

Counting unicellular maps on non-orientable surfaces

Olivier Bernardi^{1†} and Guillaume Chapuy^{2‡}

¹Department of Mathematics, MIT, 77 Massachusetts Avenue, Cambridge MA 02139, USA

²Department of Mathematics, Simon Fraser University, Burnaby, B.C. V5A 1S6, Canada

Abstract. A unicellular map is the embedding of a connected graph in a surface in such a way that the complement of the graph is a topological disk. In this paper we give a bijective operation that relates unicellular maps on a non-orientable surface to unicellular maps of a lower topological type, with distinguished vertices. From that we obtain a recurrence equation that leads to (new) explicit counting formulas for non-orientable precubic (all vertices of degree 1 or 3) unicellular maps of fixed topology. We also determine asymptotic formulas for the number of all unicellular maps of fixed topology, when the number of edges goes to infinity. Our strategy is inspired by recent results obtained for the orientable case [Chapuy, PTRF 2010], but significant novelties are introduced: in particular we construct an involution which, in some sense, “averages” the effects of non-orientability.

Résumé. Une carte unicellulaire est le plongement d’un graphe connexe dans une surface, tel que le complémentaire du graphe est un disque topologique. On décrit une opération bijective qui relie les cartes unicellulaires sur une surface non-orientable aux cartes unicellulaires de type topologique inférieur, avec des sommets marqués. On en déduit une récurrence qui conduit à de (nouvelles) formules closes d’énumération pour les cartes unicellulaires précubiques (sommets de degré 1 ou 3) de topologie fixée. On obtient aussi des formules asymptotiques pour le nombre total de cartes unicellulaires de topologie fixée, quand le nombre d’arêtes tend vers l’infini. Notre stratégie est motivée par de récents résultats dans le cas orientable [Chapuy, PTRF, 2010], mais d’importantes nouveautés sont introduites: en particulier, on construit une involution qui, en un certain sens, “moyenne” les effets de la non-orientabilité.

Keywords: One-face map, ribbon graph, non-orientable surface, bijection, involution

1 Introduction

A *map* is an embedding of a connected graph in a (2-dimensional, compact, connected) surface considered up to homeomorphism. By *embedding*, we mean that the graph is drawn on the surface in such a way that the edges do not intersect and the *faces* (connected components of the complementary of the graph) are simply connected. Maps are sometimes referred to as *ribbon graphs*, *fat-graphs*, and can be defined combinatorially rather than topologically as is recalled in Section 2. A map is *unicellular* if it has a single face. For instance, the unicellular maps on the sphere are the plane trees.

[†]Supported by ANR project A3.

[‡]Supported by a CNRS/PIMS postdoctoral fellowship.

In this paper we consider the problem of counting unicellular maps by the number of edges, when the topology of the surface is fixed. In the orientable case, this question has a respectable history. The first formula for the number $\epsilon_g(n)$ of orientable unicellular maps with n edges and $n + 1 - 2g$ vertices (hence genus g) was given by Lehman and Walsh in [WL72], as a sum over the integer partitions of size g . Independently, Harer and Zagier found a simple recurrence formula for the numbers $\epsilon_g(n)$ [HZ86]. Part of their proof relied on expressing the generating function of unicellular maps as a matrix integral. Other proofs of Harer-Zagier's formula were given in [Las01, GN05]. Recently, Chapuy [Cha09], extending previous results for cubic maps [Cha10], gave a bijective construction that relates unicellular maps of a given genus to unicellular maps of a smaller genus, hence leading to a new recurrence equation for the numbers $\epsilon_g(n)$. In particular, the construction in [Cha09] gives a combinatorial interpretation of the fact that for each g the number $\epsilon_g(n)$ is the product of a polynomial in n times the n -th Catalan number.

For non-orientable surfaces, results are more recent. The interpretation of matrix integrals over the Gaussian Orthogonal Ensemble (space of real symmetric matrices) in terms of maps was made explicit in [GJ97]. Ledoux [Led09], by means of matrix integrals and orthogonal polynomials, obtained for unicellular maps on general surfaces a recurrence relation which is similar to the Harer-Zagier one. As far as we know, no direct combinatorial nor bijective technique have successfully been used for the enumeration of a family of non-orientable maps until now.

A unicellular map is *precubic* if it has only vertices of degree 1 and 3: precubic unicellular maps are a natural generalization of binary trees to general surfaces. In this paper, we show that for all $h \in \frac{1}{2}\mathbb{N}$, the number of precubic unicellular maps of size m on the non-orientable surface of Euler Characteristic $2 - 2h$ is given by an explicit formula, which has the form of a polynomial in m times the m th Catalan number for $h \in \mathbb{N}$, and of a polynomial times 4^m if $h \notin \mathbb{N}$. These formulas, and our main results, are presented in Section 3. Our approach, which is completely combinatorial, is based on two ingredients. The first one, inspired from the orientable case [Cha10, Cha09], is to consider some special vertices called *intertwined nodes*, whose deletion reduces the topological type h of a map. The second ingredient is of a different nature: we show that, among non-orientable maps of a given topology and size, the *average number* of intertwined nodes per map can be determined explicitly. This is done thanks to an *averaging involution*, which is described in Section 4. This enables us to find a simple recurrence equation for the number of precubic unicellular maps by the number of edges and the topological type. As in the orientable case, an important feature of our recurrence is that it is recursive *only* on the topological type, contrarily to equations of the Harer-Zagier type [HZ86, Led09], where also the number of edges vary. It is then easy to iterate the recurrence, to obtain the promised counting formulas for precubic maps.

