d_{S_{\lambda_n}^1}(\gamma_\ell(i), \gamma_r(j)) < \delta n\} \\ &\quad \text{AND} \\ &\{S_{\lambda_n}^1 \text{ is the only strip of } \mathbb{T}_{\lambda_n} \text{ at height } (b+1)n\}. \end{aligned}$$

If this event occurs, the hull of radius $(b+1)n$ in \mathbb{T}_{λ_n} is the hull of radius $(b+1)n$ in $S_{\lambda_n}^1$, where the boundary geodesics γ_ℓ and γ_r have been glued together. Hence, by concatenating the geodesic in $S_{\lambda_n}^1$ between $\gamma_\ell(i)$ and $\gamma_r(j)$ and a portion of $\gamma_\ell = \gamma_r$ of length $|i - j|$, we get a cycle of length not greater than $2\delta n$, at height at least $(a - \delta)n$, and that separates ρ from infinity. In other words, for all $\delta > 0$, with probability at least $\frac{1}{2}e^{-2\sqrt{2}(b+1)}\varepsilon$, the following event occurs:

{there is a point $z \in \mathbb{T}_{\lambda_n}$ with $a - \delta \leq \frac{1}{n}d(\rho, z) \leq b + \delta$ such that for any continuous path p from ρ to infinity, there is a point y of p such that $\frac{1}{n}d(y, z) \leq 2\delta$ }.

This last property is closed for the Gromov–Hausdorff topology. To prove it properly, we would need to replace the path to infinity by a path to a point x at distance from ρ large enough so that x cannot lie in $B_{(b+1)n}^\bullet(\mathbb{T}_{\lambda_n})$, and then to replace the continuous path by a δ -chain. We omit the details here, see Appendix of [14] for something very similar.

Hence, by the main Theorem of [14], for any $\delta > 0$, with probability at least $\frac{1}{2}e^{-2\sqrt{2}(b+1)\varepsilon}$, there is a point $z \in \mathcal{P}^h$ with $a - \delta \leq d(\rho, z) \leq b + \delta$ such that any continuous path from ρ to infinity contains a point at distance at most δ from z . Since this event is nondecreasing in δ , with probability at least $\frac{1}{2}e^{-2\sqrt{2}(b+1)\varepsilon}$ it occurs for every $\delta > 0$. If it does, let z_n be such a point for $\delta = \frac{1}{n}$, and let z be a subsequential limit of (z_n) . Then we have $a \leq d(\rho, z) \leq b$ and every infinite path from ρ to infinity must contain z . Hence, there is a single point that separates the origin from infinity in \mathcal{P}^h , which is impossible by homeomorphicity of \mathcal{P}^h to the plane. \square

The next lemma shows that for any $x \in \mathbb{T}_{\lambda_n}$, there is a geodesic from x to ρ that coincides with a leftmost infinite geodesic on a quite long distance. Combined with the uniqueness of geodesics between ρ and points of $\mathbf{T}^g(\mathcal{P}^h)$, this will rule out problem (iii). We consider a strip $S_{\lambda_n}^0$ and we denote by γ_ℓ and γ_r its left and right geodesic boundaries.

Definition 4.4. Let $x \in S_{\lambda_n}^0$ and $a > 0$. We say that x is *a-close to the boundary* if there is a geodesic γ from ρ to x that contains either $\gamma_\ell(i)$ for every $0 \leq i \leq a$, or $\gamma_r(i)$ for every $0 \leq i \leq a$.

Lemma 4.5. Let $\varepsilon, r > 0$.

(i) There is $C > 0$ such that for n large enough

$$\mathbb{P}(\text{every } x \in \partial B_{Cn}^\bullet(S_{\lambda_n}^0) \text{ is } (rn)\text{-close to the boundary}) \geq 1 - \varepsilon.$$

(ii) There is $C' > 0$ such that for n large enough

$$\mathbb{P}(\text{every } x \in S_{\lambda_n}^0 \setminus B_{C'n}(S_{\lambda_n}^0) \text{ is } (rn)\text{-close to the boundary}) \geq 1 - 2\varepsilon.$$

Note that we gave two quite similar versions of the lemma. The version (i) is the most natural to prove, whereas (ii) passes more easily to the Gromov–Hausdorff limit and is the one we will use later.

