

Local limits of uniform triangulations in high genus

Thomas BUDZINSKI* and Baptiste LOUF†

June 25, 2020

Abstract

We prove a conjecture of Benjamini and Curien stating that the local limits of uniform random triangulations whose genus is proportional to the number of faces are the Planar Stochastic Hyperbolic Triangulations (PSHT) defined in [20]. The proof relies on a combinatorial argument and the Goulden–Jackson recurrence relation to obtain tightness, and probabilistic arguments showing the uniqueness of the limit. As a consequence, we obtain asymptotics up to subexponential factors on the number of triangulations when both the size and the genus go to infinity.

As a part of our proof, we also obtain the following result of independent interest: if a random triangulation of the plane T is weakly Markovian in the sense that the probability to observe a finite triangulation t around the root only depends on the perimeter and volume of t , then T is a mixture of PSHT.

Introduction

Counting maps on surfaces. The enumeration of maps or triangulations on surfaces, going back to Tutte [41] in the planar case, has proved to be connected with many different domains of mathematics and theoretical physics. Such links include the "double scaling limits" considered by physicists in the 90s, the Witten conjecture about the geometry of moduli spaces ([43, 29]), the topological recursion (see e.g. [25]), representations of the symmetric group and solutions of several integrable hierarchies such as the KdV, the KP and the 2-Toda hierarchies [35, 37, 27]. In particular, the link with the KP hierarchy has been used by Goulden and Jackson in [27] to obtain double recurrence formulas on the number of triangulations with size n and genus g (see also [16] for similar relations on quadrangulations). However, asymptotics for these numbers are only known when $n \rightarrow +\infty$ for g fixed [8], and not when both $n, g \rightarrow +\infty$.

Geometric properties of random maps. Alongside these enumerative questions, a probabilistic approach has been the object of a lot of study in the last fifteen years: the goal is then to study the geometric properties of a map picked uniformly in a certain class when the genus g and/or the size n become large. In particular, two extreme cases are now pretty well understood.

The first one is the planar case $g = 0$, which is understood both through local and scaling limits. Many natural models of finite random planar maps have been proved to converge locally as their size goes to infinity towards infinite random planar maps such as the UIPT [7] (see also

*University of British Columbia, budzinski@math.ubc.ca

†IRIF, Université Paris Diderot, blouf@irif.fr

[40] for the type-I UIPT), the UIPQ [30, 18] or infinite Boltzmann planar maps [11]. On the other hand, the Brownian map [32, 34] (see also [22] for a noncompact version) is now known to be the scaling limit of a wide class of models of random planar maps, see e.g. [1, 33, 2]. The Brownian map is also linked to the Liouville quantum gravity approach [23]. Some of these continuous models also have analogues in higher genus such as Brownian surfaces [10] or Liouville quantum gravity on complex tori [24], but the behaviour of these models when the genus goes to infinity is still poorly understood.

The other extreme case which is well understood is the case where the genus is unconstrained and the maps simply consist of uniform random gluings of polygons. Here the number of vertices tends to be very small and their degrees go to infinity [26, 19, 15], so no proper local limit exists.

Random maps with genus proportional to the size. However, much less is known about the case of maps of higher genus, and in particular when the genus is proportional to the size. The only known results so far are the identification of the local limit of uniform unicellular maps (i.e. maps with one face) [4], which is a supercritical random tree, and the calculation of their diameter [39]. One of the reasons why it is more difficult to obtain results in high genus is the lack of explicit enumeration results, which play a key role in the planar case. The goal of this work is to identify the local limit of uniform triangulations whose genus is proportional to the size.

Before describing the limiting objects that appear, let us first explain how the local limit is affected by the genus. By the Euler formula, a triangulation with $2n$ faces and genus g has $3n$ edges and $n + 2 - 2g$ vertices, which implies $g \leq \frac{n+1}{2}$. Hence, if $\frac{g}{n} \rightarrow \theta \in [0, \frac{1}{2}]$, then the average degree of the vertices goes to $\frac{6}{1-2\theta}$. In particular, if $0 < \theta < \frac{1}{2}$, this mean degree lies strictly between 6 and $+\infty$. Therefore, it is natural to expect limit objects to be hyperbolic triangulations of the plane¹. This expected relation between high genus and hyperbolic objects also echoes the construction of higher genus surfaces from the hyperbolic plane in complex geometry.

Planar Stochastic Hyperbolic Triangulations. This has motivated the introduction of random hyperbolic triangulations, first in the half-planar case by Angel and Ray [6], and then in the full-plane case by Curien [20]. More precisely, Curien built a one-parameter family $(\mathbb{T}_\lambda)_{0 < \lambda \leq \lambda_c}$ of random triangulations of the plane², where $\lambda_c = \frac{1}{12\sqrt{3}}$, and characterized them as the only random triangulations of the plane exhibiting a natural spatial Markov property. For any finite triangulation t with a hole of perimeter p and v vertices in total, we have

$$\mathbb{P}(t \subset \mathbb{T}_\lambda) = C_p(\lambda)\lambda^v,$$

where $C_p(\lambda)$ are explicit functions of λ and, by $t \subset T$, we mean that T can be obtained by filling the hole of t with an infinite triangulation. Moreover, \mathbb{T}_{λ_c} is the UIPT, whereas for $\lambda < \lambda_c$, the map \mathbb{T}_λ has hyperbolicity properties such as exponential volume growth [38, 20], positive speed of the simple random walk [20, 5] or the existence of a lot of infinite geodesics escaping quickly away from each other [14].

¹As a deterministic example, the d -regular triangulations of the plane for $d > 6$ are hyperbolic.

²To be exact, the triangulations defined in [20] are type-II triangulations, i.e. triangulations with no loop joining a vertex to itself. The type-I (with loops) analogue, which will be the one considered in this work, was defined in [13].

The PSHT as local limits. For any $g \geq 0$ and $n \geq 2g - 1$, we denote by $\mathcal{T}(n, g)$ the set of rooted type-I triangulations of genus g with $2n$ faces. By *rooted*, we mean that the triangulation is equipped with a distinguished oriented edge called the *root*. Let also $T_{n,g}$ be a uniform triangulation of $\mathcal{T}(n, g)$. We also recall that a sequence of rooted triangulations (t_n) converges locally to a triangulation T if for any $r \geq 0$, the ball of radius r around the root in t_n , seen as a map, converges to the ball of radius r in T . We refer to Section 1 for more precise definitions. For any $\lambda \in (0, \lambda_c]$, let $h \in (0, \frac{1}{4}]$ be such that $\lambda = \frac{h}{(1+8h)^{3/2}}$, and let

$$d(\lambda) = \frac{h \log \frac{1+\sqrt{1-4h}}{1-\sqrt{1-4h}}}{(1+8h)\sqrt{1-4h}}. \quad (1)$$

It can be checked that the function $d(\lambda)$ is increasing with $d(\lambda_c) = \frac{1}{6}$ and $\lim_{\lambda \rightarrow 0} d(\lambda) = 0$ (see the end of Section 4.3 for a quick proof). Then our main result is the following.

Theorem 1. Let (g_n) be a sequence such that $\frac{g_n}{n} \rightarrow \theta$ with $\theta \in [0, \frac{1}{2})$. Then we have

$$T_{n,g_n} \xrightarrow[n \rightarrow +\infty]{(d)} \mathbb{T}_\lambda$$

for the local topology, where λ is the unique solution to the equation

$$d(\lambda) = \frac{1-2\theta}{6}. \quad (2)$$

We highlight that we only prove this theorem for type-I triangulations. This result was conjectured by Benjamini and Curien [20] (in the type-II case) without an explicit formula for $d(\lambda)$, and the formula for $d(\lambda)$ was first conjectured in [12, Appendix B]. The reason why the formula (2) appears is that $d(\lambda)$ is the expected inverse of the root degree in \mathbb{T}_λ , while the corresponding quantity in T_{n,g_n} is asymptotically $\frac{1-2\theta}{6}$ by the Euler formula. While it may seem counter-intuitive that high genus objects yield planar maps in the local limit, this has already been proved for other models such as random regular graphs or unicellular maps [4]. Note that the case $\theta = 0$ corresponds to $\lambda = \lambda_c$, which proves that if $g_n = o(n)$, then T_{n,g_n} converges to the UIPT, which also seems to be a new result, even for g_n constant. On the other hand, when $\theta \rightarrow \frac{1}{2}$, we have $\lambda \rightarrow 0$, so all the range $(0, \lambda_c]$ is covered. Since the object \mathbb{T}_0 is not well defined (it corresponds to a "triangulation" where the vertex degrees are infinite), we expect that if $\theta = \frac{1}{2}$, the sequence (T_{n,g_n}) is not tight for the local topology.

Strategy of the proof. The most natural idea to prove Theorem 1 would be to obtain precise asymptotics for the numbers $\tau(n, g) = |\mathcal{T}(n, g)|$ and to adapt the ideas of [7]. In theory, these numbers are entirely characterized by the Goulden–Jackson recurrence equation [27]. However, this seems very difficult without any a priori estimate on the $\tau(n, g)$ and all our efforts to extract asymptotics when $\frac{g}{n} \rightarrow \theta > 0$ from these relations have failed. Therefore, our proof relies on more probabilistic considerations. It is however interesting to note that our probabilistic arguments allow in the end to obtain combinatorial asymptotics (Theorem 3).

The first part of the proof consists of a tightness result: we prove that (T_{n,g_n}) is tight for the local topology as long as $\frac{g_n}{n}$ stays bounded away from $\frac{1}{2}$. A key tool in the proof is the *bounded ratio lemma* (Lemma 4), which states that the ratio $\frac{\tau(n+1, g)}{\tau(n, g)}$ is bounded as long as $\frac{g}{n}$ stays bounded away from $\frac{1}{2}$. This is essentially enough to adapt the argument of Angel and

Schramm [7] for the tightness of $T_{n,0}$. Along the way, we also show that any subsequential limit is a.s. planar and one-ended. The Goulden–Jackson formula also plays an important role in the proof.

The next step is to notice that any subsequential limit T satisfies a weak Markov property: if t is a finite triangulation with a hole of perimeter p and v vertices in total, then $\mathbb{P}(t \subset T)$ only depends on p and v . From here, we deduce that T must be a mixture of PSHT, i.e. a PSHT with a random parameter Λ .

Finally, what is left to prove is that Λ is deterministic, i.e. it does not depend on T_{n,g_n} . By a surgery argument on finite triangulations which we call the *two holes argument*, we first show that if T_{n,g_n} is fixed, then Λ does not depend on the choice of the root. We conclude by using the fact that the average inverse degree of the root in T_{n,g_n} is asymptotically $\frac{1-2\theta}{6}$.

Weakly Markovian triangulations. Since one of the steps of the proof is a result of independent interest, let us highlight it right now. We call a random triangulation of the plane T *weakly Markovian* if for any finite triangulation t with a hole of perimeter p and v vertices in total, the probability $\mathbb{P}(t \subset T)$ only depends on p and v . This is strictly weaker than the spatial Markov property considered in [20] to define the PSHT, since any mixture of PSHT is weakly Markovian. The result we prove is the following.

Theorem 2. Any weakly Markovian triangulation of the plane is a mixture of PSHT.

Combinatorial asymptotics. Finally, while we were unable to obtain directly asymptotics on $\tau(n, g)$ when both n and g go to $+\infty$, Theorem 1 allows us to obtain such estimates up to sub-exponential factors. For any $\theta \in [0, \frac{1}{2})$, we denote by $\lambda(\theta)$ the value of λ given by (2).

Theorem 3. Let (g_n) be a sequence such that $0 \leq g_n \leq \frac{n+1}{2}$ for every n and $\frac{g_n}{n} \rightarrow \theta \in [0, \frac{1}{2}]$. Then we have

$$\tau(n, g_n) = n^{2g_n} \exp(f(\theta)n + o(n))$$

as $n \rightarrow +\infty$, where $f(0) = \log 12\sqrt{3}$, also $f(1/2) = \log \frac{6}{e}$ and

$$f(\theta) = 2\theta \log \frac{12\theta}{e} + \theta \int_2^{1/\theta} \log \frac{1}{\lambda(1/t)} dt \quad (3)$$

for $0 < \theta < \frac{1}{2}$.

To the best of our knowledge, these are the first asymptotic results on the number of triangulations with both large size and high genus. Note that the integral is well defined since $\lambda(\theta)$ is a continuous function and we have $\lambda(\theta) = O(1/2 - \theta)$ when $\theta \rightarrow 1/2$. Moreover, since $\lambda(\theta) \rightarrow \frac{1}{12\sqrt{3}}$ as $\theta \rightarrow 0$, it is easy to see that the function f is continuous at 0 and at $1/2$. The proof mostly relies on the observation that Theorem 1 gives the limit values of the ratio $\frac{\tau(n+1, g)}{\tau(n, g)}$.

Other types of triangulations. A natural question, which we do not answer in this paper, is to ask whether Theorem 1 can be extended to type-II (i.e. with multiple edges but no loop) or type-III (i.e. with neither multiple edges nor loops) triangulations. To adapt our argument in the type-II setting, one would need to be extra careful with the surgery operations of Section 2 and to overcome the absence of a Goulden–Jackson formula. The question seems more complicated for

type-III triangulations, since then the spatial Markov property is partly lost, and it is not even clear how to define the PSHT. Another natural strategy would be to deduce type-II (resp. type-III) results from Theorem 1. The first step would be to prove that $T_{n,g}$ has a large 2-connected (resp. 3-connected) core.

Structure of the paper. The structure of the paper is as follows. In Section 1, we review basic definitions and previous results that will be used throughout the paper. In Section 2, we prove that the triangulations T_{n,g_n} are tight for the local topology, and that any subsequential limit is a.s. planar and one-ended. In Section 3, we prove Theorem 2, which implies that any subsequential limit of T_{n,g_n} is a PSHT with random parameter Λ . In Section 4, we conclude the proof of Theorem 1 by showing that Λ is deterministic and depends only on θ . Finally, Section 5 is devoted to the proof of Theorem 3.

Acknowledgments. The authors thank Guillaume Chapuy and Nicolas Curien for helpful discussions and comments on earlier versions of this manuscript. The authors also thank the two anonymous referees for useful remarks. The first author is supported by ERC Geobrown (740943). The second author is fully supported by ERC-2016-STG 716083 "CombiTop". The authors would also like to thank the Isaac Newton Institute for Mathematical Sciences (EP-SRC grant number EP/R014604/1) for its hospitality during the Random Geometry follow-up workshop when this work was started.

Contents

1 Preliminaries	6
1.1 Definitions	6
1.2 Combinatorics	8
1.3 The PSHT	9
1.4 Local convergence and dual local convergence.	10
2 Tightness, planarity and one-endedness	11
2.1 The bounded ratio lemma	11
2.2 Planarity and one-endedness	13
2.3 Finiteness of the degrees	18
3 Weakly Markovian triangulations	20
4 Ergodicity	23
4.1 The two holes argument	24
4.2 Conclusion	27
4.3 The average root degree in the type-I PSHT via uniform spanning forests	29
5 Combinatorial asymptotics	33

1 Preliminaries

1.1 Definitions

The goal of this paragraph is to state basic definitions on triangulations that will be used throughout the paper.

