Séminaire Lotharingien de Combinatoire **78B** (2017)
Article #79, 12 pp.

Proceedings of the 29th Conference on Formal Power Series and Algebraic Combinatorics (London)

On tessellations of random maps and the t_g-recurrence

Guillaume Chapuy*

IRIF (ex-LIAFA), UMR CNRS 8243, Université Paris-Diderot, France

Abstract. The number of n-edge embedded graphs (rooted maps) on the g-torus grows as $t_g n^{5(g-1)/2} 12^n$ when n tends to infinity. The constants t_g can be computed via the non-linear "t_g-recurrence", strongly related to the KP hierarchy and the double scaling limit of the one-matrix model. The combinatorial meaning of this simple recurrence is still mysterious, and the purpose of this work is to point out an interpretation via random maps on surfaces. Namely, we show that the t_g-recurrence is equivalent, via combinatorial bijections, to the fact that $\mathbf{E}X_g^2 = \frac{1}{3}$ for any $g \geq 0$, where $X_g, 1 - X_g$ are the masses of the nearest-neighbour cells surrounding two randomly chosen points in a Brownian map of genus g. This raises the question (that we leave open) of giving an independent probabilistic or combinatorial derivation of this second moment, which would lead to a fully concrete proof of the t_g-recurrence. In fact, we conjecture that for any $g \geq 0$ and $k \geq 2$, the masses of the k nearest-neighbour cells induced by k uniform points in the genus g Brownian map form a uniform k-division of the unit interval. We leave this question open even for $(g, k) = (0, 2)$.

Résumé. Le nombre de graphes plongés à n arêtes (cartes enracinées) sur le g-tore croît comme $t_g n^{5(g-1)/2} 12^n$ quand $n \to \infty$. Les constantes t_g peuvent être calculées grâce àă la récurrence-t_g, non linéaire, fortement liée à la hiérarchie KP et à la double limit d'échelle du modèle à une matrice. Le sens combinatoire de cette récurrence simple demeure mystérieux, et le but de ce travail est d'en fournir une interprétation via les cartes aléatoires. À savoir, nous montrons que la récurrence-t_g est équivalente, via des bijections combinatoires, au fait que $\mathbf{E}X_g^2 = \frac{1}{3}$ pour tout $g \geq 0$, où $X_g, 1 - X_g$ sont les masses des cellules des plus proches voisins entourant deux points choisis au hasard dans une carte brownienne de genre g. Cela soulève la question (que nous laissons ouverte) de donner une dérivation probabiliste ou combinatoire indépendante de ce second moment, ce qui conduirait À une démonstration concrète complète de la récurrence-t_g. En fait, on conjecture que pour tout $g \geq 0$ et $k \geq 2$, les masses des k cellules des plus proches voisins induites par k points uniformes dans la carte brownienne de genre g forment une k-division uniforme de l'intervale $[0, 1]$. On laisse cette question ouverte même pour $(g, k) = (0, 2)$.

Keywords: Maps on surfaces, bijections, asymptotic enumeration, KP hierarchy

*Support from the European Research Council, grant ERC-2016-STG 716083 "CombiTop", from the City of Paris, grant "Émergences, Combinatoire à Paris", and from the grant ANR 12-JS02-001-01 "Cartaplus".

1 Introduction and results

In this note a *map* is a graph embedded without edge crossings on a closed oriented surface, in such a way that the connected components of the complement of the graph, called *faces*, are all homeomorphic to a disk. Loops and multiple edges are allowed, and maps are considered up to oriented homeomorphisms. A map is *rooted* if one edge is distinguished and oriented. The number $m_g(n)$ of rooted maps with n edges on the surface of genus g satisfies, for fixed $g \geq 0$ and $n \to \infty$:

$$m_g(n) \sim t_g n^{\frac{5(g-1)}{2}} 12^n, \text{ for } t_g > 0. \tag{1.1}$$

In genus 0, this result follows from the exact formula $m_0(n) = \frac{2 \cdot 3^n}{(n+2)(n+1)} \binom{2n}{n}$ due to Tutte [19]. In higher genus, it was proved by Bender and Canfield in [1]. A direct combinatorial interpretation of the genus 0 formula was given by Cori and Vauquelin [8] and much simplified by Schaeffer [18, 7]. A combinatorial interpretation of (1.1) was given in [6] using the Marcus-Schaeffer bijection [14] and further developed in [5].

None of the methods just mentioned enable us to say much about the sequence of constants $(t_g)_{g \geq 0}$ that appear in (1.1), and indeed these references give explicit values only for very small values of g. There is however a remarkable recurrence formula to compute these numbers, that we call the *t_g-recurrence*. It is better expressed in terms of the numbers $\tau_g = 2^{5g-2} \Gamma \left(\frac{5g-1}{2} \right) t_g$ and is given by:

$$\tau_{g+1} = \frac{(5g+1)(5g-1)}{3} \tau_g + \frac{1}{2} \sum_{g_1=1}^{g} \tau_{g_1} \tau_{g+1-g_1}, \quad g \geq 0, \tag{1.2}$$

which enables us to compute these numbers easily starting from $\tau_0 = -1$. This result was first stated in mathematical physics in relation with the *double scaling limit* of the one-matrix model, and obtained via a non-rigorous scaling of expressions involving orthogonal polynomials (we refer to [12, p201] for historical references). A more algebraic approach is based on the fact that the partition function of maps on surfaces, with infinitely many parameters marking vertex degrees, is a tau-function of the KP hierarchy. Going from the KP hierarchy to the recurrence (1.2) (or to an equivalent Painlevé-I ODE for an associated generating function) relies on a trick of elimination of variables that can be performed in different ways and whose generality is, as far as we know, yet to be fully understood (for the case of triangulations see [11, Appendix B.] or [10, 2] and for general maps see [4]).