Proof. (i) We recall that $\tau_{\lambda_n}^0$ is the skeleton of $S_{\lambda_n}^0$. We first argue that showing point (i) of the lemma is equivalent to bounding the heights of the trees grafted on the spine of τ^0 . Let $k > 0$ (we will precise its value later). We denote by x_1, \dots, x_{p+1} the vertices of $\partial B_k^\bullet(S_{\lambda_n}^0)$ that lie on the right of the spine of τ^0 , from left to right. For $1 \leq i \leq p$, let also t_i be the subtree of descendants of the edge $\{x_i, x_{i+1}\}$ in τ^0 . For $1 \leq i \leq p$, let also γ_i be the leftmost geodesic from x_i to ρ . It is clear (see Figure 6) that the distance between x_i and the point at which γ_i and γ_r merge is equal to the maximum of the heights of the trees starting between x_i and γ_r . Hence, we have $\gamma_i(j) \in \gamma_r$ as soon as j is greater than the heights of all the trees t_i, t_{i+1}, \dots, t_p . Therefore, if we denote by $H_{\lambda_n}^r(k)$ the height of the forest (t_1, \dots, t_p) , then for any i , there is a geodesic from x_i to ρ that contains $\gamma_r(j)$ for $0 \leq j \leq k - H_{\lambda_n}^r(k)$. We can do the same reasoning for vertices on the left of the spine. We denote by $H_{\lambda_n}^\ell(k)$ the height of the forest that is defined similarly on the left of the spine, and write $H_{\lambda_n}(k) = \max(H_{\lambda_n}^\ell(k), H_{\lambda_n}^r(k))$. It is then enough to find C such that for n large enough:

$$\mathbb{P}(H_{\lambda_n}(Cn) \leq (C - r)n) \geq 1 - \varepsilon. \tag{29}$$

We write $P_{\lambda_n}(Cn) = |\partial B_{Cn}^\bullet(S_{\lambda_n}^0)| - 2$ (this is the number of vertices of τ^0 at height Cn that are not on the spine). Conditionally on $P_{\lambda_n}(Cn)$, all the trees of descendants of the vertices x_i are independent Galton–Watson trees conditioned on extinction before a finite time $\lfloor Cn \rfloor$, so we have

$$\mathbb{P}(H_{\lambda_n}(Cn) \leq (C-r)n | P_{\lambda_n}(Cn) = p) = \left(\frac{g_{\lambda_n}^{\circ \lfloor (C-r)n \rfloor}(0)}{g_{\lambda_n}^{\circ \lfloor Cn \rfloor}(0)} \right)^p \geq g_{\lambda_n}^{\circ \lfloor (C-r)n \rfloor}(0)^p.$$

By Lemma 2.8, we get

$$g_{\lambda_n}^{\circ \lfloor (C-r)n \rfloor}(0) = 1 - \frac{2}{\sinh^2(\sqrt{2}(C-r))} \frac{1}{n^2} + o\left(\frac{1}{n^2}\right),$$

so $g_{\lambda_n}^{\circ \lfloor (C-r)n \rfloor}(0) \geq 1 - \frac{z}{n^2}$ for n large enough, where $z = \frac{3}{\sinh^2(\sqrt{2}(C-r))}$. Hence, we get

$$\mathbb{P}(H_{\lambda_n}(Cn) \leq (C-r)n) \geq \mathbb{E} \left[\left(1 - \frac{z}{n^2}\right)^{P_{\lambda_n}(Cn)} \right].$$

By using the distribution of P_{Cn} given by Lemma 2.16, we obtain

$$\begin{aligned} \mathbb{E} \left[\left(1 - \frac{z}{n^2}\right)^{P_{\lambda_n}(Cn)} \right] &= \frac{m_{\lambda_n}^{-\lfloor Cn \rfloor}}{\Pi_{\lambda_n}(\theta_{\lambda_n}(0))} \\ &\quad \times \left(\Pi_{\lambda_n} \left(\left(1 - \frac{z}{n^2}\right) g_{\lambda_n}^{\circ \lfloor Cn \rfloor + 1}(0) \right) - \Pi_{\lambda_n} \left(\left(1 - \frac{z}{n^2}\right) g_{\lambda_n}^{\circ \lfloor Cn \rfloor}(0) \right) \right) \\ &\geq \frac{1}{\Pi_{\lambda_n}(\theta_{\lambda_n}(0))} \times m_{\lambda_n}^{-\lfloor Cn \rfloor} \times \left(1 - \frac{z}{n^2}\right) \times \left(g_{\lambda_n}^{\circ \lfloor Cn \rfloor + 1}(0) - g_{\lambda_n}^{\circ \lfloor Cn \rfloor}(0) \right) \\ &\quad \times \Pi'_{\lambda_n} \left(\left(1 - \frac{z}{n^2}\right) g_{\lambda_n}^{\circ \lfloor Cn \rfloor}(0) \right). \end{aligned}$$

by convexity of Π_{λ_n} . We can now compute everything using Lemmas 2.8 and 2.10. As $n \rightarrow +\infty$, the first factor goes to $\frac{1}{2}$, the second one goes to $e^{2\sqrt{2}C}$, the third one goes to 1. Moreover, by Lemma 2.8 we have