In the case of *general* (not necessarily precubic) unicellular maps, our approach does not work *exactly*, but it does work, in some sense, *asymptotically*. We obtain, with the same technique, the asymptotic number of non-orientable unicellular maps of fixed topology, when the number of edges tends to infinity. As far as we know, these formulas, and the ones for precubic maps, never appeared before in the literature.

2 Topological considerations

2.1 Classical definitions of surfaces and maps

Surfaces. Our *surfaces* are compact, connected, 2-dimensional manifolds. We consider surfaces up to homeomorphism. For any non-negative integer g , we denote by \mathbb{S}_g the g -torus, that is, the orientable surface obtained by adding g *handles* to the sphere. For any positive *half-integer* h , we denote by \mathbb{N}_h

the non-orientable surface obtained by adding $2h$ *cross-caps* to the sphere. Hence, \mathbb{S}_0 is the sphere, \mathbb{S}_1 is the torus, $\mathbb{N}_{1/2}$ is the projective plane and \mathbb{N}_1 is the Klein bottle. The *type* of the surface \mathbb{S}_h or \mathbb{N}_h is the number $h \in \{0, \frac{1}{2}, 1, \frac{3}{2}, \dots\}$. By the theorem of classification, each orientable surface is homeomorphic to one of the \mathbb{S}_g and each non-orientable surface is homeomorphic to one of the \mathbb{N}_h (see e.g. [MT01]).

Maps as graphs embedding. Our *graphs* are finite and undirected; loops and multiple edges are allowed. A *map* is an embedding (without edge-crossings) of a connected graph into a surface, in such a way that the *faces* (connected components of the complement of the graph) are simply connected. Maps are always considered up to homeomorphism. A map is *unicellular* if it has a single face.

Each edge in a map is made of two *half-edges*, obtained by removing its middle-point. The *degree* of a vertex is the number of incident half-edges. A *leaf* is a vertex of degree 1. A *corner* in a map is an angular sector determined by a vertex, and two half-edges which are consecutive around it. The total number of corners in a map equals the number of half-edges which is twice the number of edges. A map is *rooted* if it carries a distinguished half-edge called the *root*, together with a distinguished side of this half-edge. The vertex incident to the root is the *root vertex*. The unique corner incident to the root half-edge and its distinguished side is the *root corner*. *From now on, all maps are rooted.*

The *type* of a map is the type of the underlying surface. If \mathbf{m} is a map, we let $v(\mathbf{m})$, $e(\mathbf{m})$, $f(\mathbf{m})$ and $h(\mathbf{m})$ be its numbers of vertices, edges, faces, and its type. These quantities satisfy the *Euler formula*:

$$e(\mathbf{m}) = v(\mathbf{m}) + f(\mathbf{m}) + 2 - 2h(\mathbf{m}). \quad (1)$$

Maps as graphs with rotation systems and twists. Let G be a graph. To each edge e of G correspond two *half-edges*, each of them incident to an endpoint of e (they are both incident to the same vertex if e is a loop). A *rotation system* for G is the choice, for each vertex v of G , of a cyclic ordering of the half-edges incident to v . We now explain the relation between maps and rotation systems. Our surfaces are locally orientable and an *orientation convention* for a map \mathbf{m} is the choice of an orientation, called *counterclockwise orientation*, in the vicinity of each vertex. Any orientation convention for the map \mathbf{m} induces a rotation system on the underlying graph, by taking the counterclockwise ordering of appearance of the half-edges around each vertex. Given an orientation convention, an edge $e = (v_1, v_2)$ of \mathbf{m} is a *twist* if the orientation conventions in the vicinity of the endpoints v_1 and v_2 are not simultaneously extendable to an orientation of a vicinity of the edge e ; this happens exactly when the two sides of e appear in the same order when crossed clockwise around v_1 and clockwise around v_2 . Therefore a map together with an orientation convention defines both a rotation system and a subset of edges (the twists). The *flip* of a vertex v consists in inverting the orientation convention at that vertex. This changes the rotation system at v by inverting the cyclic order on the half-edges incident to v , and changes the set of twists by the fact that non-loop edges incident to e become twist if and only if they were not twist (while the status of the other edges remain unchanged). The next lemma is a classical topological result (see e.g. [MT01]).

Lemma 1 *A map (and the underlying surface) is entirely determined by the triple consisting of its (connected) graph, its rotation system, and the subset of its edges which are twists. Conversely, two triples define the same map if and only if one can be obtained from the other by flipping some vertices.*

By the lemma above, we can represent maps of positive types on a sheet of paper as follows: we draw the graph (with possible edge crossings) in such a way that the rotation system at each vertex is given by the counterclockwise order of the half-edges, and we indicate the twists by marking them by a cross (see e.g. Figure 1). The faces of the map are in bijection with the *borders* of that drawing, which are obtained

by walking along the edge-sides of the graph, and using the crosses in the middle of twisted edges as "crosswalks" that change the side of the edge along which one is walking (Figure 1). Observe that the number of faces of the map gives the type of the underlying surface using Euler formula.