The main observation of this note is to relate the recurrence (1.2) to another side of the story, namely the study of random maps and their scaling limits. We refer to the introductions of the papers [15, 16, 13] for an introduction to the topic and for references. To state our main observation we first need a few more definitions. A *quadrangulation* is a map in which each face contains exactly four corners, *i.e.* is bordered by exactly four edge-sides. It is *bipartite* if its vertices can be colored in black and white in such a way that there is no monochromatic edge. For each $n, g \geq 0$, there is a classical bijection, due to Tutte, between rooted maps of genus g with n edges and rooted bipartite quadrangulations of genus g with n faces.

For $n, g \geq 0$, we let $\mathcal{Q}_n^{(g)}$ be the set of rooted bipartite quadrangulations of genus g with n faces (with the convention that there is a single quadrangulation with 0 faces, which has genus 0, no edges, and two vertices). We let $\mathbf{q}_n^{(g)} \in_u \mathcal{Q}_n^{(g)}$ be a bipartite quadrangulation of genus g with n faces chosen uniformly at random. We equip the vertex set of $\mathbf{q}_n^{(g)}$ with the graph distance, noted $\mathbf{d}_n$, and with the uniform measure, noted μ_n. This makes $\mathbf{q}_n^{(g)} \equiv (\mathbf{q}_n^{(g)}, \mathbf{d}_n, \mu_n)$ into a compact measured metric space. The set of (isometry classes of) such spaces is equipped with the Gromov-Hausdorff-Prokhorov (GHP) topology as in [15, Sec. 6]. A *Brownian map of genus g* is a random compact measured metric space $(\mathbf{q}_\infty, d_\infty, \mu_\infty)$ that is such that:

$$(\mathbf{q}_n^{(g)}, \tfrac{1}{n^{1/4}}\mathbf{d}_n, \mu_n) \longrightarrow (\mathbf{q}_\infty^{(g)}, d_\infty, \mu_\infty),$$

in distribution along some subsequence for the GHP topology. The existence of Brownian maps of genus g was proved in [15], and their uniqueness for each $g \geq 1$ has been announced by Bettinelli and Miermont [3] (in genus 0 the uniqueness is an important result proved independently by Miermont [16] and Le Gall [13]). However the uniqueness of the limit is not needed for our discussion since we will prove the convergence of all the observables we are interested in.

Theorem 1 (First observation, obtained by combinatorial means)**.** *For $g \geq 0$, let $(\mathbf{q}_\infty^{(g)}, d_\infty, \mu_\infty)$ be a Brownian map of genus g. Let $\mathbf{v}_1, \mathbf{v}_2 \in \mathbf{q}_\infty^{(g)}$ be chosen independently according to the probability measure μ_∞, and let $\mathbf{X}_g, 1 - \mathbf{X}_g$ be the masses of the corresponding cells in the nearest neighbour tessellation of $\mathbf{q}_\infty^{(g)}$ induced by $\mathbf{v}_1$ and $\mathbf{v}_2$, that is to say:*

$$\mathbf{X}_g := \mu_\infty\Big(\big\{x \in \mathbf{q}_\infty^{(g)},\ d_\infty(x, \mathbf{v}_1) < d_\infty(x, \mathbf{v}_2)\big\}\Big).$$

Then the sequence of numbers $\tau_g = 2^{5g-2}\Gamma\left(\frac{5g-1}{2}\right) t_g$ satisfies:

$$\tau_{g+1} = 2(5g+1)(5g-1)\tau_g \cdot \mathbf{E}[\mathbf{X}_g(1-\mathbf{X}_g)] + \frac{1}{2}\sum_{g_1=1}^{g} \tau_{g_1}\tau_{g+1-g_1}, \quad g \geq 0.$$

From (1.2) we immediately deduce:

Theorem 2 (Second observation, by comparing Theorem 1 with the t_g-recurrence (1.2))**.** *For any $g \geq 0$, the random variable $\mathbf{X}_g$ satisfies*

$$\mathbf{E}[\mathbf{X}_g(1-\mathbf{X}_g)] = \frac{1}{6}, \quad \textit{or equivalently } \mathbf{E}\mathbf{X}_g^2 = \frac{1}{3}.$$

The reader may be surprised that $\mathbf{E}\mathbf{X}_g^2$ does not depend on $g \geq 0$: indeed, although it is natural to expect that *local* statistics of Brownian maps do not depend on the genus, the nearest-neighbour tessellation depends *globally* of the metric space $\mathbf{q}_\infty^{(g)}$, which *is* genus dependent. This unexpected property suggests that there exists a simple probabilistic or combinatorial interpretation of this mysterious fact, based on a symmetry of the Brownian map, but we have not been able to find it. We emphasize that, via Theorem 1, such an interpretation would provide a proof of the t_g-recurrence independent of orthogonal polynomials, matrix models or integrable hierarchies.