$$g_{\lambda_n}^{\circ \lfloor Cn \rfloor}(0) = 1 - \frac{2}{\sinh^2(\sqrt{2}C)} \frac{1}{n^2} + O\left(\frac{1}{n^3}\right)$$

and, if $x_n = 1 - \frac{c}{n^2} + o\left(\frac{1}{n^2}\right)$, by (11) we have

$$g_{\lambda_n}(x_n) - x_n \underset{n \rightarrow +\infty}{\sim} \frac{2c\sqrt{c+2}}{n^3},$$

so the fourth factor is equivalent to $4\sqrt{2} \frac{\cosh(\sqrt{2}C)}{\sinh^3(\sqrt{2}C)} \frac{1}{n^3}$. Finally, by taking the derivative of Lemma 2.10, we get

$$\Pi'_{\lambda_n} \left(1 - \frac{c}{n^2} + o\left(\frac{1}{n^2}\right) \right) = \frac{1}{\sqrt{2}(\sqrt{2} + \sqrt{c+2})^2} n^3 + o(n^3).$$

By putting all these estimates together we obtain

$$\mathbb{E} \left[\left(1 - \frac{z}{n^2}\right)^{P_{\lambda_n}(Cn)} \right] \xrightarrow{n \rightarrow +\infty} e^{2\sqrt{2}C} \frac{\cosh(\sqrt{2}C)}{\sinh^3(\sqrt{2}C)} \left(1 + \sqrt{z + \coth^2(\sqrt{2}C)} \right)^{-2}.$$

This goes to 1 as $C \rightarrow +\infty$ (remember that $z = \frac{3}{\sinh^2(\sqrt{2}(C-r))}$ with r fixed), so if C is chosen large enough, then $\mathbb{E} \left[\left(1 - \frac{z}{n^2}\right)^{P_{\lambda_n}(Cn)} \right] \geq 1 - \varepsilon$ for n large enough. This proves (29) and the version (i) of the lemma.

- (ii) This is quite easy using version (i). Let C be given by point (i). Note that if any $x \in \partial B_{Cn}^\bullet(S_{\lambda_n}^0)$ is (rn) -close to the boundary, then so is any $x' \in S_{\lambda_n}^0 \setminus B_{Cn}^\bullet(S_{\lambda_n}^0)$. Indeed, any geodesic γ from x' to ρ must contain a point $x \in \partial B_{Cn}^\bullet(S_{\lambda_n}^0)$, and we can replace the portion of γ between x and ρ by a geodesic that coincides with γ_ℓ or γ_r between height 0 and rn .

Hence, it is enough to find C' such that with probability $1 - \varepsilon$, any point $x' \in S_{\lambda_n}^0$ such that $d(x', \rho) \geq C'n$ is not in $B_{C'n}^\bullet(S_{\lambda_n}^0)$. In other words, we want to prove that the radius of $\frac{1}{n}B_{C'n}^\bullet(S_{\lambda_n}^0)$ from ρ is tight as $n \rightarrow +\infty$. Since $S_{\lambda_n}^0$ can be embedded in \mathbb{T}_{λ_n} in a way that preserves the distances from ρ , this is a consequence of the local Gromov–Hausdorff tightness of $\frac{1}{n}\mathbb{T}_{\lambda_n}$. \square

Note that by Lemma 2.17, Lemma 4.3 holds if we replace $S_{\lambda_n}^1$ by $S_{\lambda_n}^0$ and Lemma 4.5 also holds if we replace $S_{\lambda_n}^0$ by $S_{\lambda_n}^1$. We will use these results for both $S_{\lambda_n}^0$ and $S_{\lambda_n}^1$.

4.3 Identification of the geodesic tree via Gromov–Hausdorff-closed events

The last two lemmas together with Lemma 4.1 and the fact that $\mathbf{T}^g(\mathcal{P}^h)$ is a tree are basically enough to prove Theorem 3. However, to prove it properly, we need to express the distribution of $\mathbf{T}^g(\mathcal{P}^h)$ in terms of closed events for the Gromov–Hausdorff topology, which turns out to be a bit technical.

Let \mathbf{t} be a (finite or infinite) plane tree with a root vertex ρ . If $v \in V(\mathbf{t}) \setminus \{\rho\}$, we write p_v for its parent. Let $(h_v)_{v \in V(\mathbf{t})}$ be a family of nonnegative numbers satisfying $h_\rho = 0$ and $h_v > h_{p_v}$ for every $v \in V(\mathbf{t}) \setminus \{\rho\}$. We write $\mathbf{t}[h]$ for the metric space obtained from \mathbf{t} by giving, for every $v \in V(\mathbf{t}) \setminus \{\rho\}$, a length $h_v - h_{p_v}$ to the edge between p_v and v . We also recall that \mathbf{B} is the infinite tree in which every vertex has two children, except the root which has only one.