Fig. 1: A representation of a map on the Klein bottle with three faces. The border of one of them is distinguished in dotted lines.




Fig. 2: (a) a twist; (b) a left corner; (c) a right corner.

2.2 Unicellular maps, tours, and canonical rotation system

Tours of unicellular maps. Let m be a unicellular map. By definition, m has a unique face. The *tour* of the map m is done by following the edges of m starting from the root corner along the distinguished side of the root half-edge, until returning to the root-corner. Since m is unicellular, every corner is visited once during the tour. An edge is said *two-ways* if it is followed in two different directions during the tour of the map (this is always the case on orientable surfaces), and is said *one-way* otherwise. The tour induces an *order of appearance* on the set of corners, for which the root corner is the least element. We denote by $c < d$ if the corner c appears before the corner d along the tour. Lastly, given an orientation convention, a corner is said *left* if it lies on the left of the walker during the tour of map, and *right* otherwise (Figure 2).

Canonical rotation-system. As explained above, the rotation system associated to a map is defined up to the choice of an orientation convention. We now explain how to choose a particular convention which will be well-suited for our purposes. A map is said *precubic* if all its vertices have degree 1 or 3, and its root-vertex has degree 1. Let m be a precubic unicellular map. Since the vertices of m all have an odd degree, there exists a unique orientation convention at each vertex such that the number of left corners is more than the number of right corners (indeed, by flipping a vertex we exchange its left and right corners). We call *canonical* this orientation convention. From now on, *we will always use the canonical orientation convention*. This defines canonically a rotation system, a set of *twists*, and a set of *left/right corners*. Observe that the root corner is a left corner (as is any corner incident to a leaf) and that vertices of degree 3 are incident to either 2 or 3 left corners. We have the following additional property.

Lemma 2 *In a (canonically oriented) precubic unicellular map, two-ways edges are incident to left corners only and are not twists.*

Proof: Let e be a two-ways edge, and let c_1, c_2 be two corners incident to the same vertex and separated by e (c_1 and c_2 coincide if v has degree 1). Since e is two-ways, the corners c_1, c_2 are either simultaneously left or simultaneously right. By definition of the canonical orientation, they have to be simultaneously left. Thus two-way edges are only incident to left corners. Therefore two-ways edges are not twists since following a twisted edge always leads from a left corner to a right corner or the converse. \square

2.3 Intertwined nodes.

We now define a notion of *intertwined node* which generalizes the definition given in [Cha10] for precubic maps on orientable surfaces.

Definition 1 Let m be a (canonically oriented) precubic unicellular map, let v be a vertex of degree 3, and let c_1, c_2, c_3 be the incident corners in counterclockwise order around v , with the convention that c_1 is the first of these corners to appear during the tour of m .

- The vertex v is an intertwined node if c_3 appears before c_2 during the tour of m .
- Moreover, we say that v has flavor **A** if it is incident to three left corners. Otherwise, v is incident to exactly one right corner, and we say that v is of flavor **B**, **C**, or **D** respectively, according to whether the right corner is c_1, c_2 or c_3 .

Observe that the definition of the canonical orientation was a prerequisite to define intertwined nodes. We will now show that intertwined nodes are exactly the ones whose deletion decreases the type of the map without disconnecting it.

The *opening* of an intertwined node of a map m is the operation consisting in splitting this vertex into three (marked) vertices of degree 1, as in Figure 3. That is, we define a rotation system and set of twists of the embedded graph n obtained in this way (we refrain from calling it a map yet, since it is unclear that it is connected) as the rotation system and set of twists inherited from the original map m .




Fig. 3: Opening an intertwined node.

Fig. 4: The tours of m and n , in the case of flavor **B**.

Proposition 1 Let n be a positive integer and let g be in $\{1, 3/2, 2, 5/2, \dots\}$. For each flavor F in $\{\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}\}$, the opening operation gives a bijection between the set of precubic unicellular maps with n edges, type h , and a distinguished intertwined node of flavor F , and the set of precubic unicellular maps with n edges, type $h - 1$ and three distinguished vertices of degree 1. The converse bijection is called the gluing of flavor F .

Moreover, if a precubic unicellular map m is obtained from a precubic unicellular map n of lower type by a gluing of flavor F , then m is orientable if and only if n is orientable and $F = \mathbf{A}$.

We omit the proof of the Proposition. However, let us give a "picture" of what happens, in the case of flavor **B**. If m is a unicellular map, and v is an intertwined node of m , then the sequence of corners appearing during the tour of m has the form $w(m) = w_1c_1w_2c_3w_3c_2w_4$, where c_1, c_2, c_3 are as in Definition 1, and w_1, w_2, w_3, w_4 are sequences of corners. Now, let n be the map obtained by opening m at v . If n has flavor **B**, then by following the edges of the map n , starting from the root, one gets the sequence of corners $w(n) = w_1d_1\bar{w}_3d_2w_2d_3w_4$, where \bar{w}_3 is the mirror word of w_3 , as can be seen from Figure 4

(we used three new letters d_1, d_2, d_3 for the three corners of degree 1 appearing after the opening). Since this sequence contains all the corners of \mathfrak{n} , we know that \mathfrak{n} is a unicellular map, and since it has two more vertices than \mathfrak{m} , its type is $h(\mathfrak{n}) = h(\mathfrak{m}) - 1$ (by Euler's formula).