It is natural to ask if other moments of the variables $\mathbf{X}_g$ or related random variables are computable, and in what way they depend on the genus. Unfortunately we will not go very far in this direction. Let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k$ be $k \geq 2$ points in $\mathbf{q}_\infty^{(g)}$ chosen independently at random according to the Lebesgue measure μ_∞. Let $(\mathbf{Y}_g^{(i:k)})_{1 \leq i \leq k}$ be the masses of the k-nearest-neighbour cells induced by the $\mathbf{v}_i$'s, *i.e.* for $i \in [1..k]$ let

$$\mathbf{Y}_g^{(i:k)} := \mu_\infty\Big(\big\{x \in \mathbf{q}_\infty^{(g)},\ \forall j \in [1..k] \setminus \{i\},\ d_\infty(x, \mathbf{v}_i) < d_\infty(x, \mathbf{v}_j)\big\}\Big).$$

We note that $\mathbf{X}_g = \mathbf{Y}_g^{(1:2)}$, but we prefer to keep the lighter notation $\mathbf{X}_g$ for $\mathbf{Y}_g^{(1:2)}$ throughout this note. The following result is similar to, and as mysterious as Theorem 2:

Theorem 3 (A similar simple formula for the case of three points)**.** *For $g \geq 0$, the masses $\mathbf{Y}_g^{(1:3)}, \mathbf{Y}_g^{(2:3)}, \mathbf{Y}_g^{(3:3)}$ of the Voronoï cells induced by three independent Lebesgue distributed points in the Brownian map of genus g satisfy, for $g \geq 0$:*

$$\mathbf{E}[\mathbf{Y}_g^{(1:3)}\mathbf{Y}_g^{(2:3)}\mathbf{Y}_g^{(3:3)}] = \frac{1}{60}.$$

As we will see, the fact that this moment is computable reflects the existence of a combinatorial device known as the "trisection lemma" [5]. The fact that it does not depend on the genus, and that it coincides[2] with the corresponding moment for a uniform three-division of the interval $[0,1]$, is as mysterious as for the previous result (or even more, since the computations leading to Theorem 3 are quite delicate).

We won't prove anything on higher moments or other values of k since we lack the tools to study them. However, numerical simulations suggest the remarkable property:

Conjecture 4. *For $k \geq 2, g \geq 0$, let $\mathbf{q}_\infty^{(g)} \equiv (\mathbf{q}_\infty^{(g)}, d_\infty, \mu_\infty)$ be a genus g Brownian map and let $\mathbf{v}_1, \dots, \mathbf{v}_k$ be chosen according to $\mu_\infty^{\otimes k}$. Then the random vector $(\mathbf{Y}_g^{(1:k)}, \mathbf{Y}_g^{(2:k)}, \dots, \mathbf{Y}_g^{(k:k)})$ has the same law as the subdivision of the unit interval induced by $k-1$ independent uniform variables. In particular, for any $g \geq 0$, $\mathbf{X}_g = \mathbf{Y}_g^{(1:2)}$ is uniform on $[0,1]$.*

[2] If U_1, U_2 are two independent uniforms on $[0,1]$ and I_1, I_2, I_3 are the lengths of the three intervals they define, then $\mathbf{E}(I_1 I_2 I_3)$ is the probability that five independent uniforms U_1, U_2, V_1, V_2, V_3 are ordered as $V_1 < U_1 \wedge U_2 < V_2 < U_1 \vee U_2 < V_3$, which is clearly equal to $\frac{2}{5!} = \frac{1}{60}$.

To conclude this introduction, we emphasize that our main observation relates the moment $\mathbf{E}\mathbf{X}_g^2$ to the g-th step of the t_g-recurrence. In particular, the fact that $\mathbf{E}\mathbf{X}_0^2 = 1/3$ for the *genus* 0 Brownian map is only "equivalent" to the computation of the genus 1 constant t_1, which can be performed by hand in several ways (and similarly, our proof of Theorem 3 for $g = 0$ relies only on the value of the constants t_1 and t_2). However, proving Conjecture 4 even for $(g, k) = (0, 2)$ would be interesting in itself. Readers familiar with Miermont's bijection [15] may try to approach this problem by exact counting of well-labelled 2-face maps (we have failed trying to do so). One could also hope that in the future purely probabilistic methods (for example using the QLE viewpoint on the Brownian map [17]) will enable to determine the full law of $\mathbf{X}_0$ or even the law of the vector $(\mathbf{Y}_0^{(i:k)})_{1\le i\le k}$ for each k. In an opposite direction, we recall that the t_g-recurrence is only a "shadow" of the fact that the generating functions of maps satisfy a set of infinitely many partial differential equations called the KP hierarchy. It is natural to expect that other joint moments of the variables $\mathbf{Y}_g^{(i:k)}$, apart from the two cases we have been able to track, are related to these equations. This may lead to a way, based on integrable hierarchies, of approaching Conjecture 4.