We will now define a large family of events, whose probability will characterize the distribution of a random tree of the form $\mathbf{B}[H]$. Let t be a finite binary tree (that is, a tree in which every vertex has 0 or 2 children, except the root which has exactly one). We write $V^*(t)$ for the set of vertices of t that are not leaves and not ρ . Let $r > 0$, and let $(a_v)_{v \in V^*(t)}$, $(b_v)_{v \in V^*(t)}$ be such that $0 < a_v < b_v < r$ for every $v \in V^*(t)$. We write $\mathcal{A}_r^t(a, b)$ for the set of unbounded trees \mathbf{T} (considered as metric spaces) such that $B_r(\mathbf{T})$ is of the form $t[h]$, where $h_\rho = 0$, $h_v = r$ if v is a leaf of t and $a_v \leq h_v \leq b_v$ for every $v \in V^*(t)$.

In order to prove Theorem 3, we will estimate the probability that $\mathbf{T}^g(\mathcal{P}^h)$ belongs to $\mathcal{A}_r^t(a, b)$. Unfortunately, for the reasons listed in of Section 4.1, the events $\{\mathbf{T}^g(X) \in \mathcal{A}_r^t(a, b)\}$ are not closed for the Gromov–Hausdorff distance. To compute $\mathbb{P}(\mathbf{T}^g(\mathcal{P}^h) \in \mathcal{A}_r^t(a, b))$ from our discrete estimates, we need to approximate the event $\{\mathbf{T}^g(X) \in \mathcal{A}_r^t(a, b)\}$ by closed events. Since such approximations are tedious to write down explicitly in the general case, we will focus on the case where $t = t_0$ is the binary tree with two leaves and one vertex of degree 3. Note that $|V^*(t_0)| = 1$, so a and b are just two real numbers.

Let $C > r$ and let $R \geq C + 1$. If $\delta, \varepsilon > 0$, we write $\mathcal{A}_{r,C,R}^{\delta,\varepsilon}(a, b)$ for the set of compact metric spaces (X, d) satisfying the following property.

"There are points x_0, x_1 and x_2 in X and geodesics γ_1 (resp. γ_2) from x_0 to x_1 (resp. to x_2) such that:

- (i) $d(\rho, x_0) = a$,
- (ii) $d(\rho, x_1) = d(\rho, x_0) + d(x_0, x_1) = r$ and $d(\rho, x_2) = d(\rho, x_0) + d(x_0, x_2) = r$,
- (iii) for every $x \in X$ with $d(\rho, x) > C$, the distance $d(\rho, x)$ is equal to $d(\rho, x_1) + d(x_1, x)$ or to $d(\rho, x_2) + d(x_2, x)$ (it may be equal to both),

- (iv) there are two points $y, z \in X$ with $d(\rho, y) \geq R$ and $d(\rho, z) \geq R$ such that $d(\rho, y) = d(\rho, x_1) + d(x_1, y)$ and $d(\rho, z) = d(\rho, x_2) + d(x_2, z)$,
- (v) if $u \in \gamma_1$ and $v \in \gamma_2$ with $d(\rho, u) > b + 2\varepsilon$ and $d(\rho, v) > b + 2\varepsilon$, then $d(u, v) \geq \delta$."

We refer to the Appendix for the proof that this event is closed for the pointed Gromov–Hausdorff topology. More precisely, it is easy to check that $\mathcal{A}_{r,C,R}^{\delta,\varepsilon}(a,b)$ is simply generated by geodesics as in Definition A.3, so by Proposition A.4 it is closed. By the convergence of $\frac{1}{n}\mathbb{T}_{\lambda_n}$ to \mathcal{P}^h , we have

$$\mathbb{P}\left(B_R(\mathcal{P}^h) \in \mathcal{A}_{r,C,R}^{\delta,\varepsilon}(a,b)\right) \geq \limsup_{n \rightarrow +\infty} \mathbb{P}\left(\frac{1}{n}B_{Rn}(\mathbb{T}_{\lambda_n}) \in \mathcal{A}_{r,C,R}^{\delta,\varepsilon}(a,b)\right). \quad (30)$$

We now try to estimate the right-hand side. By Lemma 4.1, we have

$$\lim_{n \rightarrow +\infty} \mathbb{P}\left(\frac{1}{n}\mathbf{T}_{\lambda_n}^g \in \mathcal{A}_r^{t_0}(a,b)\right) = \mathbb{P}\left(\mathbf{Y}_{2\sqrt{2}} \in \mathcal{A}_r^{t_0}(a,b)\right). \quad (31)$$

Note that to deduce (31) from Lemma 4.1, we need to show $\mathbb{P}\left(\mathbf{Y}_{2\sqrt{2}} \in \partial\mathcal{A}_r^{t_0}(a,b)\right) = 0$, where $\partial\mathcal{A}_r^{t_0}(a,b)$ is the boundary of $\mathcal{A}_r^{t_0}(a,b)$ in the space of rooted metric trees, equipped with the local Gromov–Hausdorff distance. This is true because if $T \in \partial\mathcal{A}_r^{t_0}(a,b)$, then T must have a branching point at height exactly a, b or r , which a.s. does not happen.