As in [21], we define a (finite or infinite) *map* M as a way to glue a collection of oriented polygons, called the *faces*, along their edges in a connected way that matches the orientations. Note that this definition is not restricted to maps with finitely many faces. By forgetting the faces of M and looking only at its vertices and edges, we obtain a graph G (if M is infinite, then G may have vertices with infinite degree).

If the number of polygons is finite, then M is always homeomorphic to an orientable topological surface, so we can define the genus of M as the genus of this surface. The maps that we consider will always be *rooted*, i.e. equipped with a distinguished oriented edge called the *root edge*. The face on the right of the root edge is the *root face*, and the vertex at the start of the root edge is the *root vertex*.

A *triangulation* is a rooted map where all the faces have degree 3. We will mostly be interested in *type-I triangulations*, i.e. triangulations that may contain loops and multiple edges. We mention right now that a *type-II triangulation* is a triangulation that may contain multiple edges, but no loops. In graph-theoretic terms, a type- i triangulation is a triangulation with girth (i.e. smallest cycle length) at least i . Unless specified otherwise, by *triangulation*, we will always mean *type-I triangulation*.

For every $n \geq 1$ and $g \geq 0$, we will denote by $\mathcal{T}(n, g)$ the set of triangulations of genus g with $2n$ faces (the number of faces must be even to glue the edges two by two). By the Euler formula, a triangulation of $\mathcal{T}(n, g)$ has $3n$ edges and $n + 2 - 2g$ vertices. In particular, the set $\mathcal{T}(n, g)$ is nonempty if and only if $n \geq 2g - 1$. We will also denote by $\tau(n, g)$ the cardinal of $\mathcal{T}(n, g)$ and by $T_{n, g}$ a uniform random variable on $\mathcal{T}(n, g)$.

We will also need to consider two different notions of triangulations with boundaries, that we call *triangulations with holes* and *triangulations of multi-polygons*. Basically, the first ones will be used to describe a neighbourhood of the root in a triangulation, and the second ones to describe the complementary of this neighbourhood.

For $\ell \geq 1$ and $p_1, p_2, \dots, p_\ell \geq 1$, we call a *triangulation with holes of perimeter p_1, \dots, p_ℓ* a map where all the faces have degree 3 except, for every $1 \leq i \leq \ell$, a face h_i of degree p_i . The faces h_i are called the *holes*. The boundaries of the faces h_i must be simple and edge-disjoint, but may have common vertices (see the bottom part of Figure 1). A triangulation with holes will be rooted at a distinguished oriented edge, which may lie on the boundary of a hole or not. Triangulations with holes will always be finite.

A (possibly infinite) *triangulation of the (p_1, \dots, p_ℓ) -gon* is a map where all the faces have degree 3 except, for every $1 \leq i \leq \ell$, a face f_i of degree p_i . The faces f_i are called the *external faces*, and must be simple and have vertex-disjoint boundaries. Moreover, each of the external faces comes with a distinguished edge on its boundary, such that the external face lies on the right of the distinguished edge.

We denote by $\mathcal{T}_{p_1, p_2, \dots, p_\ell}(n, g)$ the set of triangulations of the $(p_1, p_2, \dots, p_\ell)$ -gon of genus g with $2n - \sum_{i=1}^{\ell} (p_i - 2)$ triangles, and by $\tau_{p_1, p_2, \dots, p_\ell}(n, g)$ its cardinal. The reason why we choose this convention is that by the Euler formula, a triangulation of $\mathcal{T}_{p_1, p_2, \dots, p_\ell}(n, g)$ has $n + 2 - 2g$




Figure 1: A large triangulation T with genus 7 and a smaller triangulation with holes t (in the bottom) such that $t \subset T$. The holes are filled by a triangulation of the (3, 5)-gon, a triangulation of the 2-gon and a triangulation of the 4-gon. The triangles are not drawn on the picture.

vertices in total, just like a triangulation of $\mathcal{T}(n, g)$.

If t is a triangulation with holes and T a (finite or infinite) triangulation, we write $t \subset T$ if T can be obtained from t by gluing one or several triangulations of multi-polygons to the holes of t (see Figure 1). In particular, in the planar case, this definition coincides with the one used e.g. in [7]. If T is an infinite triangulation, we say that it is *one-ended* if for every finite t with $t \subset T$, only one connected component of $T \setminus t$ contains infinitely many triangles. We also say that T is *planar* if every finite t with $t \subset T$ is planar.

We also recall that to a triangulation t , we can naturally associate its *dual map* t^* : it is the map whose vertices are the faces of t and where for each edge e of t , we draw the dual edge e^* joining the two faces incident to e . If t is a triangulation of a multi-polygon, it will be more suitable to work with the convention that the external faces *do not* belong to the dual t^* . Note that triangulations of multi-polygons have simple and disjoint boundaries, so their dual t^* will always be connected.

Finally, we recall the definition of the *graph distance* in a map. For a pair of vertices (v, v') , the distance $d_t(v, v')$ is the length of the shortest path of edges of t between v and v' . We call d_t^* the graph distance in the dual³ map t^* . We also note that there is a natural way to extend d_t^* to the vertices of t . For a pair of distinct vertices (v, v') , we set

$$d_t^*(v, v') = \min(d_t^*(f, f')) + 1,$$

where the minimum is taken over all pairs (f, f') of faces such that f is incident to v and f' is incident to v' .

³In particular, if t is a triangulation of a multi-polygon, then $d_t^*(f, f')$ is the length of the smallest dual path which *avoids* the external faces.

1.2 Combinatorics

The goal of this paragraph is to summarize some previously known or basic combinatorial results about triangulations in higher genus. We start with the Goulden–Jackson recurrence formula.

Theorem 4. [27] Let $f(n, g) = (3n + 2)\tau(n, g)$, with the conventions $f(-1, 0) = \frac{1}{2}$, $f(0, 0) = 2$ and $f(-1, g) = f(0, g) = 0$ for $g \geq 1$. For every $n, g \geq 0$ with $g \leq \frac{n+1}{2}$, we have

$$f(n, g) = \frac{4(3n + 2)}{n + 1} \left(n(3n - 2)f(n - 2, g - 1) + \sum_{\substack{n_1+n_2=n-2 \\ g_1+g_2=g}} f(n_1, g_1)f(n_2, g_2) \right). \quad (4)$$

As explained in the introduction, this formula is in theory enough to compute all the cardinals $\tau(n, g)$, but efforts to extract asymptotics from here when both n and g go to $+\infty$ have failed so far. Our only use of this formula will be in Section 2.2. We will not fully use the Goulden–Jackson formula, but only the two inequalities (12) and (14) that both follow easily from (4).

We also state right now a crude inequality that bounds the number of triangulations of multi-polygons with genus g by the number of triangulations of genus g . This will be useful later.

Lemma 1. For every $n, g \geq 0$ and $p_1, \dots, p_\ell \geq 1$, we have

$$\tau_{p_1, \dots, p_\ell}(n, g) \leq (6n)^{\ell-1} \tau(n, g).$$

Proof. We describe a way to associate with each map t of $\mathcal{T}_{p_1, \dots, p_\ell}(n, g)$ a map \tilde{t} of $\mathcal{T}(n, g)$ with some marked oriented edges. For each external face f_i of t :

- if $p_i \geq 3$, we triangulate f_i by joining all the vertices of ∂f_i to the start of the distinguished edge on ∂f_i , and we mark this edge as e_i ;
- if $p_i = 2$, we simply glue together the two edges of ∂f_i , and mark the edge that we obtain as e_i ;
- if $p_i = 1$, we use the "classical" root transformation shown on Figure 2, and mark the edge obtained by the gluing as e_i .

We obtain a triangulation with the same genus as the initial one, and we root it at e_1 . Note that the above operation does not change the number of vertices, so the triangulation belongs to $\mathcal{T}(n, g)$. It is easy to see that $t \mapsto \tilde{t}$ is injective. Indeed, if we know $p_i \geq 3$ and the edge e_i , then the $p_i - 2$ triangles created by triangulating f_i are the first $p_i - 2$ triangles on the right of e_i that are incident to its starting point. If $p_i \in \{1, 2\}$, the reverse operation is straightforward. Finally, \tilde{t} is a triangulation of $\mathcal{T}(n, g)$ with $\ell - 1$ marked oriented edges (plus its root edge). Since any triangulation of $\mathcal{T}(n, g)$ has $6n$ oriented edges, we are done. □

Remark 2. The bijection of Figure 2 is classical and implies in particular $\tau_1(n, g) = \tau(n, g)$.




Figure 2: Getting rid of a boundary of size 1. The boundary faces are in blue

1.3 The PSHT

In this subsection, we recall the definition and some basic properties of the type-I Planar Stochastic Hyperbolic Triangulations, or PSHT. They were introduced in [20] in the type-II setting (no loops), but we will be more interested in the type-I PSHT defined in [13]. The PSHT $(\mathbb{T}_\lambda)_{0 < \lambda \leq \lambda_c}$ form a one-parameter family of random infinite triangulations of the plane, where $\lambda_c = \frac{1}{12\sqrt{3}}$. Their distribution is characterized as follows. There is a family of constants $(C_p(\lambda))_{p \geq 1}$ such that for every planar triangulation with a hole of perimeter p and v vertices in total, we have

$$\mathbb{P}(t \subset \mathbb{T}_\lambda) = C_p(\lambda) \times \lambda^v. \quad (5)$$

Moreover, let h be the unique solution in $(0, \frac{1}{4}]$ of

$$\lambda = \frac{h}{(1 + 8h)^{3/2}}. \quad (6)$$

Then we have

$$C_p(\lambda) = \frac{1}{\lambda} \left(8 + \frac{1}{h}\right)^{p-1} \sum_{q=0}^{p-1} \binom{2q}{q} h^q, \quad (7)$$

so the distribution of \mathbb{T}_λ is completely explicit.

A very useful consequence of (5) is the spatial Markov property of \mathbb{T}_λ : for any triangulation t with a hole of perimeter p , conditionally on $t \subset \mathbb{T}_\lambda$, the distribution of the complementary $\mathbb{T}_\lambda \setminus t$ only depends on p . Therefore, it is possible to discover it in a Markovian way by a *peeling exploration*.

Since this will be useful later, we recall basic definitions related to peeling explorations. A *peeling algorithm* \mathcal{A} is a mapping that associates with every triangulation with holes an edge on the boundary of one of the holes. Given an infinite triangulation T and a peeling algorithm \mathcal{A} , we can define an increasing sequence $(\mathcal{E}_T^{\mathcal{A}}(k))_{k \geq 0}$ of triangulations with holes such that $\mathcal{E}_T^{\mathcal{A}}(k) \subset T$ for every k in the following way:

- the map $\mathcal{E}_T^{\mathcal{A}}(0)$ is the trivial map consisting of the root edge only,
- for every $k \geq 1$, the triangulation $\mathcal{E}_T^{\mathcal{A}}(k+1)$ is obtained from $\mathcal{E}_T^{\mathcal{A}}(k)$ by adding the triangle incident to $\mathcal{A}(\mathcal{E}_T^{\mathcal{A}}(k))$ outside of $\mathcal{E}_T^{\mathcal{A}}(k)$ and, if this triangle creates a finite hole, all the triangles in this hole.

Such an exploration is called *filled-in*, because all the finite holes are filled at each step. For the PSHT, we denote by $P^\lambda(k)$ and $V^\lambda(k)$ the perimeter and volume of $\mathcal{E}_T^{\mathcal{A}}(k)$, where by *volume* we mean *total number of vertices*. The spatial Markov property ensures that $(P^\lambda(k), V^\lambda(k))_{k \geq 0}$ is

a Markov chain on \mathbb{N}^2 and that its transitions do not depend on the algorithm \mathcal{A} . We also recall the asymptotic behaviour of these two processes. We have

$$\frac{P^\lambda(k)}{k} \xrightarrow[k \rightarrow +\infty]{a.s.} \sqrt{\frac{1-4h}{1+8h}} \quad \text{and} \quad \frac{V^\lambda(k)}{k} \xrightarrow[k \rightarrow +\infty]{a.s.} \frac{1}{\sqrt{(1+8h)(1-4h)}}, \quad (8)$$

where h is given by (6). These estimates are proved in [20] in the type-II setting. For the type-I PSHT, the proofs are the same and use the combinatorial results of [31]. In particular, the asymptotic ratio between $P^\lambda(k)$ and $V^\lambda(k)$ is $1-4h$, which is a decreasing function of λ . This shows that the PSHT for different values of λ are singular with respect to each other, which will be useful later.

1.4 Local convergence and dual local convergence.

The goal of this section is to recall the definition of local convergence in a setting that is not restricted to planar maps. We also define a weaker (at least for triangulations) notion of local convergence that we call "dual local convergence".

As in the planar case, to define the local convergence, we first need to define balls of triangulations. Let t be a finite triangulation. As usual, for every $r \geq 1$, we denote by $B_r(t)$ the map formed by all the faces of t which are incident to at least one vertex at distance at most $r-1$ from the root vertex, along with all their vertices and edges. We denote by $\partial B_r(t)$ the set of edges e such that exactly one side of e is adjacent to a triangle of $B_r(t)$. The other sides of these edges form a finite number of holes, so $B_r(t)$ is a finite triangulation with holes. Note that contrary to the planar case, there is no bijection between the holes and the connected components of $t \setminus B_r(t)$ (cf. the component on the left of Figure 1). We also write $B_0(t)$ for the trivial "map" consisting of only one vertex and zero edge.

For any two finite triangulations t and t' , we write

$$d_{\text{loc}}(t, t') = (1 + \max\{r \geq 0 \mid B_r(t) = B_r(t')\})^{-1}.$$

This is the *local distance* on the set of finite triangulations. As in the planar case, its completion $\overline{\mathcal{T}}$ is a Polish space, which can be viewed as the set of (finite or infinite) triangulations in which all the vertices have finite degrees. However, this space is not compact.

In some parts of this paper, it will be more convenient to work with a weaker notion of convergence which we call the *dual local convergence*. The reason for this is that, since the degrees in the dual of a triangulation are bounded by 3, tightness for this distance will be immediate, which will allow us to work directly on infinite subsequential limits.

More precisely, we recall that d^* is the graph distance on the *dual* of a triangulation. For any finite triangulation t and any $r \geq 0$, we denote by $B_r^*(t)$ the map formed by all the faces at dual distance at most r from the root face, along with all their vertices and edges. Like $B_r(t)$, this is a finite triangulation with holes. For any two finite triangulations t and t' , we write

$$d_{\text{loc}}^*(t, t') = (1 + \max\{r \geq 0 \mid B_r^*(t) = B_r^*(t')\})^{-1}.$$

Note that in any triangulation t , since the dual graph of t is 3-regular, there are at most $3 \times 2^{r-1}$ faces at distance r from the root face. Therefore, for each r , the volume of $B_r^*(t)$ is bounded by a constant depending only on r , so $B_r^*(t)$ can only take finitely many values. It follows

from a simple diagonal extraction argument that the completion for d_{loc}^* of the set of finite triangulations is compact. We write it $\overline{\mathcal{T}}^*$. This set coincides with the set of finite or infinite triangulations, where the degrees of the vertices may be infinite.