On this extended abstract. This version is an extended abstract of a full paper available at arXiv:1603.07714, to which we refer for full proofs. Here we skip most of the technical probabilistic details related to convergence (replaced here by heuristic considerations) and focus on the combinatorial techniques and on the path of proof which consists mostly in a conjoint use of generating functions and combinatorial bijections.

2 Proof sketch of our main observation (Theorem 1)

2.1 Preliminaries

For $g \ge 0$ we let $Q_g(z)$ be the generating function of rooted bipartite quadrangulations of genus g by the number of faces, and we let $Q_g^\bullet(z)$ be the g.f. of the same objects where an additional vertex is pointed. We let $m_g(n) = [z^n]Q_g(z)$ and we use the same notation with $^\bullet$. In what follows the notation $a(n) \sim b(n)$ means (as usual) that $a(n)/b(n) \to 1$ when n tends to infinity, while the notation $F(z) \sim G(z)$ means that both F and G are algebraic functions of radius of convergence $\frac{1}{12}$, both have a unique dominant singularity at $z = \frac{1}{12}$, and we have $F(z) = G(z)(1 + o(1))$ when $z \to \frac{1}{12}$ uniformly in a neighbourhood of $z = \frac{1}{12}$ slit along the line $[\frac{1}{12}, \infty)$.

From [1] (see also [6] for purely combinatorial proofs) we have for fixed $g \ge 0$:

$$m_g(n) \sim t_g n^{\frac{5g-5}{2}} 12^n \quad , \quad m_g^\bullet(n) = (n + 2 - 2g) m_g(n) \sim t_g n^{\frac{5g-3}{2}} 12^n$$

$$Q_g^\bullet(z) \sim \Gamma(\tfrac{5g-1}{2}) t_g (1 - 12z)^{\frac{1-5g}{2}} = 2^{2-5g} \tau_g (1 - 12z)^{\frac{1-5g}{2}}. \tag{2.1}$$

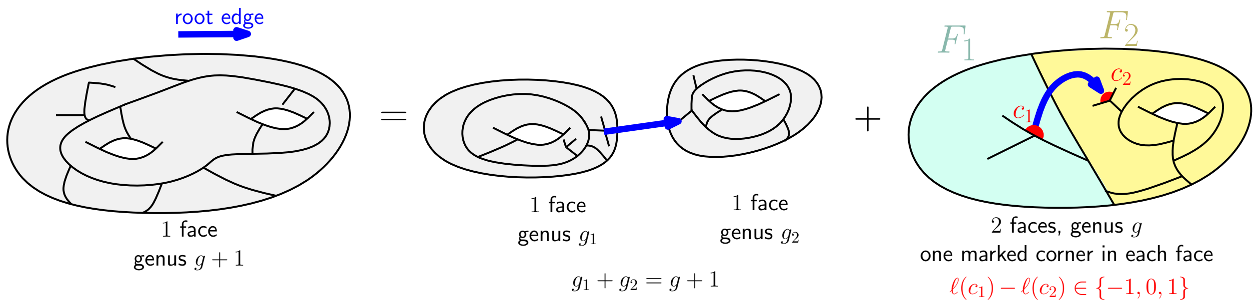

Figure 1: Illustration of the decomposition leading to (2.3)

A *labelled map* of genus g is a rooted map M of genus g equipped with a function $\ell : V(M) \to \mathbb{Z}$ such that for any edge (u,v) of M one has $\ell(u)-\ell(v) \in \{-1,0,1\}$. We consider these objects up to global translation of labels (one can fix a translation class by fixing the label of some vertex to 0). A *labelled one-face map (l.1.f.m.)* is a labelled map having only one face. We let $\mathcal{L}_n^{(g)}$ be the set of (rooted) l.1.f.m. of genus g with n edges.

The Marcus-Schaeffer bijection ([14], see also [6]) is an explicit bijection from $\mathcal{Q}_n^{(g)\bullet}$ to $\{\uparrow,\downarrow\}\times\mathcal{L}_n^{(g)}$, where $\mathcal{Q}_n^{(g)\bullet}$ is the set of rooted bipartite quadrangulations of genus g and n faces equipped with a pointed vertex. It follows that $Q_g^\bullet(z) = 2L_g(z)$ where $L_g(z)$ is the generating function of rooted l.1.f.m. of genus g by the number of edges. Moreover, in genus 0, rooted one-face maps are trees, and a standard root-edge decomposition leads to the quadratic equation $L_0(z) = 1 + 3zL_0(z)^2$, from which we get the explicit formula:

$$1 - 6zL_0(z) = \sqrt{1-12z}. \tag{2.2}$$

2.2 Decomposition equation, Miermont's bijection, proof of Theorem 1

We now come to the substance of this work, which is simply to try to write an equation for the generating function of l.1.f.m. by root-edge decomposition, and see what happens. We fix $g \geq 0$, and we consider a l.1.f.m. M of genus $g+1$. If we remove the root edge of this map, two things can happen (see Figure 1):

(i) we disconnect the map into two l.1.f.m. M_1 and M_2 whose genera sum up to $g+1$;

(ii) we do not disconnect the map; in this case we are left with a map M' of genus g with two faces. Each face of M' carries a distinguished corner, and the labels of these two corners differ by -1, 0, or 1.