If the event in the left-hand side of (31) occurs, let x_0 be the unique point of $\mathbf{T}_{\lambda_n}^g$ at height an , and let x_1 (resp. x_2) be the vertex on the left (resp. right) branch of $\mathbf{T}_{\lambda_n}^g$ at height rn . Then the vertices x_0, x_1 and x_2 and the geodesics γ_1 and γ_2 joining x_1 and x_2 to x_0 in $\mathbf{T}_{\lambda_n}^g$ satisfy assumptions (i) and (ii) in the definition of $\mathcal{A}_{r,C,R}^{\delta,\varepsilon}(a,b)$. Since the tree $\mathbf{T}_{\lambda_n}^g$ is infinite, they also satisfy assumption (iv) for any $R > 0$.

We now fix $\varepsilon > 0$ and apply Lemma 4.5 (version (ii)) to the two strips $S_{\lambda_n}^1$ and $S^0[x_0]$ (the strip whose lowest point is x_0). Lemma 4.5 shows that there is $C > 0$ such that, with probability at least $1 - 4\varepsilon$, for any x in one of the two strips $S_{\lambda_n}^1$ and $S^0(x_0)$ such that $d(x, \rho) > Cn$, there is a geodesic from ρ to x that coincides with γ_1 between ρ and x_1 or with γ_2 between ρ and x_2 . We claim that this is also the case if x does not belong to one of these two strips. Indeed, a geodesic from x to ρ must hit the boundary of one of the two strips above x_1 or x_2 (cf. Figure 13). Therefore, the probability that $\frac{1}{n}\mathbf{T}_{\lambda_n}^g \in \mathcal{A}_r^{t_0}(a,b)$ but assumption (iii) is not satisfied for C is at most 4ε .

Finally, if $\frac{1}{n}\mathbf{T}_{\lambda_n}^g \in \mathcal{A}_r^{t_0}(a,b)$ but assumption (v) is not satisfied for some $0 < \delta < \varepsilon$, there are two vertices v_1 on γ_1 and v_2 on γ_2 , at distance from ρ between $(b + 2\varepsilon)n$ and rn , such that $d_{\mathbb{T}_{\lambda_n}}(v_1, v_2) < \delta n$. A geodesic from v_1 to v_2 must cross either $S_{\lambda_n}^1$ or $S^0[x_0]$ and, since $\delta < \varepsilon$, it must stay at distance at least $(b + \varepsilon)n$ from ρ . Therefore, there are two vertices v'_1 on γ_1 and v'_2 on γ_2 , at distance from ρ between $(b + \varepsilon)n$ and $(r + \varepsilon)n$, such that $d_{S_{\lambda_n}^1}(v'_1, v'_2) < \delta n$ or $d_{S^0[x_0]}(v'_1, v'_2) < \delta n$. By applying Lemma 4.3 to $S_{\lambda_n}^1$ between heights $b + \varepsilon$ and $r + \varepsilon$, and to $S^0[x_0]$ between heights ε and $r + \varepsilon$, we can find $\delta > 0$ such that this occurs with probability at most 2ε . Hence, for every $\varepsilon > 0$, there is $\delta > 0$ such that the probability that $\frac{1}{n}\mathbf{T}_{\lambda_n}^g \in \mathcal{A}_r^{t_0}(a,b)$ but assumption (v) is not satisfied is at most 2ε .

Therefore, for all $\varepsilon > 0$, there are $C > r$ and $\delta > 0$ such that, for any $R \geq C + 1$,

$$\liminf_{n \rightarrow +\infty} \mathbb{P}\left(\frac{1}{n}B_{Rn}(\mathbb{T}_{\lambda_n}) \in \mathcal{A}_{r,C,R}^{\delta,\varepsilon}(a,b)\right) \geq \mathbb{P}\left(\mathbf{Y}_{2\sqrt{2}} \in \mathcal{A}_r^{t_0}(a,b)\right) - 6\varepsilon,$$

so by (30),

$$\mathbb{P}\left(B_R(\mathcal{P}^h) \in \mathcal{A}_{r,C,R}^{\delta,\varepsilon}(a,b)\right) \geq \mathbb{P}\left(\mathbf{Y}_{2\sqrt{2}} \in \mathcal{A}_r^{t_0}(a,b)\right) - 6\varepsilon.$$

- assumptions (i) and (ii) in the definition of $\mathcal{A}_{r,C,R}^{\delta,\varepsilon}(a,b)$ are satisfied,
- there is $C > 0$ such that assumption (iii) is satisfied,
- for every $R \geq C + 1$, assumption (iv) is satisfied,
- for every $\varepsilon > 0$, there is $\delta > 0$ such that assumption (v) is satisfied, which means that the geodesics from ρ to x_1 and x_2 are disjoint between heights b (excluded) and r (included).