Conversely, given a unicellular map \mathfrak{n} with three distinguished leaves d_1, d_2, d_3 , the gluing of flavor \mathbf{B} can be defined by identifying these three vertices to a single vertex v , and then choosing the rotation system and the twisted edges at v appropriately to ensure that the resulting map \mathfrak{m} is unicellular, and that v is an intertwined node of flavor \mathbf{B} in \mathfrak{m} .

The last statement of the Proposition is a consequence of the fact that a precubic unicellular map is orientable if and only if it has left-corners only in its canonical orientation.

3 Main results.

3.1 The number of precubic unicellular maps.

In this section, we present our main results, which rely on two facts. The first one is Proposition 1, which enables us to express the number of precubic unicellular maps of type h carrying a distinguished intertwined node in terms of the number of maps of a smaller type. The second one is the fact that, among maps of type h and fixed size, the average number of intertwined nodes in a map is $2h - 1$. This last fact, which is technically the most difficult part of this paper, relies on the existence of an ‘‘averaging involution’’, which will be described in Section 4.

Let $h \geq 1$ be an element of $\frac{1}{2}\mathbb{N}$, and let $m \geq 1$ be an integer. Given m and h , we let $n = 2m + \mathbb{1}_{h \in \mathbb{N}}$, and we let $\mathcal{O}_h(m)$ and $\mathcal{N}_h(m)$, respectively, be the sets of orientable and non-orientable precubic unicellular maps of type h with n edges. We let $\xi_h(m)$ and $\eta_h(m)$, respectively, be the cardinality of $\mathcal{O}_h(m)$ and $\mathcal{N}_h(m)$.

In order to use Proposition 1, we first need the following easy consequence of Euler's formula:

Lemma 3 *Let $l \in \frac{1}{2}\mathbb{N}$ and let \mathfrak{m} be a precubic unicellular map of type l with $n = 2m + \mathbb{1}_{l \in \mathbb{N}}$ edges. Then \mathfrak{m} has $m + (-1)^{2l} - 3\lfloor l \rfloor$ non-root leaves, where $\lfloor l \rfloor = l - \frac{1}{2}\mathbb{1}_{l \notin \mathbb{N}}$ denotes the integer part of l .*

From the lemma and Proposition 1, the number $\eta_h^{\text{inter}}(m)$ of non-orientable unicellular precubic maps of type h with n edges carrying a distinguished intertwined node equals:

$$\eta_h^{\text{inter}}(m) = 4 \binom{m' - 3\lfloor h - 1 \rfloor}{3} \eta_{h-1}(m) + 3 \binom{m' - 3\lfloor h - 1 \rfloor}{3} \xi_{h-1}(m), \quad (2)$$

where $m' = m + (-1)^{2h}$. Here, the first term accounts for intertwined nodes obtained by gluing three leaves in a non-orientable map of type $h - 1$ (in which case the flavor of the gluing can be either \mathbf{A} , \mathbf{B} , \mathbf{C} or \mathbf{D}), and the second term corresponds to the case where the starting map of type $h - 1$ is orientable (in which case the gluing has to be of flavor \mathbf{B} , \mathbf{C} or \mathbf{D} to destroy the orientability).

The keystone of this paper, which will be discussed in the next section, is the following result:

Proposition 2 *There exists an involution Φ of $\mathcal{N}_h(m)$ such that for all maps $\mathfrak{m} \in \mathcal{N}_h(m)$, the total number of intertwined nodes in the maps \mathfrak{m} and $\Phi(\mathfrak{m})$ is $4h - 2$. In particular, the average number of intertwined nodes of elements of $\mathcal{N}_h(m)$ is $(2h - 1)$, and one has $\eta_h^{\text{inter}}(m) = (2h - 1)\eta_h(m)$.*

It is interesting to compare Proposition 2 with the analogous result in [Cha10]: in the orientable case, each map of genus h has exactly $2h$ intertwined nodes, whereas here the quantity $(2h - 1)$ is only an

average value. For example, Figure 5 shows two maps on the Klein bottle ($h = 1$) which are related by the involution Φ : they have respectively 2 and 0 intertwined nodes.

As a direct corollary of Proposition 2 and Equation (2), we can state our main result:

Theorem 1 *The numbers $\eta_h(m)$ of non-orientable precubic unicellular maps of type h with $2m + 1$ edges obey the following recursion:*

$$(2h - 1) \cdot \eta_h(m) = 4 \binom{m' - 3\lfloor h - 1 \rfloor}{3} \eta_{h-1}(m) + 3 \binom{m' - 3\lfloor h - 1 \rfloor}{3} \xi_{h-1}(m), \quad (3)$$

where $m' = m + (-1)^{2h}$, and where $\xi_h(m)$ is the number of orientable precubic unicellular maps of genus h with $2m + 1$ edges, which is 0 if $h \notin \mathbb{N}$, and is given by the following formula otherwise [Cha09]:

$$\xi_h(m) = \frac{1}{(2h)!!} \binom{m+1}{3, 3, \dots, 3, m+1-3h} \text{Cat}(m) = \frac{(2m)!}{12^h h! m! (m+1-3h)!}. \quad (4)$$

The theorem implies explicit formulas for the numbers $\eta_h(m)$, as shown by the two next corollaries:

Corollary 1 (the case $h \in \mathbb{N}$) *Let $h \in \mathbb{N}$ and $m \in \mathbb{N}$, $m \geq 3h - 1$. Then the number of non-orientable precubic unicellular maps of type h with $2m + 1$ edges equals:*

$$\eta_h(m) = c_h \binom{m+1}{3, 3, \dots, 3, m+1-3h} \text{Cat}(m) = \frac{c_h \cdot (2m)!}{6^h m! (m+1-3h)!} \quad (5)$$

where $c_h = 3 \cdot 2^{3h-2} \frac{h!}{(2h)!} \sum_{l=0}^{h-1} \binom{2l}{l} 16^{-l}$.