Translating this operation into an equation for generating functions we obtain

$$L_{g+1}(z) = 3z \sum_{\substack{g_1+g_2=g+1 \\ g_1,g_2\geq 0}} L_{g_1}(z)L_{g_2}(z) + zA_g(z) \tag{2.3}$$

where:
- in the first term the factor $3z$ takes into account the choice of the increment of label along the root-edge in $\{-1,0,1\}$;

- $A_g(z)$ is the generating function by the number of edges, of unrooted labelled two-face maps, with faces numbered F_1, F_2, such that the face F_i contains a marked corner c_i for $i = 1..2$, and that $|\ell(c_1) - \ell(c_2)| \leq 1$.

Objects counted by $A_g(z)$ can be related to quadrangulations thanks to Miermont's bijection [15]. This bijection is a generalization of the Marcus-Schaeffer bijection where the l.1.f.m is replaced by a labelled map having an arbitrary number, say K, of faces. We will apply it for $K = 2$. In the following discussion, where we assume some familiarity with Miermont's bijection, we will show how to arrive informally at Lemma 5 below, and why this implies Theorem 1. As we said we will focus on key combinatorial features and stay at the intuitive level for technical probabilistic facts.

Let us consider an object counted by $[z^n]A_g(z)$. Let us fix the translation class of the labels by saying that the minimum label in face F_1 is zero, and let us call δ the minimum label in face F_2. Let $i_1 \geq 0$ and $i_2 \geq \delta$ be the labels of the two marked corners c_1 and c_2, respectively, and recall that $i_1 - i_2 \in \{-1, 0, 1\}$. Applying Miermont's bijection [15] to this object, we construct a bipartite quadrangulation Q of genus g by adding a new vertex s_1, s_2 inside each face F_1, F_2, and applying a certain closure operation. At the end of the construction, we obtain a quadrangulation such that $d(s_1, s_2) + \delta$ is even. Moreover, the two corners c_1 and c_2 of the original two-face map are naturally associated to two edges e_1 and e_2 of the quadrangulation, and the construction is such that if m_i is the endpoint of e_i closer from s_i in Q, for $i \in \{1, 2\}$ one has:

$$d(s_1, m_1) = i_1 \ \ , \ \ d(s_2, m_2) = i_2 - \delta \ \ , \ \ d(s_2, m_1) \geq i_1 - \delta \ \ , \ \ d(s_1, m_2) \geq i_2.$$

These constraints can simply be rewritten as:

$$d(s_1, m_1) \leq d(s_1, m_2) - \epsilon \ \ , \ \ d(s_2, m_2) \leq d(s_2, m_1) - \epsilon, \tag{2.4}$$

where $\epsilon = i_2 - i_1$ is such that $|\epsilon| \leq 1$. Loosely speaking, the properties in (2.4) say that, *up to an error at most 1*, s_i is (weakly) closer to m_i than to m_{3-i} for $i = 1..2$. Unfortunately these constraints do not entirely characterize these objects (see [15, Sec 2.2]) but they do, in some sense, asymptotically[3]. Thinking heuristically for a moment, we can expect that the analogue in the continuum limit of these discrete configurations is a Brownian map with four marked points $(m_1^\infty, m_2^\infty, s_1^\infty, s_2^\infty)$ such that if we subdivide the space in two nearest-neighbour cells induced by m_1^∞ and m_2^∞, the point s_i^∞ belongs to the nearest-neighbour cell induced by m_i^∞ for each $i = 1..2$. Up to technical details that we will carry out in the next section, this leads us quite naturally to the following conclusion:

Lemma 5. *The coefficient $[z^n]A_g(z)$ is such that, as n tends to ∞, with the notation of Theorem 1:*

$$\frac{[z^n]A_g(z)}{3/2 \cdot n^3 m_g(n)} \longrightarrow \mathbf{E}[\mathbf{X}_g(1 - \mathbf{X}_g)] \tag{2.5}$$

[3] Roughly speaking the only ambiguity comes from vertices where an equality is reached in (2.4), but these are in negligible proportion with high probability.

Remark 1. The reader can understand heuristically the meaning of the denominator $3/2 \cdot n^3 m_g(n)$ as follows. The tuple (Q, s_1, s_2, e_1, e_2) is a quadrangulation with two vertices and two marked edges. We can use e_1 as the root-edge of Q, and orient it by deciding that its source is at even distance from s_1. We can choose the "error" ϵ freely in $\{-1, 0, 1\}$ (since asymptotically we do not expect this error to play any role), and set $i_1 := d(s_1, m_1)$ and $\delta := i_1 + \epsilon - d(s_2, m_2)$. Since Miermont's bijection requires that $d(s_1, s_2) + \delta$ is even, we are left with a rooted quadrangulation with one marked edge e_2, and two marked vertices (s_1, s_2) subject to *two* parity constraints (that s_1 is at even distance from the root, and that $d(s_1, s_2) + \delta$ is even). Since a quadrangulation with n faces has $2n$ edges and $n + 2 - 2g$ vertices, and since it is natural to expect each parity constraint to contribute an asymptotic factor $\frac{1}{2}$, the total number of "base configurations" we obtain is $\sim 3 \times (2n)n^2/4 \cdot m_g(n)$, hence the denominator in (2.5).