Roughly speaking, the main steps of our proof for tightness will be the following. Since $(\overline{\mathcal{T}}^*, d_{\text{loc}}^*)$ is compact, the sequence (T_{n,g_n}) is tight for d_{loc}^* . We will prove that every subsequential limit is planar and one-ended, and finally that its vertices must have finite degrees. We state right now an easy, deterministic lemma that will allow us to conclude at this point.

Lemma 3. Let (t_n) be a sequence of triangulations of $\overline{\mathcal{T}}$. Assume that

$$t_n \xrightarrow[n \rightarrow +\infty]{d_{\text{loc}}^*} t,$$

with $t \in \overline{\mathcal{T}}$. Then $t_n \rightarrow t$ for d_{loc} when $n \rightarrow +\infty$.

Note that the converse is very easy: the dual ball $B_r^*(t)$ is a deterministic function of $B_{r+1}(t)$, so $d_{\text{loc}}^* \leq 2d_{\text{loc}}$, and convergence for d_{loc} implies convergence for d_{loc}^* .

Proof of Lemma 3. Let $r \geq 1$. Since $t \in \overline{\mathcal{T}}$, the ball $B_r(t)$ is finite, so we can find r^* such that $B_r(t) \subset B_{r^*}^*(t)$. By definition of d_{loc}^* , for n large enough, we have $B_{r^*}^*(t_n) = B_{r^*}^*(t)$. Therefore, we have $B_r(t) \subset t_n$, so $B_r(t_n) = B_r(t)$ for n large enough. Since this is true for any $r \geq 1$, we are done. \square

2 Tightness, planarity and one-endedness

2.1 The bounded ratio lemma

The goal of this section is to prove the following result, which will be our main new input in the proof of tightness.

Lemma 4 (Bounded ratio lemma). Let $\varepsilon > 0$. Then there is a constant $C_\varepsilon > 0$ with the following property: for every $n, g \geq 0$ satisfying $\frac{g}{n} \leq \frac{1}{2} - \varepsilon$ and for every $p \geq 1$, we have

$$\frac{\tau_p(n, g)}{\tau_p(n-1, g)} \leq C_\varepsilon.$$

In particular, by the usual bijection between $\mathcal{T}_1(n, g)$ and $\mathcal{T}(n, g)$, we have

$$\frac{\tau(n, g)}{\tau(n-1, g)} \leq C_\varepsilon.$$

For our future use, it will be important that the constant C_ε does not depend on p . The idea of the proof of Lemma 4 will be to find an "almost-injective" way to obtain a triangulation of $\mathcal{T}_p(n-1, g)$ from a triangulation of $\mathcal{T}_p(n, g)$. This will be done by merging two vertices together. For this, it will be useful to find two vertices that are quite close to each other and have a reasonable degree. This is the point of the next result. We recall that for two vertices v, v' , the distance $d^*(v, v')$ is the length of the smallest dual path from v to v' that avoids the external faces. Let us fix $\varepsilon > 0$. We will call a pair (v_1, v_2) of vertices *good* if $\deg(v_1) + \deg(v_2) \leq \frac{12}{\varepsilon}$ and $d^*(v_1, v_2) \leq \frac{24}{\varepsilon}$.

Lemma 5. In any triangulation of a polygon $t \in \mathcal{T}_p(n, g)$ with $\frac{g}{n} \leq \frac{1}{2} - \varepsilon$, there are at least $\frac{\varepsilon}{12}n$ good pairs of vertices.

Proof. Fix a triangulation $t \in \mathcal{T}_p(n, g)$ with $\frac{g}{n} \leq \frac{1}{2} - \varepsilon$. We first note that a positive proportion of the vertices have a small degree. Indeed, by the Euler formula, t has $3n + 3 - p$ edges and $n + 2 - 2g$ vertices, so the average degree of a vertex is

$$\frac{2(3n + 3 - p)}{n + 2 - 2g} \leq \frac{6n}{2\varepsilon n} = \frac{3}{\varepsilon}.$$

Therefore, at least half of the vertices have degree at most $\frac{6}{\varepsilon}$. There are $n + 2 - 2g \geq 2\varepsilon n$ vertices in t , so at least εn of them have degree not greater than $\frac{6}{\varepsilon}$.

Let $v^1, \dots, v^{\varepsilon n}$ be such vertices and, for every $1 \leq i \leq \varepsilon n$, let f^i be a face incident to v^i . Since each face is incident to only 3 vertices, the set $F = \{f^1, \dots, f^{\varepsilon n}\}$ contains at least $\frac{\varepsilon}{3}n$ faces. It is sufficient to find $\frac{\varepsilon}{12}n$ pairs $(f, f') \in F^2$ with $d^*(f, f') \leq \frac{24}{\varepsilon}$. This will follow from the fact that balls (for d^*) centered at the elements of F must strongly overlap. More precisely, let $r = \frac{12}{\varepsilon}$. Since the dual map t^* is connected, for any $f \in F$, we have⁴ $|B_r^*(f)| \geq r$. Therefore, we have

$$\sum_{f \in F} |B_r^*(f)| \geq \frac{12}{\varepsilon} \times \frac{\varepsilon}{3}n = 4n,$$

whereas $|t^*| = 2n - p + 2 < 4n$. Hence, there must be an intersection between the balls $B_r^*(f)$, so there are $f_1, f'_1 \in F$ such that $d^*(f_1, f'_1) \leq 2r = \frac{24}{\varepsilon}$, and the pair (f_1, f'_1) is good. We set $F_1 = F \setminus \{f_1\}$. We now try to find a good pair in F_1 and remove an element of this pair, and so on. Assume that F_i is the set F where i elements have been removed. If $i < \frac{\varepsilon}{12}n$, then we have

$$\sum_{f \in F_i} |B_r^*(f)| \geq \frac{12}{\varepsilon} \left(\frac{\varepsilon}{3}n - i \right) \geq 3n > |t^*|,$$

so F_i contains a good pair. Therefore, the process will not stop before $i = \frac{\varepsilon}{12}n$, so we can find $\frac{\varepsilon}{12}n$ good pairs in F^2 , which concludes the proof. \square

We are now able to prove the bounded ratio lemma.

Proof of Lemma 4. Let $g \geq 0, n \geq 2$ be such that $\frac{g}{n} \leq \frac{1}{2} - \varepsilon$. We will define an "almost-injection" Φ from $\mathcal{T}_p(n, g)$ to $\mathcal{T}_p(n - 1, g)$. The input will be a triangulation $t \in \mathcal{T}_p(n, g)$ with a marked good pair (v_1, v_2) . By Lemma 5, the number of inputs is at least

$$\frac{\varepsilon}{12}n\tau_p(n, g). \tag{9}$$

Given an input (t, v_1, v_2) , let (f_1, f_2, \dots, f_j) be the shortest path in t^* from v_1 to v_2 . Since the pair (v_1, v_2) is good, we have $j \leq \frac{24}{\varepsilon}$. For all i , let e_i be the edge separating f_i from f_{i+1} . We flip e_1 , then e_2 and so on up to e_{j-1} (see Figure 3). Note that these flips are always well defined, since the faces f_i are pairwise distinct.

As we do so, we keep track of all the flipped edges and the order in which they come. All the edges that were flipped are now incident to v_1 , and the last of them is also incident to v_2 . We then contract this edge and merge v_1 and v_2 into a vertex v , which creates two digons incident to v , that we contract into two edges. Finally, we mark the vertex v obtained by merging v_1 with v_2 , and we also mark the two edges obtained by contracting the digons.

⁴Unless the number of faces of t is smaller than r , in which case $B_r^*(f) = t^*$, so any pair of F^2 is good.




Figure 3: The injective mapping. On the left, a good pair and a path of triangles. In the center, the triangulation after the flips (we flipped the edges 1, 2, 3, 4, 5 in this order). On the right, the final map, after contraction of the brown edge. The stars indicate the contracted digons.

These operations do not change the genus and the boundary length, and remove exactly 1 vertex. Hence, the output of Φ is a triangulation of $\mathcal{T}_p(n-1, g)$, with a marked vertex v of degree at most $\frac{36}{\varepsilon}$ (since $\deg(v) < \deg(v_1) + \deg(v_2) + j$), two marked edges incident to v and an ordered list of edges incident to v . Moreover, given the triangulation $\Phi(t)$, there are at most $n+2-2g \leq n+2$ possible values of v . Since $\deg(v) \leq \frac{36}{\varepsilon}$, once v is fixed, there are at most $\left(\frac{36}{\varepsilon}\right)^2$ ways to choose the two marked edges and $\left(\frac{36}{\varepsilon}\right)^{24/\varepsilon}$ ways to choose the ordered list of edges. Hence, the number of possible outputs of Φ is at most

$$n \left(\frac{36}{\varepsilon}\right)^{24/\varepsilon+2} \tau_p(n-1, g). \quad (10)$$

Finally, it is easy to see that Φ is injective: to go backwards, one just needs to duplicate the two marked edges, split v in two between the two digons, and flip back the edges in the prescribed order. Therefore, by (9) and (10), we obtain

$$\frac{\varepsilon}{12} n \tau_p(n, g) \leq n \left(\frac{36}{\varepsilon}\right)^{24/\varepsilon+2} \tau_p(n-1, g),$$

which concludes the proof with $C_\varepsilon = \left(\frac{36}{\varepsilon}\right)^{24/\varepsilon+3}$. \square

2.2 Planarity and one-endedness

We now fix a sequence (g_n) with $\frac{g_n}{n} \rightarrow \theta \in [0, \frac{1}{2})$. As explained in Section 1.4, the tightness of (T_{n, g_n}) for d_{loc}^* is immediate. Throughout this section, we will denote by T a subsequential limit in distribution. It must be an infinite triangulation. We will first prove that T is planar and one-ended, and then that its vertices have finite degrees. To establish planarity, the idea will be to bound, for any non-planar finite triangulation t , the probability that $t \subset T_{n, g_n}$ for n large. For this, we will need the following combinatorial estimate.

Lemma 6. Fix $k \geq 1$ and $m \in \mathbb{Z}$, numbers $\ell_1, \dots, \ell_k \geq 1$ and perimeters $p_i^j \geq 1$ for $1 \leq j \leq k$ and $1 \leq i \leq \ell_j$. Then

$$\sum_{\substack{n_1 + \dots + n_k = n + m \\ h_1 + \dots + h_k = g_n - 1 - \sum_j (\ell_j - 1)}} \prod_{j=1}^k \tau_{p_1^j, \dots, p_{\ell_j}^j}(n_j, h_j) = o(\tau(n, g_n)) \quad (11)$$

when $n \rightarrow +\infty$.

Proof. By Lemma 1, the left-hand side of (11) can be bounded by

$$C \sum_{\substack{n_1+\dots+n_k=n+m \\ h_1+\dots+h_k=g_n-1-\sum_j(\ell_j-1)}} \prod_{j=1}^k n_j^{\ell_j-1} \tau(n_j, h_j)$$

for $C = 6^{\sum(\ell_j-1)}$. We now use the crude bound $\tau(n_j, h_j) \leq f(n_j, h_j)$, where the numbers $f(n, g) = (3n+2)\tau(n, g)$ are those that appear in the Goulden–Jackson formula (4). Since $\ell_j \geq 1$, we can bound $n_j^{\ell_j-1}$ by $2n^{\ell_j-1}$ for $n \geq m$. The left-hand side of (11) is then bounded by

$$Cn^{\sum_{j=1}^k(\ell_j-1)} \sum_{\substack{n_1+\dots+n_k=n+m \\ h_1+\dots+h_k=g_n-1-\sum_j(\ell_j-1)}} \prod_{j=1}^k f(n_j, h_j).$$

Moreover, the Goulden–Jackson formula implies that

$$\sum_{\substack{n_1+n_2=n \\ h_1+h_2=g}} f(n_1, h_1)f(n_2, h_2) \leq f(n+2, g) \quad (12)$$

for any $n, g \geq 0$. By an easy induction on k , we obtain

$$\sum_{\substack{n_1+\dots+n_k=n \\ h_1+\dots+h_k=g}} \prod_{j=1}^k f(n_j, h_j) \leq f(n+2k-2, g).$$

Therefore, we can bound the left-hand side of (11) by

$$Cn^{\sum_{j=1}^k(\ell_j-1)} f\left(n+m+2k-2, g_n-1-\sum_{j=1}^k(\ell_j-1)\right). \quad (13)$$

The Goulden–Jackson formula implies that

$$f(n-2, g-1) \leq n^{-2}f(n, g) \quad (14)$$

for any n and g , so $f(n, g-i) \leq n^{-2i}f(n+2i, g)$ for any n and $1 \leq i \leq g$. Therefore, from (13), we obtain the asymptotic bound

$$Cn^{\sum_{j=1}^k(\ell_j-1)} n^{-2-2\sum_{j=1}^k(\ell_j-1)} f\left(n+m+2k+2\sum(\ell_j-1), g_n\right) \leq C'n^{-2-\sum_j(\ell_j-1)} f(n, g_n),$$

where in the end we use Lemma 4 (which results in a change in the constant). In particular, the left-hand side of (11) is $o\left(\frac{f(n, g_n)}{n}\right)$, so it is $o(\tau(n, g_n))$. \square

Corollary 7. Every subsequential limit of (T_{n, g_n}) for d_{loc}^* is a.s. planar.

Proof. If a subsequential limit T is not planar, then we can find a finite triangulation t with holes and with genus 1 such that $t \subset T$. Indeed, if we explore T triangle by triangle, the genus may only increase by at most 1 at each step, so if the genus is positive at some point during the exploration, it must be 1 at some point. Therefore, it is enough to prove that for any such triangulation t , we have

$$\mathbb{P}(t \subset T_{n, g_n}) \xrightarrow{n \rightarrow +\infty} 0.$$

If $t \subset T_{n,g_n}$, let T^1, \dots, T^k be the connected components of $T_{n,g_n} \setminus t$. These components define a partition of the set of holes of t , where a hole h is in the j -th class if T^j is the connected component glued to h (for example, on Figure 1, the three classes have sizes 2, 1 and 1). Note that the number of possible partitions is finite and depends only on t (and not on n). Therefore, it is enough to prove that for any partition π of the set of holes of t , we have

$$\mathbb{P}(t \subset T_{n,g_n} \text{ and the partition defined by } T_{n,g_n} \text{ is } \pi) \xrightarrow{n \rightarrow +\infty} 0. \quad (15)$$

If this occurs, for each j , let ℓ_j be the number of holes of T glued to T^j and let $p_1^j, \dots, p_{\ell_j}^j$ be the perimeters of these holes. Then the connected component T^j is a triangulation of the $(p_1^j, \dots, p_{\ell_j}^j)$ -gon (see Figure 1). Moreover, if T_j has genus h_j , then the total genus of T_{n,g_n} is equal to

$$1 + \sum_{j=1}^k h_j + \sum_{j=1}^k (\ell_j - 1),$$

so this sum must be equal to g_n , so

$$\sum_{j=1}^k h_j = g_n - 1 - \sum_{j=1}^k (\ell_j - 1).$$

Moreover, let n_j be such that T^j belongs to $\mathcal{T}_{p_1^j, \dots, p_{\ell_j}^j}(n_j, h_j)$. An easy computation shows that

$$\sum_{j=1}^k n_j = n + m$$

with

$$m = \frac{1}{2} \left(-|F(t)| + \sum_{j=1}^k \sum_{i=1}^{\ell_j} (p_i^j - 2) \right) \in \mathbb{Z},$$

where $|F(t)|$ is the number of triangles of t .