Corollary 2 (the case $h \notin \mathbb{N}$) *Let $h \in \frac{1}{2} + \mathbb{N}$ and $m \in \mathbb{N}$, $m \geq 3\lfloor h \rfloor + 1$. Then the number of non-orientable precubic unicellular maps of type h with $2m$ edges equals:*

$$\begin{aligned} \eta_h(m) &= \frac{4^{\lfloor h \rfloor}}{(2h-1)(2h-3)\dots 1} \binom{m-1}{3, 3, \dots, 3, m-1-3\lfloor h \rfloor} \times \eta_{1/2}(m) \\ &= \frac{4^{m+\lfloor h \rfloor-1} (m-1)!}{6^{\lfloor h \rfloor} (2h-1)!! (m-1-3\lfloor h \rfloor)!} \end{aligned}$$

Proof of Corollary 1: It follows by induction and Equations (3) and (4) that the statement of Equation (5) holds, with the constant c_h defined by the recurrence $c_0 = 0$ and $c_h = a_{h-1} + b_{h-1}c_{h-1}$, with $a_{h-1} = \frac{3}{2^{h-1}(h-1)!(2h-1)}$ and $b_{h-1} = \frac{4}{2h-1}$. The solution of this recurrence is $c_h = \sum_{l=0}^{h-1} a_l b_{l+1} b_{l+2} \dots b_{h-1}$.

Now, by definition, we have $a_l b_{l+1} b_{l+2} \dots b_{h-1} = \frac{3 \cdot 4^{h-1-l}}{2^l l! (2l+1)(2l+3)(2l+5) \dots (2h-1)}$. Using the expression $\frac{1}{(2l+1)(2l+3) \dots (2h-1)} = \frac{2^h h! (2l)!}{(2h)! 2^l l!}$ and reporting it in the sum gives the expression of c_h given in Corollary 1. \square

Proof of Corollary 2: Since for non-integer h we have $\xi_{h-1}(m) = 0$, the first equality is a direct consequence of an iteration of the theorem. Therefore the only thing to prove is that $\eta_{1/2}(m) = 4^{m-1}$.

This can be done easily by induction via an adaptation of Rémy's bijection [Rém85], as follows. For $m = 1$, we have $\eta_{1/2}(m) = 1$, since the only precubic projective unicellular map with two edges is "the twisted loop with a hanging leaf". For the induction step, observe that precubic projective unicellular maps with one distinguished non-root leaf are in bijection with precubic projective unicellular maps with one leaf less and a distinguished edge-side: too see that, delete the distinguished leaf, transform the remaining vertex of degree 2 into an edge, and remember the side of that edge on which the original leaf was attached. Since a projective precubic unicellular map with $2k$ edges has $k - 1$ non-root leaves and $4k$ edge-sides, we obtain for all $m \geq 1$ that $m \eta_{1/2}(m + 1) = 4m \eta_{1/2}(m)$, and the result follows. \square




Fig. 5: Two maps on the Klein Bottle. (a) $T_{LR}(m) = 1$, $T_{RL}(m) = 1$; (b) $T_{LR}(m) = 2$, $T_{RL}(m) = 0$.




Fig. 6: The opening, in the case of dominant unicellular maps.

3.2 The asymptotic number of rooted unicellular maps.

Though our results do not apply to the general case of *all* unicellular maps of given type (i.e., not necessarily precubic), they do hold, in some sense, *asymptotically*. This is what we explain in this section.

If m is a unicellular map, its *core* is the map obtained by deleting recursively all the leaves of m , until having only vertices of degree 2 or more left. Therefore the core is a unicellular map formed by chains of vertices of degree 2 attached together at vertices of degree at least 3. The *scheme* of m is the map obtained by replacing each of these chains by an edge. Hence, in the scheme, all vertices have degree at least 3. We say that a unicellular map is *dominant* if all the vertices of its scheme have degree 3. This terminology, borrowed from [Cha10], comes from the next proposition.

Proposition 3 ([CMS09, BR09]) *Let $h \in \frac{1}{2}\mathbb{N}$. Then, among non-orientable unicellular maps of type h with n edges, the proportion of maps which are dominant tends to 1 when n tends to infinity.*

The idea behind that proposition is the following. Given a scheme \mathfrak{s} , one can easily compute the generating series of all unicellular maps of scheme \mathfrak{s} , by observing that these maps are obtained by substituting each edge of the scheme with a "branch of tree". From that, it follows that this generating series has a unique principal singularity at $z = \frac{1}{4}$, with dominating term $(1 - 4z)^{-e(\mathfrak{s})/2-1}$, up to a multiplicative constant. Therefore, the schemes with the greatest contribution are those which have the maximal number of edges, which for a given type, is achieved by schemes whose all vertices have degree 3.