We now conclude the proof of Theorem 1. The main idea is to express in two ways the "unknown" series $A_g(z)$ (at leading order). First, rewrite the decomposition (2.3) as:

$$(1 - 6zL_0(z))L_{g+1}(z) - 3z \sum_{\substack{g_1+g_2=g+1,\\ g_1,g_2>0}} L_{g_1}(z)L_{g_2}(z) = zA_g(z), \tag{2.6}$$

which expresses the generating function $L_{g+1}(z)$ in terms of the lower genus functions $L_i(z)$ for $i = 1..g$, and of the "unknown" quantity $A_g(z)$. We recall that $Q_g^\bullet(z) = 2L_g(z)$ and (2.1), from which we observe that each term in the left side of (2.6) has a dominant singularity at $z = \frac{1}{12}$ with the same order of magnitude. More precisely, for the first term, using (2.2), we obtain $(1 - 6zL_0(z))L_{g+1}(z) \sim 2^{1-5(g+1)}\tau_{g+1}(1 - 12z)^{1-\frac{5}{2}(g+1)}$. For product terms we have $L_{g_1}(z)L_{g_2}(z) \sim 2^{2-5(g+1)}\tau_{g_1}\tau_{g_2}(1 - 12z)^{1-\frac{5}{2}(g_1+g_2)}$. It follows, using standard transfer theorems for algebraic functions [9] that when n tends to infinity:

$$[z^{n-1}]A_g(z) \sim 12^n n^{\frac{5g+1}{2}} 2^{1-5(g+1)} \Gamma\left(\frac{5g+3}{2}\right)^{-1} \Big(\tau_{g+1} - \frac{1}{2} \sum_{\substack{g_1+g_2=g+1\\ g_1,g_2>0}} \tau_{g_1}\tau_{g_2}\Big). \tag{2.7}$$

But Lemma 5 gives another expansion of the "unknown" coefficient $[z^{n-1}]A_g(z)$, namely:

$$[z^{n-1}]A_g(z) = \mathbf{E}\mathbf{X}_g(1 - \mathbf{X}_g) \cdot 3/2 \cdot n^3 m_g(n-1) \sim \mathbf{E}\mathbf{X}_g(1 - \mathbf{X}_g) \cdot 3 \cdot 2^{1-5g}\Gamma(\tfrac{5g-1}{2})^{-1}\tau_g n^{\frac{5g+1}{2}} 12^{n-1}.$$

Theorem 1 follows by comparing the last two expansions of the "unknown" quantity $[z^{n-1}]A_g(z)$.

3 Three marked points (proof of Theorem 3)

We will now use our remaining space to sketch the proof of Theorem 3. We will be even quicker than in the previous section only giving an idea of why the third moment of Theorem 3 shows up in the argument. We refer again to the arxiv version for full details.

We first need some definitions from [6, 5]. If L is a one-face map, its *skeleton* is the map obtained by removing all vertices of degree 1 in L, and continuing to do so recursively until only vertices of degree at least 2 remain. Vertices of a one-face map that are vertices of degree at least 3 of its skeleton are called *nodes*. A node v that has degree k in the skeleton is called a k-node (note that its degree as a vertex in the one-face map can be larger than k). A one-face map is *dominant* if all vertices of its skeleton have degree at most 3, *i.e.* if all its nodes are 3-nodes. It is proved in [6] that for fixed g, as n tends to infinity, a proportion at least $1 - O(n^{-1/4})$ of l.1.f.m. of genus g with n edges are dominant. By Euler's formula, a dominant one-face map has $4g - 2$ nodes.

Following [5], we introduce the operation of *opening*. If L is a one-face map and v is a 3-node of L, the *opening* of v is the operation that consists in replacing v by three new vertices, each linked to one edge of the skeleton, and distributing the three (possibly empty) subtrees attached to v among these new vertices as in the following figure:

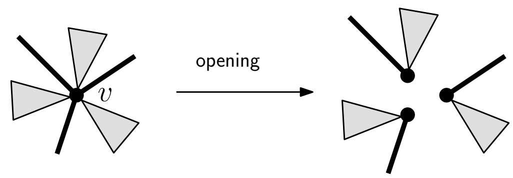

Following [5], we distinguish two types of 3-nodes in a one-face map: *intertwined nodes*, which are such that their opening results in a one-face map of genus $g - 1$ with three marked vertices; and *non-intertwined nodes*, which are such that their opening results in a map of genus $g - 2$ with three faces, and one marked vertex inside each face (here the map can be disconnected, and its genus and number of faces are defined additively on connected components). The *trisection lemma* [5, Lemma 5], which is the key result underlying this section, asserts that any dominant map of genus $g \geq 1$ has exactly $2g$ intertwined nodes, hence $2g - 2$ non-intertwined ones.