Therefore, the number of triangulations $T \in \mathcal{T}(n, g_n)$ such that $t \subset T$ and the resulting partition of the holes is equal to π is the number of ways to choose, for each j , a triangulation of the $(p_1^j, \dots, p_{\ell_j}^j)$ -gon, such that the total genus of these triangulations is $g_n - 1 - \sum_{j=1}^k (\ell_j - 1)$, and their total size is $n + m$. This is equal to the left-hand side of Lemma 6, so (15) is a consequence of Lemma 6, which concludes the proof. \square

The proof of one-endedness will be similar, but the combinatorial estimate that is needed is slightly different.

Lemma 8. • Fix $k \geq 1$, $m \in \mathbb{Z}$, numbers $\ell_1, \dots, \ell_k \geq 1$ that are not all equal to 1, and perimeters $p_i^j \geq 1$ for $1 \leq j \leq k$ and $1 \leq i \leq \ell_j$. Then

$$\sum_{\substack{n_1 + \dots + n_k = n + m \\ h_1 + \dots + h_k = g_n - \sum_j (\ell_j - 1)}} \prod_{j=1}^k \tau_{p_1^j, \dots, p_{\ell_j}^j}(n_j, h_j) = o(\tau(n, g_n)). \quad (16)$$

- Fix $k \geq 2$, $m \in \mathbb{Z}$ and perimeters p_1, \dots, p_k . There is a constant C such that, for every a and n , we have

$$\sum_{\substack{n_1 + \dots + n_k = n+m \\ h_1 + \dots + h_k = g_n \\ n_1, n_2 > a}} \prod_{j=1}^k \tau_{p_j}(n_j, h_j) \leq \frac{C}{a} \tau(n, g_n). \quad (17)$$

Proof. We start with the first point. The proof is very similar to the proof of Lemma 6, with the following difference: here, the sum of the genera differs by one, so we will lose a factor n^2 in the end of the computation. This forces us to be more careful in the beginning and to use the assumption that the ℓ_j are not all equal to 1.

More precisely, by using Lemma 1 and the bound $\tau(n, g) \leq \frac{1}{n} f(n, g)$, as in the proof of Lemma 6, the left-hand side of (16) can be bounded by

$$C \sum_{\substack{n_1 + \dots + n_k = n+m \\ h_1 + \dots + h_k = g_n - \sum_j (\ell_j - 1)}} \left(\prod_{j=1}^k n_j^{\ell_j - 2} \right) \left(\prod_{j=1}^k f(n_j, h_j) \right),$$

where C does not depend on n . Without loss of generality, assume that $\ell_1 \geq 2$. Then we have $n_1^{\ell_1 - 2} \leq (n+m)^{\ell_1 - 2}$. Moreover, for every $j \geq 2$, we have $n_j^{\ell_j - 2} \leq n_j^{\ell_j - 1} \leq (n+m)^{\ell_j - 1}$ since $\ell_j \geq 1$. Therefore, we obtain for $n \geq m$:

$$\prod_{j=1}^k n_j^{\ell_j - 2} \leq 2n^{\sum_j (\ell_j - 1) - 1}.$$

By using this and the Goulden–Jackson formula in the same way as in the proof of Lemma 6, we obtain the bound

$$C n^{-1 - \sum_j (\ell_j - 1)} f(n+m+2k-2, g_n) \leq C' n^{-1 - \sum_j (\ell_j - 1)} f(n, g_n),$$

where the last inequality follows from the bounded ratio lemma. Since $\ell_1 \geq 2$, we have $\sum_j (\ell_j - 1) \geq 1$, so this is $o\left(\frac{f(n, g_n)}{n}\right)$ and we get the result.

We now prove the second point. As in the first case (but with $\ell_j = 1$ for every j), the left-hand side can be bounded by

$$C \sum_{\substack{n_1 + \dots + n_k = n+m \\ h_1 + \dots + h_k = g_n \\ n_1, n_2 > a}} \left(\prod_{j=1}^k \frac{1}{n_j} \right) \left(\prod_{j=1}^k f(n_j, h_j) \right).$$

Moreover, if $n_1, n_2 > a$, then at least one of the n_j is larger than $\frac{n+m}{k}$ and two are larger than a , so $\prod_{j=1}^k n_j \geq \frac{(n+m)a}{k}$, so we obtain the bound (for n large enough, with C' and C'')

independent of n and a)

$$\begin{aligned}
\frac{C'}{an} \sum_{\substack{n_1+\dots+n_k=n+m \\ h_1+\dots+h_k=g_n \\ n_1, n_2 > a}} \prod_{j=1}^k f(n_j, h_j) &\leq \frac{C'}{an} \sum_{\substack{n_1+\dots+n_k=n+m \\ h_1+\dots+h_k=g_n}} \prod_{j=1}^k f(n_j, h_j) \\
&\leq \frac{C'}{an} f(n+m+2k-2, g_n) \\
&\leq \frac{C''}{an} f(n, g_n) \\
&\leq \frac{C''}{a} \tau(n, g_n),
\end{aligned}$$

where we use the Goulden–Jackson formula to reduce the sum and finally the bounded ratio lemma, in the same way as previously. \square

Corollary 9. Every subsequential limit T of (T_{n, g_n}) for d_{loc}^* is a.s. one-ended in the sense that, for every finite triangulation t with holes such that $t \subset T$, only one hole of t is filled with infinitely many faces⁵.

Proof. The proof is quite similar to the proof of Corollary 7, but with Lemma 8 playing the role of Lemma 6.

More precisely, if a subsequential limit T is not one-ended with positive probability, it contains a finite triangulation t such that two of the connected components of $T \setminus t$ are infinite. This means that we can find $\varepsilon > 0$, a triangulation t and two holes h_1, h_2 of t such that, for every $a > 0$,

$$P(t \subset T \text{ and } T \setminus t \text{ has two connected components with at least } a \text{ faces}) \geq \varepsilon. \quad (18)$$

By Corollary 7, we can assume that t is planar. If this holds, then T contains a finite triangulation obtained by starting from t and adding a faces in the hole h_1 and a faces in the hole h_2 . We denote by $t^{a,a}$ the set of such triangulations. Then (18) means that for any $a > 0$, for n large enough, we have

$$\mathbb{P}(T_{n, g_n} \text{ contains a triangulation of } t^{a,a}) \geq \varepsilon. \quad (19)$$

This can occur in two different ways, which will correspond to the two items of Lemma 8:

- (i) either at least one connected component of $T_{n, g_n} \setminus t$ is glued to at least two holes of t ,
- (ii) or the k holes of t correspond to k connected components T^1, \dots, T^k , where T^1 and T^2 have size at least a .

In case (i), the connected components of T_{n, g_n} are triangulations of multi-polygons, at least one of which has two boundaries. The proof that the probability of this case goes to 0 is now the same as the proof of Corollary 7, but we use the first point of Lemma 8. Note that the assumption the ℓ_j are not all 1 comes from the fact that one of the connected components is glued to two holes. Moreover, the sum of the genres of the T^j is $g - \sum_j (\ell_j - 1)$ and not $g - 1 - \sum_j (\ell_j - 1)$ because this time t has genus 0 and not 1.

⁵This is a "weak" definition of one-endedness, since it does not prevent T to be the dual of a tree. However, once we will have proved that T has finite vertex degrees, this will be equivalent to the usual definition.

Similarly, in case (ii), the k holes of t must be filled with k triangulations of a single polygon, two of which have at least a faces, so they belong to a set of the form $\mathcal{T}_{p_j}(n_j, h_j)$ with $n_j \geq \frac{a}{2}$ if a is large enough compared to the perimeters of the holes. Hence, the second point of Lemma 8 allows to bound the number of ways to fill these holes. We obtain that, for a large enough, we have

$$\mathbb{P}(T_{n, g_n} \text{ contains a triangulation of } t^{a, a}) \leq o(1) + \frac{2C}{a}$$

as $n \rightarrow +\infty$, where $o(1)$ comes from case (i) and $\frac{2C}{a}$ from case (ii). This contradicts (19), so T is a.s. one-ended. \square

2.3 Finiteness of the degrees

Our goal is now to prove tightness for d_{loc} . As before, let (g_n) be a sequence with $\frac{g_n}{n} \rightarrow \theta \in [0, \frac{1}{2}]$.

Proposition 10. The sequence (T_{n, g_n}) is tight for d_{loc} .

Let T be a subsequential limit of (T_{n, g_n}) for d_{loc}^* . By Lemma 3, to finish the proof of tightness for d_{loc} , we only need to show that almost surely, all the vertices of T have finite degrees. As in [7], we will first study the degree of the root vertex, and then extend finiteness by using invariance under the simple random walk. The main difference with [7] is that, while [7] uses exact enumeration results, we will rely on the bounded ratio lemma.

Lemma 11. The root vertex of T has a.s. finite degree.

Proof. We follow the approach of [7] and perform a filled-in peeling exploration of T . Before expliciting the peeling algorithm that we use, note that we already know by Corollary 7 that the explored part will always be planar, so no peeling step will merge two different existing holes. Moreover, by Corollary 9, if a peeling step separates the boundary into two holes, then one of them has finitely many faces inside, so it will be filled with a finite triangulation. Therefore, at each step, the explored part will be a triangulation with a single hole.

The peeling algorithm \mathcal{A} that we use is the following: if the root vertex ρ belongs to ∂t , then $\mathcal{A}(t)$ is the edge on ∂t on the left of ρ . If $\rho \notin \partial t$, then the exploration is stopped. Since only finitely many edges incident to ρ are added at each step, it is enough to prove that the exploration will a.s. eventually stop. We recall that $\mathcal{E}_T^A(i)$ is the explored part at time i .

We will prove that at each step, conditionally on $\mathcal{E}_T^A(i)$, the probability to swallow the root and finish the exploration at time $i+1$ or $i+2$ is bounded from below by a positive constant. For every triangulation t with one hole such that $\rho \in \partial t$, we denote by t^+ the triangulation constructed from t as follows (see Figure 4):

- we first glue a triangle to the edge of ∂t on the left of ρ , in such a way that the third vertex of this triangle does not belong to t , to obtain a triangulation with perimeter at least 2;
- we then glue a second triangle to the two edges of the boundary incident to ρ .

Note that t^+ is a planar map with the same perimeter as t but one more vertex and two more triangles. By the choice of our peeling algorithm, if $\mathcal{E}_T^A(i)^+ \subset T$, then we have $\mathcal{E}_T^A(i+2) = \mathcal{E}_T^A(i)^+$. Moreover, if this is the case, the exploration is stopped at time $i+2$. Hence, it is enough to prove that the quantity

$$\mathbb{P}(t^+ \subset T | t \subset T)$$




Figure 4: The construction of t^+ from t . In gray, the triangulation t . In red, the root vertex. In blue, the new triangles. In the bottom, the case $|\partial t| = 1$.

is bounded from below for finite, planar triangulations t with a single hole and $\rho \in \partial t$.

We fix such a t , with perimeter p . Let also (n_k) be a sequence of indices such that $T_{n_k, g_{n_k}}$ converges in distribution to T . We have

$$\mathbb{P}(t^+ \subset T | t \subset T) = \lim_{k \rightarrow +\infty} \frac{\mathbb{P}(t^+ \in T_{n_k, g_{n_k}})}{\mathbb{P}(t \in T_{n_k, g_{n_k}})} = \lim_{k \rightarrow +\infty} \frac{\tau_p(n_k + m - 1, g_{n_k})}{\tau_p(n_k + m, g_{n_k})},$$

where $m = \frac{p-2-|F(t)|}{2}$ and $|F(t)|$ is the number of triangles of t . Moreover, there is $\varepsilon > 0$ such that $g_n \leq (\frac{1}{2} - 2\varepsilon)n$ for n large enough, so $\frac{g_{n_k}}{n_k + m} \leq \frac{1}{2} - \varepsilon$ for k large enough. By the bounded ratio lemma, we obtain

$$\mathbb{P}(t^+ \subset T | t \subset T) \geq \frac{1}{C_\varepsilon}$$

for every t , which concludes the proof. \square

Remark 12. Our proof shows that the number of steps needed to swallow the root has exponential tail. However, since we do not control the finite triangulations filling the holes that may appear, it does not give any quantitative bound on the root degree.

Proof of Proposition 10. Let T be a subsequential limit of (T_{n, g_n}) for d_{loc}^* . By Lemma 3, to obtain tightness, it is enough to prove that almost surely, all the vertices of T have finite degrees. The argument is essentially the same as in [7] and relies on Lemma 11 and invariance under the simple random walk.

More precisely, for every n , let \vec{e}_0^n be the root edge of T_{n, g_n} and let \vec{e}_0 be the root of T . We first note that the distribution of T_{n, g_n} is invariant under reversing the orientation of the root, so this is also the case of T . By Lemma 11, this implies that the endpoint of \vec{e}_0 has a.s. finite degree.

We then denote by \vec{e}_1^n the first step of the simple random walk on T_{n, g_n} : its starting point is the endpoint of \vec{e}_0^n and its endpoint is picked uniformly among all the neighbours of the starting point. Since the endpoint of \vec{e}_0 has finite degree, we can also define the first step \vec{e}_1 of the simple random walk on T . For the same reason as in the planar case (see Theorem 3.2 of [7]), the triangulations $(T_{n, g_n}, \vec{e}_1^n)$ and $(T_{n, g_n}, \vec{e}_0^n)$ have the same distribution, so (T, \vec{e}_1) has the same distribution as (T, \vec{e}_0) . In particular, all the neighbours of the endpoint of \vec{e}_0 must have

finite degrees. From here, an easy induction on i shows that for every $i \geq 0$, we can define the i -th step $\vec{\mathcal{E}}_i$ of the simple random walk on T , that $(T, \vec{\mathcal{E}}_i)$ has the same distribution as $(T, \vec{\mathcal{E}}_0)$ and that all vertices at distance i from the root in T are finite. This proves the tightness for d_{loc} .

Planarity is proved by Corollary 7. Finally, it is easy to check that for triangulations with finite vertex degrees, the weak version of one-endedness proved in Corollary 9 implies the usual one. For example, if V is a finite set of vertices of T , one can consider a finite, connected triangulation $t \subset T$ containing all the faces and edges incident to vertices of V . \square

3 Weakly Markovian triangulations

The goal of this section is to prove Theorem 2. We first recall the definition of a weakly Markovian triangulation.