Now, most of the combinatorics defined in this paper still apply to dominant unicellular maps. Given a dominant map m of type h and scheme \mathfrak{s} , and v an intertwined node of \mathfrak{s} , we can define the opening

operation of m at v by splitting the vertex v in three, and deciding on a convention on the redistribution of the three "subtrees" attached to the scheme at this point (Figure 6): one obtains a dominant map n of type $h - 1$ with three distinguished vertices. These vertices are not *any* three vertices: they have to be in "general position" in n (i.e., they cannot be part of the core, and none can lie on a path from one to another), but again, in the asymptotic case this does not make a big difference: when n tends to infinity, the proportion of vertices which are in "general position" tends to 1. We do not state here the asymptotic estimates that can make the previous claims precise (they can be copied almost verbatim from the orientable case [Cha10]), but rather we state now our asymptotic theorem:

Theorem 2 *Let $\kappa_h(n)$ be the number of non-orientable rooted unicellular maps of type h with n edges. Then one has, when n tends to infinity:*

$$(2h - 1)\kappa_h(n) \sim 4\frac{n^3}{3!}\kappa_{h-1}(n) + 3\frac{n^3}{3!}\epsilon_{h-1}(n)$$

where $\epsilon_h(n)$ denotes the number of orientable rooted unicellular maps of genus h with n edges.

Using that $\epsilon_h(n) = 0$ if $h \notin \mathbb{N}$, that $\epsilon_h(n) \sim \frac{1}{12^h h! \sqrt{\pi}} n^{3h - \frac{3}{2}}$ otherwise, and that $\kappa_{1/2}(n) \sim \frac{1}{2} 4^n$ [BCR88], we obtain:

Corollary 3 *Let $h \in \frac{1}{2}\mathbb{N}$. Then one has, when n tends to infinity:*

$$\kappa_h(n) \sim \frac{c_h}{\sqrt{\pi} 6^h} n^{3h - \frac{3}{2}} 4^n \text{ if } h \in \mathbb{N} \quad , \quad \kappa_h(n) \sim \frac{4^{\lfloor h \rfloor}}{2 \cdot 6^{\lfloor h \rfloor} (2h - 1)!!} n^{3h - \frac{3}{2}} 4^n \text{ if } h \notin \mathbb{N}.$$

where the constant c_h is defined in Corollary 1.

4 The average number of intertwined nodes

In this section we prove Proposition 2, and in particular the key result that the average number of intertwined nodes per map, among precubic unicellular maps of type h and size m is $(2h - 1)$:

$$\eta_h^{\text{inter}}(m) = (2h - 1)\eta_h(m). \quad (6)$$

Let us emphasize the fact that the number of intertwined nodes is not a constant over the set of unicellular precubic maps of given type and number of edges. For instance among the six maps with 5 edges on the Klein bottle \mathbb{N}_1 , three maps have 2 intertwined nodes, and three maps have none; see Figure 7. As stated in Proposition 2, our strategy to prove Equation (6) is to exhibit a bijection Φ from the set $\mathcal{N}_h(m)$ to itself, such that for any given map m , the total number of intertwined nodes in the maps m , $\Phi(m)$ is $4h(m) - 2$. Observe from Figure 7 that the involution Φ cannot be a simple re-rooting of the map m .

Before defining the mapping Φ , we relate the number of intertwined nodes of a map to certain properties of its twists. Let m be a (canonically oriented) precubic map, and let e be an edge of m which is a twist. Let c be the corner incident to e which appears first in the tour of m . We say that e is *left-to-right* if c is a left-corner, and that it is *right-to-left* otherwise (see Figure 5). In other words, the twist e is left-to-right if it changes the side of the corners from left, to right, when it is crossed for the first time in the tour of the map (and the converse is true for right-to-left twists). We omit the proof of the next lemma:




Fig. 7: The precubic unicellular maps with 5 edges on the Klein bottle (the root in the unique leaf corner). Intertwined nodes are indicated as white vertices.

Lemma 4 *Let \mathfrak{m} be a precubic unicellular map of type $h(\mathfrak{m})$, considered in its canonical orientation. Then its numbers $\tau(\mathfrak{m})$ of intertwined nodes, $T_{\text{LR}}(\mathfrak{m})$ of left-to-right twists, and $T_{\text{RL}}(\mathfrak{m})$ of right-to-left twists are related by the formula:*

$$2h(\mathfrak{m}) = \tau(\mathfrak{m}) + T_{\text{LR}}(\mathfrak{m}) - T_{\text{RL}}(\mathfrak{m}). \quad (7)$$

We now define the promised mapping Φ averaging the number of intertwined nodes. Let \mathfrak{m} be a unicellular precubic map on a non-orientable surface. We consider the canonical orientation convention for the map \mathfrak{m} , which defines a rotation system and set of twists. The set of twists is non-empty since the map \mathfrak{m} lives on a non-orientable surface. By cutting every twist of \mathfrak{m} at their middle point, one obtains a graph together with a rotation system and some *dangling half-edges* that we call *buds*. The resulting embedded graph with buds, which we denote by $\hat{\mathfrak{m}}$, can have several connected components and each component (which is a map with buds) can have several faces; see Figure 8. We set a convention for the direction in which one *turns around a face* of $\hat{\mathfrak{m}}$: the edges are followed in such a way that every corner is left (this is possible since $\hat{\mathfrak{m}}$ has no twist). For any bud b of $\hat{\mathfrak{m}}$, we let $\sigma(b)$ be the bud following b when turning around the face of $\hat{\mathfrak{m}}$ containing b . Clearly, the mapping σ is a permutation on the set of buds. We now define $\Phi(\mathfrak{m})$ to be the graph with rotation system and twists obtained from $\hat{\mathfrak{m}}$ by gluing together into a twist the buds $\sigma(b)$ and $\sigma(b')$ for every pair of buds b, b' forming a twist of \mathfrak{m} . The mapping Φ is represented in Figure 8.