Since assumption (iv) is satisfied for any R large enough, there are arbitrarily large geodesics started from ρ passing through x_1 . By a compactness argument, there are infinite geodesics started from ρ and passing through x_1 , and the same is true for x_2 . By assumption (iii), the points x_1 and x_2 are the only ones with this property. By assumptions (i) and (v), the branching point between the geodesics from ρ to x_1 and x_2 lies between heights a and b , so $\mathbf{T}^g(\mathcal{P}^h) \in \mathcal{A}_r^{t_0}(a,b)$. \square

The end of the proof of Theorem 3 is now easy. By Lemma 4.6 and (32), we get

$$\mathbb{P}\left(\mathbf{T}^g(\mathcal{P}^h) \in \mathcal{A}_r^{t_0}(a,b)\right) \geq \mathbb{P}\left(\mathbf{Y}_{2\sqrt{2}} \in \mathcal{A}_r^{t_0}(a,b)\right).$$

The general case for t can be treated along the same lines. This shows that the distribution of $\mathbf{T}^g(\mathcal{P}^h)$ dominates that of $\mathbf{Y}_{2\sqrt{2}}$. Since they are both probability measures, they are the same.

A A Gromov–Hausdorff closedness result

The goal of this appendix is to prove Proposition A.4. It shows that a wide class of events related to geodesics are closed for the Gromov–Hausdorff distance. We believe it might be of interest in other settings. We write \mathcal{G} for the space of pointed compact metric spaces, equipped with the Gromov–Hausdorff distance. We will be interested in some events depending on a metric space $(X, d, \rho) \in \mathcal{G}$.

Definition A.1. We say that a subset \mathcal{A} of \mathcal{G} is *simply generated by points* if it has the following form. Let $k \geq 1$, and let $F \subset \mathbb{R}^{(k+2)^2}$ be closed. Then \mathcal{A} is the set of those $(X, d, \rho) \in \mathcal{G}$ for which there are $(x_i)_{0 \leq i \leq k}$ in X with $x_0 = \rho$ such that, for any $x_{k+1} \in X$, the matrix

$$(d(x_i, x_j))_{0 \leq i, j \leq k+1}$$

lies in F .

A simple example of such a subset would be the set of metric spaces that can be covered by k balls of radius r for some fixed $k \geq 1$ and r .

Lemma A.2. Any subset of \mathcal{G} that is simply generated by points is closed.

Proof. Assume that \mathcal{A} is simply generated by points and let k and F be as above. Let (X_n, d_n, ρ_n) converge to a space (X, d, ρ) with $X_n \in \mathcal{A}$ for every n . By Gromov–Hausdorff convergence, we can embed X and all the X_n isometrically in a space (Z, d_Z) such that the Hausdorff distance D_n between X_n and X goes to 0.

For every n , let $x_0^n, \dots, x_k^n \in X_n$ satisfy the condition given by Definition A.1. We take $y_0^n, \dots, y_k^n \in X$ such that $d_Z(x_i^n, y_i^n) \leq 2D_n$. For all $0 \leq i \leq k$, let y_i be a subsequential limit of $(y_i^n)_{n \geq 0}$ in X (which exists by compactness). To complete the proof that $X \in \mathcal{A}$, all we need to show is that $y_0 = \rho$ and that for any $y_{k+1} \in X$, we have

$$(d(y_i, y_j))_{0 \leq i, j \leq k+1} \in F.$$

The first point is easy because the distances $d_Z(\rho, \rho_n)$, $d_Z(\rho_n, y_0^n)$ and $d(y_0^n, y_0)$ all go to 0 along some subsequence. Moreover, let $y_{k+1} \in X$. There is $x_{k+1}^n \in X_n$ such that $d_Z(x_{k+1}^n, y_{k+1}) \leq 2D_n$. For every $0 \leq i, j \leq k+1$, we then have

$$d(y_i, y_j) = \lim_{n \rightarrow +\infty} d(x_i^n, x_j^n)$$

along some subsequence. But we know that for all n we have $(d(x_i^n, x_j^n))_{0 \leq i, j \leq k+1} \in F$, so we can conclude. \square