Definition 13. Let T be a random infinite triangulation of the plane. We say that T is *weakly Markovian* if there is a family $(a_v^p)_{v \geq p \geq 1}$ of numbers with the following property: for every triangulation t with a hole of perimeter p and v vertices in total, we have

$$\mathbb{P}(t \subset T) = a_v^p.$$

By their definition, the PSHT \mathbb{T}_λ are weakly Markovian. We denote by $a_v^p(\lambda) = C_p(\lambda)\lambda^v$ the associated constants, where $C_p(\lambda)$ is given by (7). This implies that any mixture of these is also weakly Markovian. Indeed, for any random variable Λ with values in $(0, \lambda_c]$, we denote by \mathbb{T}_Λ the PSHT with random parameter Λ . Let also μ be the distribution of Λ . Then for every triangulation t with a hole of perimeter p and v vertices in total, we have

$$\mathbb{P}(t \subset \mathbb{T}_\Lambda) = \int_0^{\lambda_c} \mathbb{P}(t \subset \mathbb{T}_\lambda) \mu(d\lambda) = \int_0^{\lambda_c} C_p(\lambda)\lambda^v \mu(d\lambda) =: a_v^p[\mu]. \quad (20)$$

Note that the last integral always converges since $C_p(\lambda)\lambda^v$ is bounded by 1. Therefore, the triangulation \mathbb{T}_Λ is weakly Markovian. Our goal here is to prove Theorem 2, which states that the converse is true.

In the remainder of this section, we will denote by T some weakly Markovian random triangulation, and by $(a_v^p)_{v \geq p \geq 1}$ the associated constants. Before giving an idea of the proof, let us start with a remark that will be very useful in all that follows. The numbers a_v^p are linked to each other by linear equations that we call the *peeling equations*. In this section, for every $p \geq 1$ and $j \geq 0$, we denote by $|\mathcal{T}_p(j)|$ the number of planar triangulations of a p -gon (rooted on the boundary) with exactly j inner vertices⁶. We also adopt the convention that $|\mathcal{T}_2(0)| = 1$ (i.e. a hole of perimeter 2 can be filled by simply gluing the two edges to each other). On the other hand $|\mathcal{T}_1(0)| = 0$.

Lemma 14. For every $v \geq p \geq 1$, we have

$$a_v^p = a_{v+1}^{p+1} + 2 \sum_{i=0}^{p-1} \sum_{j=0}^{+\infty} |\mathcal{T}_{i+1}(j)| a_{v+j}^{p-i}. \quad (21)$$

⁶In order to have nicer formulas, the convention we use here differs from the rest of the paper, in which the parameter n is related to the *total* number of vertices. This is why we do not use the τ notation. We insist that this holds only in Section 3.

In particular, the sum in the right-hand side must converge. Note also that if $v \geq p \geq 1$, then $v + j \geq p - i \geq 1$ in all the terms of the sum, so all the terms make sense.

Proof of Lemma 14. The proof just consists of making one peeling step. Assume that $t \subset T$ for some triangulation t with a hole of perimeter p and volume v . Fix an edge $e \in \partial T$, and consider the face f out of t that is adjacent to e . Then we are in exactly one of the three following cases:

- the third vertex of f does not belong to ∂t ,
- the third vertex of f belongs to ∂t , and f separates on its left i edges of ∂t from infinity,
- the third vertex of f belongs to ∂t , and f separates on its right i edges of ∂t from infinity.

In the first case, the triangulation we obtain by adding f to t has perimeter $p + 1$ and $v + 1$ vertices in total. In the other two cases, f separates $T \setminus t$ in two components, one finite and one infinite. The finite component has perimeter $i + 1$. If it has j inner vertices, then after filling the finite component, we obtain a triangulation with perimeter $p - i$ and volume $v + j$. \square

Our main task will be to prove the following result: it states that the numbers a_v^1 for $v \geq 1$ are compatible with some mixture of PSHT.

Proposition 15. There is a probability measure μ on $(0, \lambda_c]$ such that, for every $v \geq 1$, we have

$$a_v^1 = a_v^1[\mu].$$

Once Proposition 15 is proved, Theorem 2 easily follows. Indeed, Lemma 14 can be rewritten

$$a_{v+1}^{p+1} = a_v^p - 2 \sum_{i=0}^{p-1} \sum_{j=0}^{+\infty} |\mathcal{T}_{i+1}(j)| a_{v+j}^{p-i} \quad (22)$$

for $v+1 \geq p+1 \geq 2$. Hence, numbers of the form a_v^{p+1} can be expressed in terms of the numbers a_v^i with $i \leq p$. On the other hand, the numbers $a_v^p[\mu]$ also satisfy (22). Therefore, by induction on p , we can prove that for every $v \geq p \geq 1$, we have

$$a_v^p = a_v^p[\mu].$$

Since the numbers a_v^p characterize entirely the distribution of T , we are done.

Therefore, we only need to prove Proposition 15. As a particular case of (20), for every measure μ and any $v \geq 0$, we have

$$a_{v+1}^1[\mu] = \int_0^{\lambda_c} C_1(\lambda) \lambda^{v+1} \mu(d\lambda) = \int_0^{\lambda_c} \lambda^v \mu(d\lambda).$$

The right-hand side can be interpreted as the v -th moment of a random variable with distribution μ . Therefore, all we need to prove is that there is a random variable Λ with support in $(0, \lambda_c]$ such that, for every $v \geq 0$, we have

$$a_{v+1}^1 = \mathbb{E}[\Lambda^v].$$

We first show that there exists such a variable with support in $[0, 1]$. Since $a_1^1 = 1$, this is precisely the Hausdorff moment problem [28], so all we need to show is that the sequence $(a_{v+1}^1)_{v \geq 0}$ is completely monotonic. More precisely, let Δ be the discrete derivative operator:

$$(\Delta u)_n = u_n - u_{n+1}.$$

Then, to prove that $(a_{v+1}^1)_{v \geq 0}$ is the moment sequence of a probability measure on $[0, 1]$, it is sufficient to prove the following lemma.

Lemma 16. For every $k \geq 0$ and $v \geq p \geq 1$, we have

$$\left(\Delta^k a^p\right)_v \geq 0.$$

Note that the case $p = 1$ is sufficient. However, it will be more convenient to prove the lemma in the general case.

Proof. We prove the lemma by induction on k . The case $k = 0$ just means that $a_v^p \geq 0$ for every $v \geq p \geq 1$. The case $k = 1$ means that a_v^p is non-decreasing in v , which is a straightforward consequence of the peeling equation: in the right-hand side of (21), the term corresponding to $i = 0$ and $j = 1$ is $|\mathcal{T}_1(1)|a_{v+1}^p$, so

$$a_v^p \geq 2a_{v+1}^p \geq a_{v+1}^p.$$

Now assume that $k \geq 1$ and that the lemma is proved for k . We will use the result for k and $p + 1$ to prove it for $k + 1$ and p . Let $v \geq p \geq 1$. By using the induction hypothesis and (22) for $(p, v), (p, v + 1), \dots, (p, v + k)$, we obtain

$$\begin{aligned} 0 &\leq \Delta^k(a^{p+1})_{v+1} \\ &= \sum_{\ell=0}^k (-1)^\ell \binom{k}{\ell} a_{v+1+\ell}^{p+1} \\ &= \sum_{\ell=0}^k (-1)^\ell \binom{k}{\ell} a_{v+\ell}^p - 2 \sum_{i=0}^{p-1} \sum_{j=0}^{+\infty} |\mathcal{T}_{i+1}(j)| \sum_{\ell=0}^k (-1)^\ell \binom{k}{\ell} a_{v+\ell+j}^{p-i} \\ &= \left(\Delta^k a^p\right)_v - 2 \sum_{i=0}^{p-1} \sum_{j=0}^{+\infty} |\mathcal{T}_{i+1}(j)| \left(\Delta^k a^{p-i}\right)_{v+j}. \end{aligned}$$

By the induction hypothesis, all the terms $\left(\Delta^k a^{p-i}\right)_{v+j}$ in the sum are nonnegative. Therefore, the above sum does not decrease if we remove some of the terms. In particular, we may remove all the terms except the one for which $i = 0$ and $j = 1$, and we may also remove the factor 2. Since $|\mathcal{T}_1(1)| = 1$, we obtain

$$0 \leq \left(\Delta^k a^p\right)_v - \left(\Delta^k a^p\right)_{v+1} = \left(\Delta^{k+1} a^p\right)_v,$$

which proves the lemma by induction. \square

Remark 17. Many bounds used in the last proof may seem very crude. This is due to the fact that we prove the existence of a variable Λ with support in $[0, 1]$, whereas its support is actually in $[0, \lambda_c]$. For example, if we had not removed the factors 2, we would have obtained that Λ has support in $[0, \frac{1}{2}]$ instead of $[0, 1]$.

End of the proof of Theorem 2. By Lemma 16, there is a random variable Λ with support in $[0, 1]$ such that, for every $v \geq 1$, we have $a_{v+1}^1 = \mathbb{E}[\Lambda^v]$. All we need to show is that $\Lambda \in (0, \lambda_c]$ almost surely. We first explain why $\Lambda \leq \lambda_c$. The peeling equation for $p = v = 1$ shows that

$$\sum_{j=0}^{+\infty} |\mathcal{T}_1(j)| a_{j+1}^1 \leq a_1^1 < +\infty.$$




Figure 5: The "triangle chain" triangulation t_p with p vertices (for $p = 8$).

On the other hand, we know that $|\mathcal{T}_1(j)| \sim c\lambda_c^{-j}j^{-5/2}$ as $j \rightarrow +\infty$ for some constant $c > 0$. Therefore, if $\mathbb{P}(\Lambda \geq \lambda_c + \varepsilon) \geq \varepsilon$ for some $\varepsilon > 0$, then $a_{v+1}^1 \geq \varepsilon(\lambda_c + \varepsilon)^v$ for every v and the series above diverge, so we get a contradiction.

We finally prove that $\mathbb{P}(\Lambda = 0) = 0$. As explained above, we already know that

$$a_v^p = \int_0^{\lambda_c} a_v^p(\lambda)\mu(d\lambda),$$

where μ is the distribution of Λ . We also have $a_v^p(0) = \mathbb{1}_{p=v}$. Therefore, if $\delta = \mathbb{P}(\Lambda = 0)$, for every $p \geq 1$, we have $a_p^p \geq \delta$. Now let t_p be the triangulation with p vertices and a hole of perimeter p represented on Figure 5. For every $p \geq 1$, we have $\mathbb{P}(t_p \subset T) = a_p^p \geq \delta$. Since the events $t_p \subset T$ are nonincreasing in p , we have

$$\mathbb{P}(\forall p \geq 1, t_p \subset T) \geq \delta.$$

On the other hand, if $t_p \subset T$, then the degree of the root vertex is at least $p + 2$, so we have $\mathbb{P}(\deg(\rho) = +\infty) \geq \delta$. Since the degrees are finite, we must have $\delta = 0$, so Λ is supported in $(0, \lambda_c]$, which concludes the proof. \square

4 Ergodicity

Throughout this section, we denote by (g_n) a sequence such that $\frac{g_n}{n} \rightarrow \theta \in [0, \frac{1}{2})$ and by T a subsequential limit of (T_{n,g_n}) for d_{loc} . In order to keep the notation light, we will always implicitly restrict ourselves to values of n along which T_{n,g_n} converges to T in distribution.

For any n , the triangulation T_{n,g_n} is weakly Markovian, so this is also the case of T . Therefore, by Theorem 2, we know that T must be a mixture of PSHT, i.e. there is a random variable $\Lambda \in (0, \lambda_c]$ such that T has the same distribution as the PSHT with random parameter \mathbb{T}_Λ .

We also note right now that by the discussion in the end of Section 1.3 (Equation (8)), the parameter Λ is a measurable function of the triangulation \mathbb{T}_Λ . More precisely, if P^Λ and V^Λ are the perimeter and volume processes associated with the peeling exploration of \mathbb{T}_Λ , then Λ can be defined as $f^{-1}\left(\lim_{i \rightarrow +\infty} \frac{P^\Lambda(i)}{V^\Lambda(i)}\right)$, where $f(\lambda) = 1 - 4h$ is injective. In particular Λ is defined without any ambiguity. Our goal is now to prove that Λ is deterministic, and given by (2).




Figure 6: The idea of the two-holes argument: the inner parts around e_n^1 and e_n^2 are "swapped". The root edges e_n^1 and e_n^2 are swapped at the same time as their neighbourhoods.

4.1 The two holes argument

Roughly speaking, we know that T_{n,g_n} seen from a typical point e_n looks like a PSHT with random parameter Λ . We would like to show that Λ does not depend on (T_{n,g_n}, e_n) . The first step is essentially to prove that Λ does not depend on e_n .

More precisely, conditionally on T_{n,g_n} , we pick two independent uniform oriented edges e_n^1, e_n^2 of T_{n,g_n} . The pairs (T_{n,g_n}, e_n^1) and (T_{n,g_n}, e_n^2) have the same distribution, and both converge in distribution to \mathbb{T}_Λ . It follows that the pairs

$$((T_{n,g_n}, e_n^1), (T_{n,g_n}, e_n^2))$$

for $n \geq 1$ are tight, so up to further extraction, they converge to a pair $(\mathbb{T}_{\Lambda_1}^1, \mathbb{T}_{\Lambda_2}^2)$, where both marginals have the same distribution as \mathbb{T}_Λ . By the above discussion, the variables Λ_1 and Λ_2 are well defined. By the Skorokhod representation theorem, we may assume the convergence in distribution is almost sure. Our goal in this subsection is to prove the following lemma.

Proposition 18. Under the convergence

$$((T_{n,g_n}, e_n^1), (T_{n,g_n}, e_n^2)) \xrightarrow[n \rightarrow +\infty]{a.s.} (\mathbb{T}_{\Lambda_1}^1, \mathbb{T}_{\Lambda_2}^2),$$

we have $\Lambda_1 = \Lambda_2$ almost surely.

Proof. The idea of the proof is the following: if $\Lambda_1 \neq \Lambda_2$, consider two large neighbourhoods N_n^1 and N_n^2 (say, of size 100) with the same perimeter around e_n^1 and e_n^2 . Then there is a natural way to "swap" N_n^1 and N_n^2 in T_{n,g_n} (cf. Figure 6), without changing its distribution. If we do this, then T_{n,g_n} around e_n^1 looks like \mathbb{T}_{Λ_1} in a neighbourhood of the root of size 100, and like \mathbb{T}_{Λ_2} out of this neighbourhood. However, if $\Lambda_1 \neq \Lambda_2$, such a configuration is highly unlikely for a mixture of PSHT.

More precisely, in the following proof, we consider a deterministic peeling algorithm \mathcal{A} , according to which we will explore T_{n,g_n} around e_n^1 and around e_n^2 . All the explorations we will consider will be filled-in: everytime the peeled face separates the undiscovered part in two, we discover completely the smaller part. While the computations in the beginning of the proof hold for any choice of \mathcal{A} , we will need later to specify the choice of the algorithm.

We denote by $\mathcal{E}_n^1(i)$ the explored part at time i when the exploration is started from e_n^1 . If at some point $\mathcal{E}_n^1(i)$ is non-planar, then the exploration is stopped. We write respectively $P_n^1(i)$

and $V_n^1(i)$ for its perimeter and its volume (i.e. its total number of vertices). For any $p \geq 2$, let $\tau_n^1(p)$ be the smallest i such that $P_n^1(i) = p$. We denote by $\mathcal{E}_n^2(i)$, $P_n^2(i)$, $V_n^2(i)$ and $\tau_n^2(p)$ the analog quantities for the exploration started from e_n^2 . We also denote by $\mathcal{E}_\infty^1(i)$, $P_\infty^1(i)$, $V_\infty^1(i)$ and $\tau_\infty^1(p)$ the analog quantities for $\mathbb{T}_{\Lambda_1}^1$, and similarly for $\mathbb{T}_{\Lambda_2}^2$.