Before proving that $\Phi(\mathfrak{m})$ is a unicellular map, we set some additional notations. We denote by k the number of twists of \mathfrak{m} and we denote by $w(\mathfrak{m}) = w_1 w_2 \cdots w_{2k+1}$ the sequence of corners encountered during the tour of \mathfrak{m} , where the subsequences w_i and w_{i+1} are separated by the traversal of a twist for $i = 1 \dots 2k$. Observe that corners in w_i are left corners of \mathfrak{m} if i is odd, and right corners if i is even (since following a twist leads from a left to a right corner or the converse). Hence, the sequence of corners encountered between two buds around a face of $\hat{\mathfrak{m}}$ are one of the sequences $w'_1, w'_2, \dots, w'_{2k}$, where $w'_1 = w_{2k+1} w_1$, and for $i > 1$, $w'_i = w_i$ if i is odd and $w'_i = \bar{w}_i$ otherwise (where \bar{w}_i is the *mirror* sequence of w_i obtained by reading w_i backwards). We identify the buds of $\hat{\mathfrak{m}}$ (i.e. the half-twists of \mathfrak{m} or $\hat{\mathfrak{m}}$) with the integers in $\{1, \dots, 2k\}$ by calling i the bud following the sequence of corners w'_i around the faces of $\hat{\mathfrak{m}}$. The permutation σ can then be considered as a permutation of $\{1, \dots, 2k\}$ and we denote $r = \sigma^{-1}(1)$. The map in Figure 8 gives $\sigma = (1, 8, 13, 2, 9, 14, 3, 10)(4, 11, 6, 5)(7, 12)$ and $r = 10$.

Lemma 5 *The embedded graph $\Phi(\mathfrak{m})$ is a unicellular map. Moreover, the rotation system and set of twists of $\Phi(\mathfrak{m})$ inherited from \mathfrak{m} correspond to the canonical orientation convention of $\Phi(\mathfrak{m})$. Lastly, the sequence of corners encountered during the tour of $\Phi(\mathfrak{m})$ reads $v_1 v_2 \dots v_{2k+1}$, where the subsequences v_i separated twist traversals are given by $v_i = w_{\sigma(r+1-i)}$ for $i = 1, \dots, r$, $v_i = w_{\sigma(2n+r+1-i)}$ for $i = r+1, \dots, 2k$, and $v_{2k+1} = w_{2k+1}$.*




Fig. 8: A unicellular map m and its image by the mapping Φ . The twists are indicated by (partially) dotted lines, while the map \hat{m} is represented in solid lines.

Lemma 6 *Let m be a positive integer and h be in $\{1/2, 1, 3/2, \dots\}$. The mapping Φ is a bijection from the set $\mathcal{N}_h(m)$ to itself. Moreover, for every map m in $\mathcal{N}_h(m)$, the total number of intertwined nodes in the maps m and $\Phi(m)$ is $4h - 2$.*

Proof of Lemma 6: Clearly, the maps m and $\Phi(m)$ have the same number of edges and vertices. Hence, they have the same type by Euler formula. Moreover, they both have $k > 0$ twists (for their canonical convention) hence are non-orientable. Thus, Φ maps the set $\mathcal{N}_h(m)$ to itself. To prove the bijectivity (i.e. injectivity) of Φ , observe that for any map m , the embedded graphs \hat{m} and $\widehat{\Phi(m)}$ are equal; this is because the canonical rotation system and set of twists of m and $\Phi(m)$ coincide. In particular, the permutation σ on the half-twists of m can be read from $\Phi(m)$. Hence, the twists of m are easily recovered from those of $\Phi(m)$: the buds i and j form a twist of m if $\sigma(i)$ and $\sigma(j)$ form a twist of $\Phi(m)$.

We now proceed to prove that the total number of intertwined nodes in m and $\Phi(m)$ is $4h - 2$. By Lemma 4, this amounts to proving that $T_{LR}(m) - T_{RL}(m) + T_{LR}(\Phi(m)) - T_{RL}(\Phi(m)) = 2$. Since m and $\Phi(m)$ both have k twists, $T_{LR}(m) - T_{RL}(m) + T_{LR}(\Phi(m)) - T_{RL}(\Phi(m)) = 2(T_{LR}(m) + T_{LR}(\Phi(m)) - k)$. Hence, we have to prove $T_{LR}(m) + T_{LR}(\Phi(m)) = k + 1$.