Definition A.3. We say that a subset \mathcal{A} of \mathcal{G} is *simply generated by geodesics* if it has the following form. Let $k \geq 1$, and let $F \subset \mathbb{R}^{(2k+2)^2}$ be closed. Then \mathcal{A} is the set of those $(X, d, \rho) \in \mathcal{G}$ for which there are $(x_i)_{0 \leq i \leq k}$ in X with $x_0 = \rho$ and geodesics $(\gamma_i)_{1 \leq i \leq k}$ from ρ to x_i , satisfying the following property. For any $(x_{k+i})_{1 \leq i \leq k}$ such that $x_{k+i} \in \gamma_i$ for every i , and for every $x_{2k+1} \in X$, the matrix

$$(d(x_i, x_j))_{0 \leq i, j \leq 2k+1}$$

lies in F .

For example, the events studied in Section 4.3 are simply generated by geodesics. A simpler example of such an event would be "there are k geodesics such that any point lies at distance at most r from one of these geodesics".

Proposition A.4. Any subset of \mathcal{G} that is simply generated by geodesics is closed.

To go from Lemma A.2 to Proposition A.4, we will need the following definition.

Definition A.5. Let $\ell \geq 0$, and let x, y be two points of a metric space (X, d) . An ℓ -geodesic chain from x to y is a finite sequence $(x(i))_{0 \leq i \leq 2^\ell}$ of points of X such that

- (i) $x(0) = x$ and $x(2^\ell) = y$,
- (ii) $d(x(i), x(i+1)) = \frac{1}{2^\ell} d(x, y)$ for any $0 \leq i \leq 2^\ell - 1$.

Proof of Proposition A.4. Let \mathcal{A} be a subset of \mathcal{G} that is simply generated by geodesics. For $\ell > 0$, we write \mathcal{A}^ℓ for the subset of \mathcal{G} we obtain if we replace continuous geodesics by ℓ -geodesic chains in Definition A.3. Then \mathcal{A}^ℓ is simply generated by points (because the conditions in the definition of an ℓ -geodesic chain are closed), so \mathcal{A}^ℓ is closed by Lemma A.2. Hence, to conclude, it is enough to show

$$\mathcal{A} = \bigcap_{\ell \geq 0} \mathcal{A}^\ell.$$

The inclusion from left to right is immediate since any continuous geodesic contains an ℓ -geodesic chain. Now let $(X, d, \rho) \in \bigcap_{\ell \geq 0} \mathcal{A}^\ell$. For every $\ell \geq 0$ and $1 \leq i \leq k$, let (x_i^ℓ) in X and let $(\gamma_i^\ell(j))_{0 \leq j \leq 2^\ell}$ be ℓ -geodesic chains from ρ to x_i^ℓ satisfying the assumptions of definition A.3. Up to extraction, we may assume that for every $1 \leq i \leq k$, the points x_i^ℓ converge to a point $x_i \in X$. Up to further extraction, by a diagonal argument, for every t of the form $\frac{j}{2^m}$ with $0 \leq j \leq 2^m$, the sequence $(\gamma_i^{m+\ell}(2^\ell j))_{\ell \geq 0}$ converges to a point $\gamma_i(t)$. Moreover, for all such t, t' , we have $d(\gamma_i(t), \gamma_i(t')) = |t - t'| d(\rho, x_i)$, so we can extend $(\gamma_i(t))_{0 \leq t \leq 1, t=j/2^m}$ to a continuous geodesic from ρ to x_i . It is then easy to check that the geodesics γ_i satisfy the required hypothesis. \square