We fix $\varepsilon > 0$. We recall that

$$\frac{P_\infty^1(i)}{V_\infty^1(i)} \xrightarrow[i \rightarrow +\infty]{a.s.} f(\Lambda_1),$$

where $f : (0, \lambda_c] \rightarrow [0, 1)$ is bijective and decreasing. Note also that the times $\tau_\infty^1(p)$ are a.s. finite and $\tau_\infty^1(p) \rightarrow +\infty$ when $p \rightarrow +\infty$, so for p large enough, we have

$$\mathbb{P} \left(\left| \frac{p}{V_\infty^1(\tau_\infty^1(p))} - f(\Lambda_1) \right| \geq \varepsilon \right) \leq \varepsilon.$$

We fix such a p until the end of the proof. Moreover, we have

$$\frac{q}{V_\infty^1(\tau_\infty^1(q)) - V_\infty^1(\tau_\infty^1(p))} \xrightarrow[q \rightarrow +\infty]{a.s.} f(\Lambda_1),$$

so for q large enough, we have

$$\mathbb{P} \left(\left| \frac{q}{V_\infty^1(\tau_\infty^1(q)) - V_\infty^1(\tau_\infty^1(p))} - f(\Lambda_1) \right| \geq \varepsilon \right) \leq \varepsilon.$$

We fix a such a $q > p$ until the end of the proof. Moreover, the local convergence of (T_{n,g_n}, e_n^1) to $\mathbb{T}_{\Lambda_1}^1$ implies the convergence of the peeling exploration. Hence, the probability that $\tau_n^1(p)$ is finite goes to 1 as $n \rightarrow +\infty$. Therefore, for n large enough, we have

$$\mathbb{P} \left(\left| \frac{p}{V_n^1(\tau_n^1(p))} - f(\Lambda_1) \right| \geq 2\varepsilon \right) \leq 2\varepsilon, \quad (23)$$

$$\mathbb{P} \left(\left| \frac{q}{V_n^1(\tau_n^1(q)) - V_n^1(\tau_n^1(p))} - f(\Lambda_1) \right| \geq 2\varepsilon \right) \leq 2\varepsilon. \quad (24)$$

By combining (23) and (24), we also obtain

$$\mathbb{P} \left(\left| \frac{q}{V_n^1(\tau_n^1(q)) - V_n^1(\tau_n^1(p))} - \frac{p}{V_n^1(\tau_n^1(p))} \right| \geq 4\varepsilon \right) \leq 4\varepsilon \quad (25)$$

for n large enough. Moreover, the same is true if we replace the exploration from e_n^1 by the exploration from e_n^2 (with of course Λ_2 playing the role of Λ_1).

We now specify the properties that we need our peeling algorithm \mathcal{A} to satisfy. It roughly means that the edges that we peel after time $\tau(p)$ should not depend on what happened before time $\tau(p)$. More precisely, before the time $\tau(p)$, the peeled edge can be any edge of the boundary. At time $\tau(p)$, we color in red an edge of the boundary of $\mathcal{E}(\tau(p))$ according to some deterministic convention. At time $i \geq \tau(p)$, the choice of the edge to peel at time $i + 1$ must only depend on the triangulation $\mathcal{E}(i) \setminus \mathcal{E}(\tau(p))$, which is rooted at the red edge (see Figure 7). It is easy to see that such an algorithm exists: we only need to fix a peeling algorithm \mathcal{A}' for triangulations of the sphere or the plane, and a peeling algorithm \mathcal{A}'' for triangulations of the p -gon, and to use \mathcal{A}' before time $\tau(p)$ and \mathcal{A}'' after $\tau(p)$. Note also that we can know, by only looking at $\mathcal{E}(i)$, if $i \leq \tau(p)$ or not, so at each step the peeled edge will indeed depend only on the explored part.




Figure 7: Illustration of the choice of the peeling algorithm: in gray, the part $\mathcal{E}(\tau(p))$. In cyan, the part $\mathcal{E}(i) \setminus \mathcal{E}(\tau(p))$, rooted at the red edge. The peeled edge (in blue) may only depend on the cyan part (which is rooted at the red edge).

We can now define precisely our surgery operation on $(T_{n,g_n}, e_n^1, e_n^2)$. We say that e_n^1 and e_n^2 are *well separated* if $\tau_n^1(p), \tau_n^2(p) < +\infty$ and if the regions $\mathcal{E}_n^1(\tau_n^1(p))$ and $\mathcal{E}_n^2(\tau_n^2(p))$ have no common vertex. If e_n^1 and e_n^2 are well separated, we remove the triangulations $\mathcal{E}_n^1(\tau_n^1(p))$ and $\mathcal{E}_n^2(\tau_n^2(p))$, which creates two holes of perimeter p . We then glue each of the two triangulations to the hole where the other was, in such a way that the red edges of our peeling algorithm coincide (see Figure 6). If the two roots are not well-separated, then $(T_{n,g_n}, e_n^1, e_n^2)$ is left unchanged. This operation is an involution on the set of bi-rooted triangulations with fixed size and genus. Therefore, if we denote by T'_{n,g_n} the new triangulation we obtain, it is still uniform, and $(T'_{n,g_n}, e_n^1, e_n^2)$ has the same distribution as $(T_{n,g_n}, e_n^1, e_n^2)$.

Lemma 19. The probability that e_n^1 and e_n^2 are well separated goes to 1 as $n \rightarrow +\infty$.

Proof. The local convergence of the exploration implies that the sequence $(\tau_n^1(p))_{n \geq 0}$ is tight, and so is $(V_n^1(\tau_n^1(p)))_{n \geq 0}$. Since the diameter of a triangulation is bounded by its volume, the diameters of the maps $\mathcal{E}_n^1(\tau_n^1(p))$ are tight (when $n \rightarrow +\infty$), and the same is true around e_n^2 . On the other hand, by local convergence (for d_{loc}), for every fixed r , the volume of the ball of radius r around e_n^1 is tight. Since e_n^2 is uniform and independent of e_n^1 , we have

$$\mathbb{P}(d_{T_{n,g_n}}(e_n^1, e_n^2) \leq r) \xrightarrow{n \rightarrow +\infty} 0$$

for every $r \geq 0$. From here, we obtain

$$\mathbb{P}(d_{T_{n,g_n}}(e_n^1, e_n^2) \leq \text{diam}(\mathcal{E}_n^1(\tau_n^1(p))) + \text{diam}(\mathcal{E}_n^2(\tau_n^2(p)))) \xrightarrow{n \rightarrow +\infty} 0.$$

If this last event occurs, then the regions $\mathcal{E}_n^1(\tau_n^1(p))$ and $\mathcal{E}_n^2(\tau_n^2(p))$ do not intersect, which proves the lemma. \square

Now, if we perform a peeling exploration on T'_{n,g_n} from e_n^1 using algorithm \mathcal{A} , then the part explored at time $\tau(p)$ will be exactly the triangulation $\mathcal{E}_n^1(\tau_n^1(p))$. Moreover, the red edge on $\partial\mathcal{E}_n^1(\tau_n^1(p))$ is glued at the same place as the red edge on $\partial\mathcal{E}_n^2(\tau_n^2(p))$ and our peeling algorithm "forgets" the interior of $\mathcal{E}_n^1(\tau_n^1(p))$. Hence, the part explored between $\tau(p)$ and $\tau(q)$ is exactly $\mathcal{E}_n^2(\tau_n^2(q)) \setminus \mathcal{E}_n^2(\tau_n^2(p))$ (this is where we need the algorithm \mathcal{A} to satisfy the property described

above). Therefore, the part of T'_{n,g_n} discovered between times $\tau(p)$ and $\tau(q)$ has perimeter q and volume

$$V_n^2(\tau_n^2(q)) - V_n^2(\tau_n^2(p)).$$

Now, since (T'_{n,g_n}, e_n^1) has the same distribution as (T_{n,g_n}, e_n^1) , we can apply (25) to the exploration of T'_{n,g_n} from e_n^1 . We obtain

$$\mathbb{P}\left(\left|\frac{q}{V_n^2(\tau_n^2(q)) - V_n^2(\tau_n^2(p))} - \frac{p}{V_n^1(\tau_n^1(p))}\right| \geq 4\varepsilon\right) \leq 4\varepsilon.$$

By combining this with (23) for the exploration of T_{n,g_n} started from e_n^1 and with (24) for the exploration of T_{n,g_n} started from e_n^2 , we obtain

$$\mathbb{P}(|f(\Lambda_1) - f(\Lambda_2)| \geq 8\varepsilon) \leq 8\varepsilon.$$

This is true for any $\varepsilon > 0$, so we have $f(\Lambda_1) = f(\Lambda_2)$ a.s. Since f is strictly decreasing on $(0, \lambda_c]$, we are done. \square

4.2 Conclusion

In order to conclude the proof of Theorem 1, we will need to compute the mean inverse root degree in the PSHT. This is the only observable that is easy to compute in T_{n,g_n} , so it is not surprising that such a result is needed to link θ to λ .

Proposition 20. For $\lambda \in (0, \lambda_c]$, let

$$d(\lambda) = \mathbb{E}\left[\frac{1}{\deg_{\mathbb{T}_\lambda}(\rho)}\right].$$

Then we have

$$d(\lambda) = \frac{h \log \frac{1+\sqrt{1-4h}}{1-\sqrt{1-4h}}}{(1+8h)\sqrt{1-4h}},$$

where h is given by (6). In particular, the function d is strictly increasing on $(0, \lambda_c]$, with $d(\lambda_c) = \frac{1}{6}$ and $\lim_{\lambda \rightarrow 0} d(\lambda) = 0$.

We postpone the proof of Proposition 20 until Section 4.3, and finish the proof of Theorem 1. We recall that we work with a subsequential limit of (T_{n,g_n}) , and we know that it has the same distribution as \mathbb{T}_Λ for some random Λ . Our goal is to prove that Λ is the solution to $d(\Lambda) = \frac{1-2\theta}{6}$. Note that by Proposition 20, the solution to this equation exists and is unique.

The idea is the following: by Proposition 18, the random parameter Λ only depends on the triangulation T_{n,g_n} and not on its root. On the one hand, for any triangulation, the mean inverse root degree over all possible choices of the root is $\frac{1-2\theta}{6}$. Therefore, the mean inverse root degree over all triangulations corresponding to a fixed value of Λ is $\frac{1-2\theta}{6}$. On the other hand, the mean inverse degree conditionally on Λ is $d(\Lambda)$, so we should have $d(\Lambda) = \frac{1-2\theta}{6}$.

To prove this properly, we need a way to "read" Λ on finite maps, which is the goal of the next (easy) technical lemma. As earlier, we restrict ourselves to a subset of values of n along which $T_{n,g_n} \rightarrow \mathbb{T}_\Lambda$.

Lemma 21. There is a function $\ell : \bigcup_{n \geq 1} \mathcal{T}(n, g_n) \rightarrow [0, \lambda_c]$ such that we have the convergence

$$(T_{n,g_n}, \ell(T_{n,g_n})) \xrightarrow[n \rightarrow +\infty]{(d)} (\mathbb{T}_\Lambda, \Lambda).$$

Note that a priori $\ell(t)$ may depend on the choice of the root of t , although Proposition 18 will guarantee that this is asymptotically not the case.

Proof. By the Skorokhod representation theorem, we may assume the convergence $T_{n,g_n} \rightarrow \mathbb{T}_\Lambda$ is almost sure. As explained in the beginning of Section 4, the parameter Λ is a measurable function of \mathbb{T}_Λ , so there is a measurable function $\tilde{\ell}$ from the set of all (finite or infinite) triangulations such that $\tilde{\ell}(\mathbb{T}_\Lambda) = \Lambda$ a.s. Therefore, for any $\varepsilon > 0$, we can find a continuous function (for the local topology) ℓ_ε such that

$$\mathbb{P} \left(|\ell_\varepsilon(\mathbb{T}_\Lambda) - \Lambda| \geq \frac{\varepsilon}{2} \right) \leq \frac{\varepsilon}{2}.$$

For n larger than some N_ε , we have

$$\mathbb{P} (|\ell_\varepsilon(T_{n,g_n}) - \Lambda| \geq \varepsilon) \leq \varepsilon.$$

Now let $\varepsilon_k \rightarrow 0$ as $k \rightarrow +\infty$. We may assume that (N_{ε_k}) is strictly increasing. For any triangulation $t \in \mathcal{T}(n, g_n)$, we set $\ell(t) = \ell_{\varepsilon_k}(t)$, where k is such that $N_{\varepsilon_k} < n \leq N_{\varepsilon_{k+1}}$. This ensures $\ell(T_{n,g_n}) \rightarrow \Lambda$ almost surely, so

$$(T_{n,g_n}, \ell(T_{n,g_n})) \xrightarrow[n \rightarrow +\infty]{(d)} (\mathbb{T}_\Lambda, \Lambda).$$

□

We are now ready to prove our main Theorem.

Proof of Theorem 1. Let f be a continuous, bounded function on $[0, \lambda_c]$. For any n , let e_n^1 and e_n^2 be two oriented edges chosen independently and uniformly in T_{n,g_n} , and let ρ_n^1 and ρ_n^2 be their starting points. We will estimate in two different ways the limit of the quantity

$$\mathbb{E} \left[\frac{1}{\deg(\rho_n^1)} f(\ell(T_{n,g_n}, e_n^1)) \right]. \quad (26)$$

On the one hand, the quantity in the expectation is bounded and converges in distribution to $\frac{1}{\deg(\rho)} f(\Lambda)$, where ρ is the root vertex of \mathbb{T}_Λ , so (26) goes to

$$\mathbb{E} \left[\frac{1}{\deg(\rho)} f(\Lambda) \right] = \mathbb{E} [d(\Lambda) f(\Lambda)]$$

as $n \rightarrow +\infty$. On the other hand, by Proposition 18 and Lemma 21, we have

$$(\ell(T_{n,g_n}, e_n^1), \ell(T_{n,g_n}, e_n^2)) \xrightarrow[n \rightarrow +\infty]{(d)} (\Lambda, \Lambda),$$

so $\ell(T_{n,g_n}, e_n^1) - \ell(T_{n,g_n}, e_n^2)$ goes to 0 in probability. Since f is bounded and uniformly continuous, the difference $f(\ell(T_{n,g_n}, e_n^1)) - f(\ell(T_{n,g_n}, e_n^2))$ goes to 0 in L^1 , so we have

$$\mathbb{E} \left[\frac{1}{\deg(\rho_n^1)} f(\ell(T_{n,g_n}, e_n^1)) \right] = \mathbb{E} \left[\frac{1}{\deg(\rho_n^1)} f(\ell(T_{n,g_n}, e_n^2)) \right] + o(1).$$

Moreover, the expectation of $\frac{1}{\deg(\rho_n^1)}$ conditionally on (T_{n,g_n}, e_n^2) is equal to $\frac{|V(T_{n,g_n})|}{2|E(T_{n,g_n})|} = \frac{n+2-2g_n}{6n}$, so we can write

$$\mathbb{E} \left[\frac{1}{\deg(\rho_n^1)} f(\ell(T_{n,g_n}, e_n^1)) \right] = \frac{n+2-2g_n}{6n} \mathbb{E} [f(\ell(T_{n,g_n}, e_n^2))] + o(1) \xrightarrow[n \rightarrow +\infty]{} \frac{1-2\theta}{6} \mathbb{E} [f(\Lambda)].$$

Therefore, for any bounded, continuous f on $[0, \lambda_c]$, we have

$$\mathbb{E}[d(\Lambda)f(\Lambda)] = \frac{1-2\theta}{6}\mathbb{E}[f(\Lambda)],$$

so $d(\Lambda) = \frac{1-2\theta}{6}$ a.s. By injectivity of d , this fixes a deterministic value for Λ that depends only on θ , which concludes the proof. \square

4.3 The average root degree in the type-I PSHT via uniform spanning forests

Our goal is now to prove Proposition 20. Since we have not been able to perform a direct computation, our strategy will be to use the similar computation for the type-II PSHT that is done in [12, Appendix B]. To link the mean degree in the type-I and in the type-II PSHT, we use the core decomposition of the type-I PSHT [12, Appendix A] and an interpretation in terms of the Wired Uniform Spanning Forest.