Let i be a bud of \hat{m} , let t be the twist of m containing i , and let c, c' be the corners preceding and following i in counterclockwise order around the vertex incident to i . By definition, the twist t of m is left-to-right if and only if c appears before c' during the tour of m . Given that the corners c and c' belong respectively to the subsequences w_i and $w_{\sigma(i)}$ (except if $i = r$ in which case $\sigma(i) = 1$ and c' is in w_{2k+1}), the twist t is left-to-right if and only if $i < \sigma(i)$ or $i = r$.

Before going on, let us introduce a notation: for an integer i we denote by \bar{i} the representative of i modulo $2k$ belonging to $\{1, \dots, 2k\}$. Let us now examine under which circumstances the bud $\sigma(i)$ is part of a left-to-right twist of $\Phi(m)$. The corners d and d' preceding and following the bud $\sigma(i)$ in counterclockwise order around the vertex incident to $\sigma(i)$ belong respectively to $w_{\sigma(i)}$ and $w_{\sigma\sigma(i)}$ (except if $\sigma(i) = r$, in which case $\sigma\sigma(i) = 1$ and d' belongs to w_{2k+1}). By Lemma 5, $w_{\sigma(i)} = \overline{v_{r+1-i}}$ for $i = 1 \dots 2k$. Therefore, the twist t' of $\Phi(m)$ containing $\sigma(i)$ is left-to-right if and only if $r + 1 - \bar{i} < r + 1 - \sigma(i)$ or $\sigma(i) = r$.

The two preceding points gives the number $T_{LR}(m) + T_{LR}(\Phi(m))$ of left-to right twists as

$$T_{LR}(m) + T_{LR}(\Phi(m)) = 1 + \frac{1}{2} \sum_{i=1}^{2k} \delta(i),$$

where $\delta(i) = \mathbb{1}_{i < \sigma(i)} + \frac{1}{2} \mathbb{1}_{r+1-i < r+1-\sigma(i)}$ is the sum of two indicator functions (the factor $1/2$ accounts for the fact that a twist has two halves). The contribution $\delta(i)$ is equal to 2 if $i \leq r < \sigma(i)$, 0 if $\sigma(i) \leq r < i$, and 1 otherwise. Finally, there are as many integers i such that $i \leq r < \sigma(i)$ as integers such that $\sigma(i) \leq r < i$ (true for each cycle of σ). Thus, $\sum_{i=1}^{2k} \delta(i) = 2k$, and $T_{LR}(\mathbf{m}) + T_{LR}(\Phi(\mathbf{m})) = k + 1$. \square

The last lemma is sufficient to establish Equation (6), and the enumerative results of Section 3. However, Proposition 2 was saying a little bit more, namely that the bijection Φ can be chosen as an *involution*:

Proof of Proposition 2: Observe that, as we defined it, the bijection Φ is not an involution. But one can easily define an involution from Φ , as the mapping acting as Φ on elements \mathbf{m} of $\mathcal{N}_h(m)$ such that $\tau(\mathbf{m}) > 2h - 1$, acting as Φ^{-1} if $\tau(\mathbf{m}) < 2h - 1$, and as the identity if $\tau(\mathbf{m}) = 2h - 1$. \square

References

- [BCR88] E. A. Bender, E. R. Canfield, and R. W. Robinson. The enumeration of maps on the torus and the projective plane. *Canad. Math. Bull.*, 31:257–271, 1988.
- [BR09] O. Bernardi and J. Rué. Enumerating simplicial decompositions of surfaces with boundaries. *arXiv:0901.1608*, 2009.
- [Cha09] G. Chapuy. A new combinatorial identity for unicellular maps, via a direct bijective approach. *DMTCS (FPSAC '09) – Preliminary long version at www.lix.polytechnique.fr/chapuy*, 2009.
- [Cha10] G. Chapuy. The structure of unicellular maps, and a connection between maps of positive genus and planar labelled trees. *Probability Theory and Related Fields*, 147(3):415–447, 2010.
- [CMS09] G. Chapuy, M. Marcus, and G. Schaeffer. A bijection for rooted maps on orientable surfaces. *SIAM J. Discrete Math.*, 23(3):1587–1611, 2009.
- [GJ97] I. Goulden and D.M. Jackson. Maps in locally orientable surfaces and integrals over real symmetric matrices. *Canad. J. Math*, 48, 1997.
- [GN05] I. Goulden and A. Nica. A direct bijection for the Harer-Zagier formula. *JCTA*, 2005.
- [HZ86] J. Harer and D. Zagier. The Euler characteristic of the moduli space of curves. *Invent. Math.*, 85(3):457–485, 1986.
- [Las01] B. Lass. Démonstration combinatoire de la formule de Harer-Zagier. *C. R. Acad. Sci. Paris*, 333:155–160, 2001.
- [Led09] M. Ledoux. A recursion formula for the moments of the Gaussian orthogonal ensemble. *Ann. Inst. Henri Poincaré Probab. Stat.*, 45(3):754–769, 2009.
- [MT01] B. Mohar and C. Thomassen. *Graphs on surfaces*. J. Hopkins Univ. Press, 2001.
- [Ré85] J.-L. Rémy. Un procédé itératif de dénombrement d’arbres binaires et son application à leur génération aléatoire. *RAIRO Inform. Théor.*, 19(2):179–195, 1985.
- [WL72] T. R. S. Walsh and A. B. Lehman. Counting rooted maps by genus. I. *J. Combinatorial Theory Ser. B*, 13:192–218, 1972.