It follows that the number $K_{g+2}(n)$ of l.1.f.m. of genus $g + 2$ with n edges whose root edge is a skeleton-edge leaving a *non-intertwined* 3-node satisfies:

$$K_{g+2}(n) \sim \frac{6(g+1)}{2n}[z^n]L_{g+2}(z). \tag{3.1}$$

Indeed, the first-order contribution is given by dominant l.1.f.m., and in a dominant l.1.f.m. of genus $g + 2$ we can choose $3(2(g + 2) - 2) = 6(g + 1)$ edges outgoing from a non-intertwined node as a new root edge, but we obtain each map $2n$ times in this way (since maps counted by $L_{g+2}(z)$ are already rooted at one of their $2n$ oriented edges).

We are now going to obtain another expression for the number $K_{g+2}(n)$ by performing a combinatorial decomposition. Comparing the two will lead us to Theorem 3. Let L be a dominant l.1.f.m of genus $g + 2$ whose root edge is a skeleton-edge leaving a non-intertwined 3-node v. We distinguish three cases, according to what happens when we perform the opening of the node v (see Figure 2):

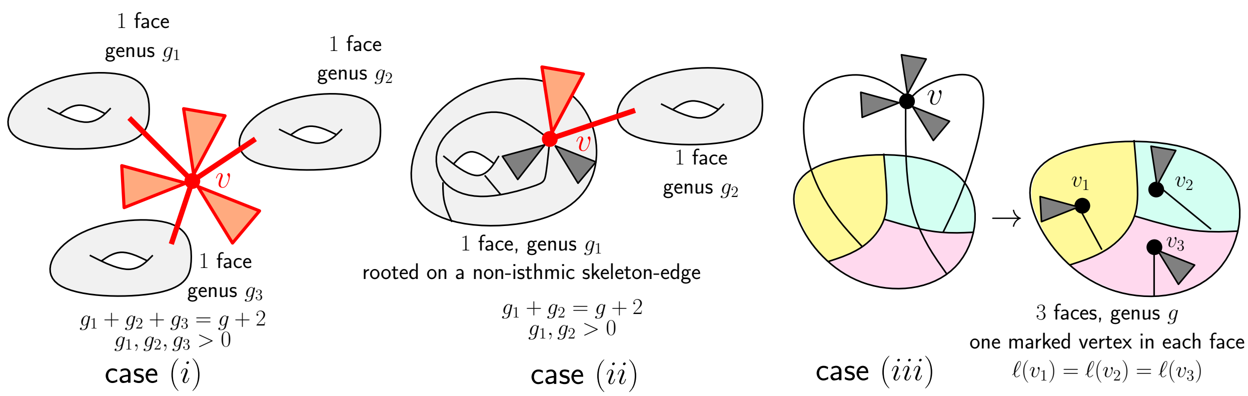

Figure 2: The three cases for a one-face map of genus $g + 2$ rooted at skeleton edge leaving a non-intertwined 3-node v.

(i) we disconnect the map into three components;
(ii) we disconnect the map into two components;
(iii) we do not disconnect the map.

It should be no surprise to our experienced reader that the generating functions for the first two cases (which we denote by $C^{(i)}_{g+2}(z)$ and $C^{(ii)}_{g+2}(z)$) can be expressed in terms of generating functions of maps rooted in several ways. Filling in the details, one easily sees that the leading-order contribution for the sum of the first two cases is:

$$C^{(i)}_{g+2}(z) + C^{(ii)}_{g+2}(z) \sim 2^{-5(g+2)+1} \left(3 \cdot \sum_{\substack{g_1+g_2\\=g+2}} \tau_{g_1}\tau_{g_2} - \sum_{\substack{g_1+g_2+g_3\\=g+2}} \tau_{g_1}\tau_{g_2}\tau_{g_3} \right) \cdot (1-12z)^{\frac{3}{2}-\frac{5}{2}(g+2)} \tag{3.2}$$

where the sums are taken over positive indices (*i.e.* $g_1, g_2, g_3 > 0$). Observe the *cubic* convolution (triple sum), that naturally reflects the fact that the removal of the marked node in case (i) disconnects the map into three components.

We now study case (iii) and for this, we apply again the Miermont's bijection. If we disconnect the three endpoints belonging to the skeleton and the root vertex in a map from case (iii), we obtain a labelled map of genus g with three faces, with one marked vertex inside each face, subject to the constraint that those three vertices have the same label (see Figure 2-Right). Miermont's bijection transforms this object into a bipartite quadrangulation of genus g with *six* marked vertices $(s_1, s_2, s_3, v_1, v_2, v_3)$, such that for $i = 1..3$ the source v_i is closer from the vertex s_i than from the two other vertices s_j (to see this, write precisely the inequalities analogue to (2.4) in previous section). As in the previous section, up to subdominating cases, this property asymptotically characterizes those configurations, and we get:

Lemma 6. *The number $c_{g+2}^{(iii)}(n)$ of configurations in case (iii) satisfies:*

$$\frac{c_{g+2}^{(iii)}(n)}{n^5/4 \cdot 2^{-2} m_g(n)} \sim \mathbf{E}[\mathbf{Y}_g^{(1:3)} \mathbf{Y}_g^{(2:3)} \mathbf{Y}_g^{(3:3)}].$$

The reader can understand heuristically the denominator in the previous expression as follows. The factor $n^5/4$ comes from the fact that we have $\sim n^6$ ways to mark 6 vertices (among $n + 2 - 2g$) but the quadrangulation is unrooted so we divide by $4n$. The factor 2^{-2} corresponds to the fact that we have two parity constraints relating the distances of the six points together (these constraints enable us to choose the delays in such a way that the target vertices get the same label while respecting the parity constraints on delays required by Miermont's bijection). Now that we have seen how to make the quantity $\mathbf{E}[\mathbf{Y}_g^{(1:3)} \mathbf{Y}_g^{(2:3)} \mathbf{Y}_g^{(3:3)}]$ appear, let us sketch the end of the proof.