The type-II case. We will need to use the computation of the same quantity in the type-II case, which is performed in [12, Appendix B]. We first recall briefly the definition of the type-II PSHT [20]. These are the type-II analogues of the type-I PSHT, i.e. they may contain multiple edges but no loop. They form a one-parameter family $(\mathbb{T}_\kappa^{II})_{0 < \kappa \leq \kappa_c}$ of random infinite triangulations of the plane, where $\kappa_c = \frac{2}{27}$. Their distribution is characterized as follows: for any type-II triangulation with a hole of perimeter p and v vertices in total, we have

$$\mathbb{P}(t \subset \mathbb{T}_\kappa^{II}) = C_p^{II}(\kappa) \times \kappa^v,$$

where the numbers $C_p^{II}(\kappa)$ are explicitly determined by κ .

Proposition 22. [12, Appendix B] For $\kappa \in (0, \kappa_c]$, let

$$d_{II}(\kappa) = \mathbb{E}\left[\frac{1}{\deg_{\mathbb{T}_\kappa^{II}}(\rho)}\right].$$

Then we have

$$d_{II}(\kappa) = -\frac{1-\alpha}{2} + \frac{(1-\alpha)\sqrt{\alpha}}{2\sqrt{3\alpha-2}} \log \frac{\sqrt{\alpha} + \sqrt{3\alpha-2}}{\sqrt{\alpha} - \sqrt{3\alpha-2}},$$

where $\alpha \in [\frac{2}{3}, 1)$ is such that $\kappa = \frac{\alpha^2(1-\alpha)}{2}$.

In [12], Proposition 22 is proved by using a peeling exploration to obtain an equation on the generating function of the mean inverse root degree in PSHT with boundaries. While this approach could theoretically also work in the type-I setting, some technical details make the formulas much more complicated. Although everything remains in theory completely solvable, our efforts to push the computation until the end have failed, even with a computer algebra software. Therefore, we will rely on Proposition 22 and on the correspondence between type-I PSHT and type-II PSHT given in [12, Appendix A].




Figure 8: The random block A_λ (here we have $B = 4$). The green parts are i.i.d. λ -Boltzmann triangulations of the 1-gon.

The type-I–type-II correspondence. We now introduce the correspondence between the type-I and the type-II PSHT stated in [12, Appendix A]. We write $\beta = 1 - \frac{1+2h}{\sqrt{1+8h}}$, and $\kappa = \frac{h}{(1+2h)^3}$, where h is given by (6). We also define a random triangulation A_λ of the 2-gon as follows. We start from two vertices x and y linked by B edges, where $B - 1$ is a geometric variable of parameter β . In each of the $B - 1$ slots between these edges, we insert a loop (the loops are glued to x or y according to independent fair coin flips). Finally, these $B - 1$ loops are filled with i.i.d. Boltzmann triangulations of the 1-gon with parameter λ (see Figure 8). Copies of A_λ will be called *blocks* in what follows.

Let $|A_\lambda|$ be the number of inner vertices of A_λ (i.e. x and y excluded), and let $E(A_\lambda)$ be its set of edges (including the boundary edges). The Euler formula implies $|E(A_\lambda)| = 1 + 3|A_\lambda|$. Let \tilde{A}_λ be a random triangulation with the same distribution as A_λ , biased by $|E(A_\lambda)|$. If T is an infinite type-II triangulation of the plane and $(a^e)_{e \in E(T)}$ are blocks, we denote by $\Phi(T, (a^e))$ the triangulation obtained from T by replacing each edge e of T by a^e . After elementary computations, the correspondence shown in [12, Appendix A] can be reformulated as follows.

Proposition 23. The triangulation \mathbb{T}_λ has the same law as $\Phi\left(\mathbb{T}_\kappa^{II}, (A_\lambda^e)_{e \in E(\mathbb{T}_\kappa^{II})}\right)$, where conditionally on \mathbb{T}_κ^{II} :

- the blocks A_λ^e are independent;
- the block $A_\lambda^{e_0^{II}}$, where e_0^{II} is the root edge of \mathbb{T}_κ^{II} , has the law of \tilde{A}_λ , rooted at a uniform oriented edge of \tilde{A}_λ ;
- for $e \neq e_0^{II}$, the block A_λ^e has the law of A_λ .

In the remainder of this section, we will consider that \mathbb{T}_λ and \mathbb{T}_κ^{II} are coupled together as described by Proposition 23. It will also be useful to consider the triangulation obtained from \mathbb{T}_κ^{II} in a similar way, but where $A_\lambda^{e_0^{II}}$ has the law of A_λ instead of \tilde{A}_λ . We denote this triangulation by \mathbb{T}'_λ . We note right now that \mathbb{T}_λ is absolutely continuous with respect to \mathbb{T}'_λ .

The wired uniform spanning forest. The last ingredient that we need to prove Proposition 20 is the oriented wired uniform spanning forest (OWUSF) [9], which may be defined on any transient graph. We first recall its definition. We recall that if (γ_i) is a transient path, its *loop-erasure* is the path (γ_{i_j}) , where $i_0 = 0$ and $i_{j+1} = 1 + \max\{i \geq i_j | \gamma(i) = \gamma(i_j)\}$. If G is

an infinite transient graph, its OWUSF \vec{F} is generated by the following generalization of the Wilson algorithm [42].

1. We order the vertices of G as $(v_n)_{n \geq 1}$.
2. We run a simple random walk (X_n) on G started from v_1 , and consider its loop-erasure, which is an oriented path from v_1 to ∞ . We add this path to \vec{F} .
3. We consider the first i such that v_i does not belong to any edge of \vec{F} yet. We run a SRW from v_i , stopped when it hits \vec{F} , and add its loop-erasure to \vec{F} .
4. We repeat step 3 infinitely many times.

It is immediate that \vec{F} is an oriented forest such that for any vertex v of G , there is exactly one oriented edge of \vec{F} going out of v . It can be checked that the distribution of \vec{F} does not depend on the way the vertices are ordered.

Stationarity and reversibility. We recall that a random rooted graph G is *reversible* if its distribution is invariant under reversing the orientation of the root edge. It is *stationary* if its distribution is invariant under re-rooting G at the first step of the simple random walk. In particular, we recall from [20] that the type-II PSHT are stationary and reversible. Moreover, the proof of [20] also works for the type-I PSHT. Similarly, a random rooted, forest-decorated graph (G, \vec{e}_0, \vec{F}) is called *reversible* (resp. *stationary*) if its distribution is invariant under reversing the root edge (resp. re-rooting at the first step of the SRW).

If two random, rooted transient graphs (G, \vec{e}_0) and (G', \vec{e}'_0) have the same distribution and if \vec{F} and \vec{F}' are their respective OWUSF, then (G, \vec{e}_0, \vec{F}) and $(G', \vec{e}'_0, \vec{F}')$ also have the same distribution. It follows that if (G, \vec{e}_0) is stationary and reversible, so is the triplet (G, \vec{e}_0, \vec{F}) . The link between the OWUSF and the mean inverse root degree is given by the next lemma.

Lemma 24. Let (G, \vec{e}_0) be a stationary, reversible and transient random rooted graph, and let ρ be the starting point of the root edge \vec{e}_0 . We denote by \vec{F} the OWUSF of G . Then

$$\mathbb{P}(\vec{e}_0 \in \vec{F}) = \mathbb{E} \left[\frac{1}{\deg_G(\rho)} \right].$$

Proof. Conditionally on G , let \vec{e}_1 be a random oriented edge chosen uniformly (independently of \vec{F}) among all the edges leaving the root vertex. By invariance and reversibility of (G, \vec{e}_0) , we know that (G, \vec{e}_1, \vec{F}) has the same distribution as (G, \vec{e}_0, \vec{F}) . Therefore, we have

$$\mathbb{P}(\vec{e}_0 \in \vec{F}) = \mathbb{P}(\vec{e}_1 \in \vec{F}) = \mathbb{E} \left[\mathbb{P}(\vec{e}_1 \in \vec{F} | G, \vec{e}_0, \vec{F}) \right] = \mathbb{E} \left[\frac{1}{\deg_G(\rho)} \right],$$

where in the end we use the fact that, by construction, \vec{F} contains exactly one oriented edge going out of the root vertex. \square

Remark 25. Lemma 24 is the analogue for stationary, reversible graphs of the fact that the WUSF on any unimodular random graph has expected root degree 2 [3, Theorem 7.3]. These two properties can also be deduced from one another.

Proof of Proposition 20. By an easy uniform integrability argument, the function d is continuous in λ , so it is enough to prove the formula for $\lambda < \lambda_c$. In this case, all variants of \mathbb{T}_λ that we will consider are a.s. transient by results from [20], so the OWUSF on these graphs is well defined. We now fix $\lambda \in (0, \lambda_c)$. Let β and κ be given by Proposition 23. We will denote by \vec{e}_0 , \vec{e}'_0 and \vec{e}_0^{II} the respective root edges of \mathbb{T}_λ , \mathbb{T}'_λ and \mathbb{T}_κ^{II} , and by \vec{F} , \vec{F}' and \vec{F}^{II} their respective OWUSF.

By Proposition 23, conditionally on \mathbb{T}_κ^{II} and on the blocks $(A_\lambda^e)_{e \in E(\mathbb{T}_\kappa^{II})}$, the edge \vec{e}_0 is picked uniformly among the oriented edges of $A_\lambda^{e_0^{II}}$. Moreover, this choice is also independent of \vec{F} . Therefore, we have

$$\mathbb{P}(\vec{e}_0 \in \vec{F}) = \mathbb{E} \left[\mathbb{P}(\vec{e}_0 \in \vec{F} \mid \mathbb{T}_\kappa^{II}, (A_\lambda^e), \vec{F}) \right] = \mathbb{E} \left[\frac{|E(\tilde{A}_\lambda) \cap \vec{F}|}{2|E(\tilde{A}_\lambda)|} \right],$$

where we recall that $E(\tilde{A}_\lambda)$ is the set of edges of \tilde{A}_λ . Moreover, the pair $(\mathbb{T}_\lambda, \vec{F})$ has the same distribution as $(\mathbb{T}'_\lambda, \vec{F}')$ biased by $|E(A_\lambda^{e_0})|$. Therefore, we can write

$$\mathbb{P}(\vec{e}_0 \in \vec{F}) = \frac{\mathbb{E} \left[|E(A_\lambda^{e_0}) \cap \vec{F}'| \right]}{2\mathbb{E} [|E(A_\lambda)|]}. \quad (27)$$

Since the denominator can easily be explicitly computed by using the definition of the block A_λ , we first focus on the numerator. In a block, we call *principal edges* the edges joining the two boundary vertices of the block. We know that there is exactly one edge of \vec{F}' going out of every inner vertex of A_λ . Moreover, if a non-principal edge of A_λ belongs to \vec{F}' , then it goes out of an inner vertex (it cannot go out of a boundary vertex since the edges of \vec{F}' are "directed towards infinity"). Therefore, the number of non-principal edges of A_λ that belong to \vec{F}' is exactly $|A_\lambda|$. On the other hand, since \vec{F}' is a forest, it contains at most one principal edge of A_λ . Therefore, we have

$$\mathbb{E} \left[|E(A_\lambda^{e_0}) \cap \vec{F}'| \right] = \mathbb{E}[|A_\lambda|] + \mathbb{P}(\vec{F}' \text{ contains a principal edge of } A_\lambda^{e_0}).$$

To finish the computation, we claim that

$$\mathbb{P}(\vec{F}' \text{ contains a principal edge of } A_\lambda^{e_0}) = 2\mathbb{E} \left[\frac{1}{\deg_{\mathbb{T}_\kappa^{II}}(\rho)} \right] = 2d_{II}(\kappa). \quad (28)$$

Indeed, let ρ^{II} be the starting vertex of \vec{e}_0^{II} . By reversibility of \mathbb{T}_κ^{II} , the left-hand side of (28) is twice the probability that \vec{F}' contains a principal edge of $A_\lambda^{e_0}$ going out of ρ^{II} . Moreover, since \mathbb{T}_κ^{II} is stationary, the distribution of \mathbb{T}'_λ is stable under the following operation:

- first, we resample the root edge of \mathbb{T}_κ^{II} uniformly among all the oriented edges going out of ρ^{II} ;
- then, we pick the new root edge of \mathbb{T}'_λ uniformly among all the oriented edges of the block corresponding to the new root edge of \mathbb{T}_κ^{II} .

Note that this operation is different from resampling the root of \mathbb{T}'_λ uniformly among the edges going out of its root vertex, and that \mathbb{T}'_λ is *not* stationary.

Since \mathbb{T}'_λ is stationary for this operation, so is $(\mathbb{T}'_\lambda, \vec{F}')$. On the other hand, there is exactly one block incident to ρ^{II} in which a principal edge going out of ρ^{II} belongs to \vec{F}' . By the same argument as in the proof of Lemma 24, this implies (28). By combining Lemma 24 with (27) and (28), we obtain

$$d(\lambda) = \frac{\mathbb{E}[|A_\lambda|] + 2d_{II}(\kappa)}{2(1 + 3\mathbb{E}[|A_\lambda|])}.$$

By the definition of A_λ , we also have

$$\mathbb{E}[|A_\lambda|] = \frac{\beta}{1 - \beta} \times \frac{\lambda Z'_1(\lambda)}{Z_1(\lambda)} = \frac{2h}{1 + 2h},$$

where $Z_1(\lambda)$ is the partition function of λ -Boltzmann type-I triangulations of a 1-gon. Finally, the value of $d_{II}(\kappa)$ is given by Proposition 22. It is easy to obtain that the α of Proposition 22 is equal to $\frac{1}{1+2h}$, and we get the formula for $d(\lambda)$.