References

- [1] C. Abraham. Rescaled bipartite planar maps converge to the Brownian map. *Ann. Inst. H. Poincaré Probab. Statist.*, 52(2):575–595, 05 2016.
- [2] L. Addario-Berry and M. Albenque. The scaling limit of random simple triangulations and random simple quadrangulations. *Ann. Probab.*, 45(5):2767–2825, 09 2017.
- [3] O. Angel. Growth and percolation on the uniform infinite planar triangulation. *Geom. Funct. Anal.*, 13(5):935–974, 2003.
- [4] O. Angel, T. Hutchcroft, A. Nachmias, and G. Ray. Unimodular hyperbolic triangulations: circle packing and random walk. *Inventiones mathematicae*, 206(1):229–268, Oct 2016.
- [5] O. Angel, B. Kolesnik, and G. Miermont. Stability of geodesics in the Brownian map. *Ann. Probab.*, 45(5):3451–3479, 09 2017.
- [6] O. Angel, A. Nachmias, and G. Ray. Random walks on stochastic hyperbolic half planar triangulations. *Random Structures and Algorithms*, 49(2):213–234, 2016.
- [7] O. Angel and G. Ray. Classification of half planar maps. *Ann. Probab.*, 43(3):1315–1349, 2015.
- [8] O. Angel and G. Ray. The half plane UIPT is recurrent. *Probability Theory and Related Fields*, 170(3):657–683, Apr 2018.
- [9] O. Angel and O. Schramm. Uniform infinite planar triangulations. *Comm. Math. Phys.*, 241(2-3):191–213, 2003.
- [10] K. B. Athreya and P. E. Ney. *Branching processes*, volume 196 of *Die Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, 1972.
- [11] J. Beltran and J.-F. Le Gall. Quadrangulations with no pending vertices. *Bernoulli*, 19:1150–1175, 2013.
- [12] I. Benjamini and R. Tessera. First passage percolation on a hyperbolic graph admits bi-infinite geodesics. *Electron. Commun. Probab.*, 22:8 pp., 2017.
- [13] J. Bettinelli, E. Jacob, and G. Miermont. The scaling limit of uniform random plane maps, via the Ambjørn–Budd bijection. *Electronic J. Probab.*, 19, 2014.
- [14] T. Budzinski. The hyperbolic Brownian plane. *Probability Theory and Related Fields*, 171(1):503–541, Jun 2018.
- [15] T. Budzinski. Supercritical causal maps: geodesics and simple random walk. *arxiv:1806.10588*, 2018.
- [16] N. Curien. Planar stochastic hyperbolic triangulations. *Probability Theory and Related Fields*, 165(3):509–540, 2016.
- [17] N. Curien and J.-F. Le Gall. The Brownian plane. *J. Theoret. Probab.*, 27(4):1249–1291, 2014.
- [18] N. Curien and J.-F. Le Gall. First-passage percolation and local modifications of distances in random triangulations. *arXiv:1511.04264*, 2015.
- [19] N. Curien and J.-F. Le Gall. Scaling limits for the peeling process on random maps. *Ann. Inst. H. Poincaré Probab. Statist.*, 53(1):322–357, 02 2017.

- [20] N. Curien, L. Ménard, and G. Miermont. A view from infinity of the uniform infinite planar quadrangulation. *Lat. Am. J. Probab. Math. Stat.*, 10(1):45–88, 2013.
- [21] N. Curien and L. Ménard. The skeleton of the UIPT, seen from infinity. *Annales Henri Lebesgue*, 1, 03 2018.
- [22] W. Esty. The reverse Galton-Watson process. *J. of Applied Prob.*, 12:574–580, 1975.
- [23] M. Gromov. Hyperbolic manifolds, groups and actions. In *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978)*, volume 97 of *Ann. of Math. Stud.*, pages 183–213. Princeton Univ. Press, Princeton, N.J., 1981.
- [24] T. Hutchcroft and Y. Peres. Boundaries of planar graphs: a unified approach. *Electron. J. Probab.*, 22:20 pp., 2017.
- [25] H. Kesten. Aspects of first passage percolation. In *École d’été de probabilités de Saint-Flour, XIV—1984*, volume 1180 of *Lecture Notes in Math.*, pages 125–264. Springer, Berlin, 1986.
- [26] H. Kesten. Subdiffusive behavior of random walk on a random cluster. *Ann. Inst. H. Poincaré Probab. Statist.*, 22(4):425–487, 1986.
- [27] M. Krikun. Local structure of random quadrangulations. *arXiv:0512304*.
- [28] M. Krikun. A uniformly distributed infinite planar triangulation and a related branching process. *Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI)*, 307(Teor. Predst. Din. Sist. Komb. i Algoritm. Metody. 10):141–174, 282–283, 2004.
- [29] M. Krikun. Explicit enumeration of triangulations with multiple boundaries. *Electron. J. Combin.*, 14(1):Research Paper 61, 14 pp. (electronic), 2007.
- [30] J.-F. Le Gall. Geodesics in large planar maps and in the Brownian map. *Acta Math.*, 205:287–360, 2010.
- [31] J.-F. Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.*, 41:2880–2960, 2013.
- [32] J.-F. Le Gall and F. Paulin. Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere. *Geom. Funct. Anal.*, 18(3):893–918, 2008.
- [33] C. Marzouk. Scaling limits of random bipartite planar maps with a prescribed degree sequence. *Random Structures and Algorithms*, 2018.
- [34] G. Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.*, 210(2):319–401, 2013.
- [35] L. Ménard. Volumes in the Uniform Infinite Planar Triangulation: from skeletons to generating functions. *Combinatorics, Probability and Computing*, 27(6):946–973, 2018.
- [36] R. Stephenson. Local convergence of large critical multi-type Galton–Watson trees and applications to random maps. *Journal of Theoretical Probability*, 31(1):159–205, Mar 2018.