Sketch of the path to the proof of Theorem 3. Summing contributions given by the last lemma and by (3.2) enables us to obtain an expression of the leading order of the sum of the three cases (i), (ii), (iii), therefore to obtain an alternative expression for the quantity (3.1), at leading order. Comparing the two, we therefore obtain an expression of $\mathbf{E}[\mathbf{Y}_g^{(1:3)} \mathbf{Y}_g^{(2:3)} \mathbf{Y}_g^{(3:3)}]$ as a *trilinear* expression in the $(t_h)_{h \le g+2}$. So far, we have not used the t_g-recurrence. We will do it now, and a miracle happens. By doing a suitable bootstrap of the t_g-recurrence, we can obtain a (complicated) trilinear relation satisfied by the numbers $(t_g)_{g \ge 0}$. It turns out that this trilinear recurrence (which is equivalent to the t_g-recurrence) enables us to simplify considerably the obtained expression. Indeed, as was the case in the previous section for the case of two points, the algebra works perfectly and after simplifications, one obtains that $\mathbf{E}[\mathbf{Y}_g^{(1:3)} \mathbf{Y}_g^{(2:3)} \mathbf{Y}_g^{(3:3)}]$ is in fact equal to $1/60$. We do not have any explanation for this remarkable simplification, and refer the reader to the full version for details of the tedious computations.

Conclusion

We could be disappointed by the seemingly "magical" simplification that occurs at the very end of the proof. Instead, we believe that the fact that computations match so well is just a shadow of the deep links between combinatorial bijections, random maps, and integrable hierarchies. Our results strongly suggest that a unified understanding of these links is possible, although much remains to be done. We hope that Conjecture 4 can play a stimulating role in this direction.

References

[1] E. A. Bender and E. R. Canfield. "The asymptotic number of rooted maps on a surface". *J. Combin. Theory Ser. A* **43** (1986), pp. 244–257. DOI.

[2] E. A. Bender, Z. Gao, and L. B. Richmond. "The map asymptotics constant t_g". *Electron. J. Combin.* **15** (2008), Art. #R51. URL.

[3] J. Bettinelli and G. Miermont. In preparation.

[4] S. R. Carrell and G. Chapuy. "Simple recurrence formulas to count maps on orientable surfaces". *J. Combin. Theory Ser. A* **133** (2015), pp. 58–75. DOI.

[5] G. Chapuy. "The structure of unicellular maps, and a connection between maps of positive genus and planar labelled trees". *Probab. Theory Related Fields* **147** (2010), pp. 415–447. DOI.

[6] G. Chapuy, M. Marcus, and G. Schaeffer. "A bijection for rooted maps on orientable surfaces". *SIAM J. Discrete Math.* **23** (2009), pp. 1587–1611. DOI.

[7] P. Chassaing and G. Schaeffer. "Random planar lattices and integrated superBrownian excursion". *Probab. Theory Related Fields* **128** (2004), pp. 161–212. DOI.

[8] R. Cori and B. Vauquelin. "Planar maps are well labeled trees". *Canad. J. Math.* **33** (1981), pp. 1023–1042. DOI.

[9] P. Flajolet and R. Sedgewick. *Analytic Combinatorics*. Cambridge University Press, 2009.

[10] I. P. Goulden and D. M. Jackson. "The KP hierarchy, branched covers, and triangulations". *Adv. Math.* **219** (2008), pp. 932–951. DOI.

[11] V. A. Kazakov, I. K. Kostov, and N. Nekrasov. "D-particles, matrix integrals and KP hierarchy". *Nuclear Phys. B* **557** (1999), pp. 413–442. DOI.

[12] S. K. Lando and A. K. Zvonkin. *Graphs on Surfaces and their Applications*. Encyclopaedia of Mathematical Sciences, Vol. 141. Springer-Verlag, 2004.

[13] J.-F. Le Gall. "Uniqueness and universality of the Brownian map". *Ann. Probab.* **41** (2013), pp. 2880–2960. DOI.

[14] M. Marcus and G. Schaeffer. "Une bijection simple pour les cartes orientables". Manuscript. 2001.

[15] G. Miermont. "Tessellations of random maps of arbitrary genus". *Ann. Sci. Éc. Norm. Supér. (4)* **42** (2009), pp. 725–781. DOI.

[16] G. Miermont. "The Brownian map is the scaling limit of uniform random plane quadrangulations". *Acta Math.* **210** (2013), pp. 319–401. DOI.

[17] J. Miller and S. Sheffield. "Liouville quantum gravity and the Brownian map I: The QLE(8/3,0) metric". 2015. arXiv:1507.00719.

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

[19] W. T. Tutte. "A census of planar maps". *Canad. J. Math.* **15** (1963), pp. 249–271. DOI.