It remains to prove that this is an increasing function of λ (or equivalently of h). For this, set $x = \sqrt{1 - 4h}$. We want to prove that the function

$$x \rightarrow \frac{1 - x^2}{4x(3 - 2x^2)} \log \frac{1 + x}{1 - x}$$

is decreasing on $[0, 1]$. By computing the derivative with respect to x , this is equivalent to

$$(2x^4 - 3x^2 + 3)x \log \frac{1 + x}{1 - x} \geq 2x^2(3 - 2x^2)$$

for $0 < x < 1$. The power series expansion of the left-hand side is

$$6x^2 - 4x^4 + \sum_{k \geq 3} \left(\frac{6}{2k-1} - \frac{6}{2k-3} + \frac{4}{2k-5} \right) x^{2k} \geq 6x^2 - 4x^4.$$

Since all the terms in the sum are positive, this proves monotonicity. \square

5 Combinatorial asymptotics

The goal of this section is to prove Theorem 3. The proof relies on the lemma below, which is an easy consequence of Theorem 1. We recall that $\tau(n, g)$ is the number of triangulations of genus g with $2n$ faces, and that $\lambda(\theta)$ is the solution to (2).

Lemma 26. Let (g_n) be such that $\frac{g_n}{n} \rightarrow \theta$, with $0 \leq \theta < \frac{1}{2}$. Then

$$\frac{\tau(n-1, g_n)}{\tau(n, g_n)} \rightarrow \lambda(\theta)$$

Proof. This is a simple consequence of Theorem 1. Indeed, we have $(T_{n, g_n}) \rightarrow \mathbb{T}_{\lambda(\theta)}$ locally. Let t_1 be the triangulation represented on Figure 9. We have on the one hand

$$\mathbb{P}(t_1 \subset T_{n, g_n}) = \frac{\tau_1(n-1, g_n)}{\tau(n, g_n)} = \frac{\tau(n-1, g_n)}{\tau(n, g_n)}$$

by the usual root transformation (see Remark 2). On the other hand, we have

$$\mathbb{P}(t_1 \subset \mathbb{T}_{\lambda(\theta)}) = C_1(\lambda(\theta))\lambda(\theta)^2 = \lambda(\theta).$$

Since $\{t_1 \subset T\}$ is closed and open for d_{loc} , the lemma follows. \square




Figure 9: The triangulation t_1 with perimeter 1 and volume 2.

The idea of the proof of Theorem 3 is then to write

$$\tau(n, g_n) = \frac{\tau(n, g_n)}{\tau((2 + \varepsilon)g_n, g_n)} \times \tau((2 + \varepsilon)g_n, g_n) \quad (29)$$

for some small $\varepsilon > 0$. The first factor can be turned into a telescopic product and estimated by Lemma 26. The trickier part will be to estimate the second. We will prove the following bounds.

Proposition 27. There is a function h with $h(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$ such that

$$e^{o(g)} \left(\frac{6}{e}\right)^{2g} (2g)^{2g} \leq \tau((2 + \varepsilon)g, g) \leq e^{h(\varepsilon)g + o(g)} \left(\frac{6}{e}\right)^{2g} (2g)^{2g}$$

as $g \rightarrow +\infty$. Moreover, if $\varepsilon_g \rightarrow 0$, then

$$\tau((2 + \varepsilon_g)g, g) = e^{o(g)} \left(\frac{6}{e}\right)^{2g} (2g)^{2g}.$$

We delay the proof of Proposition 27, and first explain how to finish the proof of Theorem 3.

Proof of Theorem 3. We first note that $f(1/2) = \frac{6}{e}$, so the case $\theta = 1/2$ is exactly the second part of Proposition 27. We now assume $0 \leq \theta < 1/2$. Let $\varepsilon > 0$ be such that $\theta < \frac{1}{2+\varepsilon}$. We estimate the first factor of (29). We write

$$\frac{1}{n} \log \left(\frac{\tau(n, g_n)}{\tau((2 + \varepsilon)g_n, g_n)} \right) = \frac{1}{n} \sum_{i=(2+\varepsilon)g_n}^n \log \left(\frac{\tau(i, g_n)}{\tau(i-1, g_n)} \right).$$

By Lemma 26, when $\frac{g_n}{i} \rightarrow \frac{1}{t}$, so $\log \frac{\tau(i, g_n)}{\tau(i-1, g_n)} \rightarrow \log \frac{1}{\lambda(1/t)}$. Moreover, by the bounded ratio lemma (Lemma 4), each of the terms is bounded by some constant C_ε . Hence, by dominated convergence, we have

$$\frac{1}{g_n} \log \left(\frac{\tau(n, g_n)}{\tau((2 + \varepsilon)g_n, g_n)} \right) \xrightarrow{n \rightarrow +\infty} \int_{(2+\varepsilon)}^{1/\theta} \log \frac{1}{\lambda(1/t)} dt. \quad (30)$$

On the other hand, if we replace g by g_n with $\frac{g_n}{n} \rightarrow \theta$, then the left-hand side of Proposition 27 becomes $n^{2g_n} \exp\left(\left(2\theta \log \frac{12\theta}{e}\right)n + o(n)\right)$, so Proposition 27 gives

$$\left(2\theta \log \frac{12\theta}{e}\right)n + o(n) \leq \log \frac{\tau((2 + \varepsilon)g_n, g_n)}{n^{2g_n}} \leq \left(2\theta \log \frac{12\theta}{e} + h(\varepsilon)\right)n + o(n),$$

where $h(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$. By combining this with (29) and (30) and finally letting $\varepsilon \rightarrow 0$, we get the result. \square




Figure 10: An example of the cutting operation on a cubic map. The distinguished spanning tree is in red. We picked a planar example but in the general case the "opened" map on the right is unicellular.

We finally prove Proposition 27. The lower bound can be deduced easily from the Goulden–Jackson formula. For the upper bound, we will bound crudely the number of triangulations by the number of tree-rooted triangulations, i.e. triangulations decorated with a spanning tree. These can be counted by adapting classical operations going back to [36] in the planar case.

Proof of Proposition 27. We start with the lower bound. The Goulden–Jackson formula (4) implies the following (crude) inequality:

$$\tau(n, g) \geq (36 + o(1))n^2\tau(n - 2, g - 1),$$

where the o is uniform in g , when $n \rightarrow +\infty$. By an easy induction, we have

$$\tau((2 + \varepsilon)g, g) \geq (36 + o(1))^g \frac{((2 + \varepsilon)g)!}{(\varepsilon g)!} \tau(\varepsilon g, 0) \geq (36 + o(1))^g \frac{((2 + \varepsilon)g)!}{(\varepsilon g)!},$$

and the lower bound follows from the Stirling formula.

For the upper bound, we adapt a classical argument about tree-rooted maps. We denote by $\tilde{\mathcal{T}}((2 + \varepsilon)g, g)$ the set of triangulations of genus g with $2(2 + \varepsilon)g$ faces and a distinguished spanning tree, and by $\tilde{\tau}((2 + \varepsilon)g, g)$ its cardinal. We recall that such triangulations have $\varepsilon g + 2$ vertices, so the spanning tree has $\varepsilon g + 1$ edges.

If $t \in \tilde{\mathcal{T}}((2 + \varepsilon)g, g)$, we consider the dual cubic map of t and "cut in two" the edges that are crossed by the spanning tree, as on Figure 10. The map that we obtain is unicellular (i.e. it has one face), and we denote it by $U(t)$. Moreover, it is *precubic* (i.e. its vertices have only degree 1 or 3), has genus g and $(6 + 4\varepsilon)g + 1$ edges (the number of edges of the original triangulation, plus one for each edge of the spanning tree). The number of precubic unicellular maps with fixed genus and number of edges is computed exactly in [17, Corollary 7] and is equal to

$$\frac{2((6 + 4\varepsilon)g + 1)!}{12^g g! (2\varepsilon g + 2)! ((3 + 2\varepsilon)g)!} = e^{h_U(\varepsilon)g + o(g)} \left(\frac{6}{e}\right)^{2g} (2g)^{2g},$$

where

$$h_U(\varepsilon) = 2\varepsilon \log \frac{6}{e} + (3 + 2\varepsilon) \log \left(1 + \frac{2\varepsilon}{3}\right) \xrightarrow{\varepsilon \rightarrow 0} 0$$

by the Stirling formula. Finally, to go back from $U(t)$ to t , we also need to remember how to match the leaves two by two in the face of $U(t)$ without any crossing. The number of ways to do so is $\text{Catalan}(\varepsilon g + 1) \leq 4^{\varepsilon g}$, so

$$\tau((2 + \varepsilon)g, g) \leq \tilde{\tau}((2 + \varepsilon)g, g) \leq 4^{\varepsilon g} e^{h_U(\varepsilon)g + o(g)} \left(\frac{6}{e}\right)^{2g} (2g)^{2g},$$

which is enough to conclude. The proof for the second part of the proposition is exactly the same, but where we replace ε by $\varepsilon_g \rightarrow 0$. \square

Remark 28. Bounding the number of triangulations by the number of tree-rooted triangulations may seem very crude. The reason why this is sufficient is that the spanning trees have only $\varepsilon g + 1$ edges, so the number of spanning trees of a triangulation can be bounded by $\binom{3(2+\varepsilon)g}{\varepsilon g + 1}$, which is of the form $e^{h(\varepsilon)g + o(g)}$ with $h(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$.

References

- [1] 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.
- [2] L. Addario-Berry and M. Albenque. Convergence of odd-angulations via symmetrization of labeled trees. *arXiv:1904.04786*, 2019.
- [3] D. Aldous and R. Lyons. Processes on unimodular random networks. *Electron. J. Probab.*, 12:no. 54, 1454–1508 (electronic), 2007.
- [4] O. Angel, G. Chapuy, N. Curien, and G. Ray. The local limit of unicellular maps in high genus. *Electron. Commun. Probab.*, 18(86):1–8, 2013.
- [5] O. Angel, A. Nachmias, and G. Ray. Random walks on stochastic hyperbolic half planar triangulations. *Random Structures and Algorithms*, 49(2):213–234, 2016.
- [6] O. Angel and G. Ray. Classification of half planar maps. *Ann. Probab.*, 43(3):1315–1349, 2015.
- [7] O. Angel and O. Schramm. Uniform infinite planar triangulations. *Comm. Math. Phys.*, 241(2-3):191–213, 2003.
- [8] E. A. Bender and E. Canfield. The asymptotic number of rooted maps on a surface. *Journal of Combinatorial Theory, Series A*, 43(2):244 – 257, 1986.
- [9] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Uniform spanning forests. *Ann. Probab.*, 29(1):1–65, 02 2001.
- [10] J. Bettinelli. Geodesics in Brownian surfaces (Brownian maps). *Ann. Inst. Henri Poincaré Probab. Stat.*, 52(2):612–646, 2016.
- [11] T. Budd. The peeling process of infinite Boltzmann planar maps. *Electronic Journal of Combinatorics*, 23, 06 2015.
- [12] T. Budzinski. *Cartes aléatoires hyperboliques*. PhD thesis, Université Paris-Sud, 2018.

- [13] T. Budzinski. The hyperbolic Brownian plane. *Probability Theory and Related Fields*, 171(1):503–541, Jun 2018.
- [14] T. Budzinski. Infinite geodesics in hyperbolic random triangulations. *Ann. Inst. H. Poincaré Probab. Statist.*, 56(2):1129–1161, 05 2020.
- [15] T. Budzinski, N. Curien, and B. Petri. Universality for random surfaces in unconstrained genus. *Electr. J. Combinatorics*, 26(4), 2019.
- [16] S. R. Carrell and G. Chapuy. Simple recurrence formulas to count maps on orientable surfaces. *Journal of Combinatorial Theory, Series A*, 133:58 – 75, 2015.
- [17] G. Chapuy. A new combinatorial identity for unicellular maps, via a direct bijective approach. *Advances in Applied Mathematics*, 47(4):874 – 893, 2011.
- [18] P. Chassaing and B. Durhuus. Local limit of labeled trees and expected volume growth in a random quadrangulation. *Ann. Probab.*, 34(3):879–917, 2006.
- [19] S. Chmutov and B. Pittel. On a surface formed by randomly gluing together polygonal discs. *Advances in Applied Mathematics*, 73:23–42, 2016.
- [20] N. Curien. Planar stochastic hyperbolic triangulations. *Probability Theory and Related Fields*, 165(3):509–540, 2016.
- [21] N. Curien. Peeling random planar maps. *Saint-Flour lecture notes*, 2019.
- [22] N. Curien and J.-F. Le Gall. The Brownian plane. *J. Theoret. Probab.*, 27(4):1249–1291, 2014.
- [23] F. David, A. Kupiainen, R. Rhodes, and V. Vargas. Liouville quantum gravity on the Riemann sphere. *Communications in Mathematical Physics*, 342(3):869–907, Mar 2016.
- [24] F. David, R. Rhodes, and V. Vargas. Liouville quantum gravity on complex tori. *Journal of Mathematical Physics*, 57(2):022302, 2016.
- [25] B. Eynard. *Counting surfaces*, volume 70 of *Progress in Mathematical Physics*. Birkhäuser/Springer, [Cham], 2016. CRM Aisenstadt chair lectures.
- [26] A. Gamburd. Poisson–Dirichlet distribution for random Belyi surfaces. *Ann. Probab.*, 34(5):1827–1848, 2006.
- [27] I. P. Goulden and D. M. Jackson. The KP hierarchy, branched covers, and triangulations. *Adv. Math.*, 219(3):932–951, 2008.
- [28] F. Hausdorff. Summationsmethoden und Momentfolgen I. *Mathematische Zeitschrift*, 9:74–109, 1921.
- [29] M. Kontsevich. Intersection theory on the moduli space of curves and the matrix Airy function. *Comm. Math. Phys.*, 147(1):1–23, 1992.
- [30] M. Krikun. Local structure of random quadrangulations. *arXiv:0512304*.

- [31] M. Krikun. Explicit enumeration of triangulations with multiple boundaries. *Electron. J. Combin.*, 14(1):Research Paper 61, 14 pp. (electronic), 2007.
- [32] J.-F. Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.*, 41:2880–2960, 2013.
- [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] T. Miwa, M. Jimbo, and E. Date. *Solitons*, volume 135 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2000.
- [36] R. Mullin. On the enumeration of tree-rooted maps. *Canad. J. Math.*, 19:174–183, 1967.
- [37] A. Okounkov. Toda equations for Hurwitz numbers. *Math. Res. Lett.*, 7(4):447–453, 2000.
- [38] G. Ray. Geometry and percolation on half planar triangulations. *Electron. J. Probab.*, 19(47):1–28, 2014.
- [39] G. Ray. Large unicellular maps in high genus. *Ann. Inst. H. Poincaré Probab. Statist.*, 51(4):1432–1456, 11 2015.
- [40] 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.
- [41] W. T. Tutte. A census of planar maps. *Journal canadien de mathématiques*, 15(0):249–271, jan 1963.
- [42] D. B. Wilson. Generating random spanning trees more quickly than the cover time. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing, STOC '96*, pages 296–303, New York, NY, USA, 1996. ACM.
- [43] E. Witten. Two-dimensional gravity and intersection theory on moduli space. In *Surveys in differential geometry (Cambridge, MA, 1990)*, pages 243–310. Lehigh Univ., Bethlehem, PA, 1991.