

Counting unicellular maps on non-orientable surfaces

Olivier Bernardi*

Department of Mathematics,
Massachusetts Institute of Technology,
77 Massachusetts Avenue,
Cambridge MA 02139, USA.

Guillaume Chapuy†

Department of Mathematics,
Simon Fraser University,
8888 University Drive,
Burnaby B.C. V5A 1S6, Canada.

August 10, 2010

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 present a bijective link between unicellular maps on a non-orientable surface and 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 unicellular maps of fixed topology. In particular, we give exact formulas for the pre-cubic case (all vertices of degree 1 or 3), and asymptotic formulas for the general case, when the number of edges goes to infinity. Our strategy is inspired by recent results obtained by the second author for the orientable case, but significant novelties are introduced: in particular we construct an involution which, in some sense, “averages” the effects of non-orientability.

Keywords: One-face map, one-vertex triangulation, bijection, Harer-Zagier formula.

Mathematics Subject Classification: 05A15, 05A19, 05C30.

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.

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

*Supported by ANR project A3 and by European project ExploreMaps - ERC StG 208471.

†Supported by a CNRS/PIMS postdoctoral fellowship.

with n edges and genus g (hence $n + 1 - 2g$ vertices) was given by Lehman and Walsh in [12], as a sum over the integer partitions of size g . Independently, Harer and Zagier found a simple recurrence formula for the numbers $\epsilon_g(n)$ [7]. 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 [8, 6]. Recently, Chapuy [2], extending previous results for cubic maps [3], 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 [2] 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 $\text{Cat}(n) = \frac{1}{n+1} \binom{2n}{n}$.

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 [5]. Ledoux [9], 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 give for all $h \in \frac{1}{2}\mathbb{N}$ an explicit formula for the number $\eta_h(m)$ of precubic unicellular maps of *size* m ($2m + \mathbb{1}_{h \in \mathbb{N}}$ edges) on the non-orientable surface of Euler Characteristic $2 - 2h$. These formulas (Corollaries 8 and 9) take the form $\eta_h(m) = P_h(m)\text{Cat}(m)$ if h is an integer, and $\eta_h(m) = P_h(m)4^m$ otherwise, where P_h is a polynomial of degree $3\lfloor h \rfloor$. Our approach, which is completely combinatorial, is based on two ingredients. The first one, inspired from the orientable case [3, 2], 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 numbers $\eta_h(m)$. As in the orientable case, an important feature of our recurrence is that it is recursive *only* on the topological type h , contrarily to equations of the Harer-Zagier type [7, 9], where also the number of edges vary. It is then easy to iterate the recurrence in order to obtain an explicit formula for $\eta_h(m)$.

In the case of *general* (not necessarily precubic) unicellular maps, our approach does not work *exactly*, but it does work *asymptotically*. That is, 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 (Theorem 11). As far as we know, all the formulas obtained in this paper are new.

2 Topological considerations

In this section we recall some definitions on maps and gather the topological tools needed for proving our results. One of these tools is a canonical way to represent non-orientable maps combinatorially which will prove very useful for our purposes.

We denote $\mathbb{N} = \{0, 1, 2, 3, \dots\}$ and $\frac{1}{2}\mathbb{N} = \{0, \frac{1}{2}, 1, \frac{3}{2}, \dots\}$. For a non-negative real number x , we denote by $\lfloor x \rfloor$ the integer part of x . For a non-negative integer n , we denote $n!! = n \cdot (n - 2)!!$ if $n > 1$, and $0!! = 1!! = 1$.

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 h , we denote by \mathcal{S}_h the torus of genus h , that is, the orientable surface obtained by adding h *handles* to the sphere. For any h in $\frac{1}{2}\mathbb{N}$, we denote by \mathcal{N}_h the non-orientable surface obtained by adding $2h$ *cross-caps* to the sphere. Hence, \mathcal{S}_0 is the sphere, \mathcal{S}_1 is the torus, $\mathcal{N}_{1/2}$ is the projective plane and \mathcal{N}_1 is the Klein bottle. The *type* of the surface \mathcal{S}_h or \mathcal{N}_h is the number h . By the theorem of classification, each orientable surface is homeomorphic to one of the \mathcal{S}_h and each non-orientable surface is homeomorphic to one of the \mathcal{N}_h (see e.g. [10]).

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* $h(\mathbf{m})$ of a map \mathbf{m} is the type of the underlying surface, that is to say, the Euler characteristic of the surface is $2 - 2h(\mathbf{m})$. If \mathbf{m} is a map, we let $v(\mathbf{m})$, $e(\mathbf{m})$ and $f(\mathbf{m})$ be its numbers of vertices, edges and faces. 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 counterclockwise around v_1 and counterclockwise 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. [10]).

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) with the convention that the counterclockwise order of the half-edges around each vertex on the paper gives the rotation system, 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’s formula.




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




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

2.2 Unicellular maps, tour, and canonical rotation system

Tour of a unicellular map. Let \mathbf{m} be a unicellular map. By definition, \mathbf{m} has a unique face. The *tour* of the map \mathbf{m} is done by following the edges of \mathbf{m} starting from the root corner along the distinguished side of the root half-edge, until returning to the root-corner. Since \mathbf{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 \mathbf{m} be a precubic unicellular map. Since the vertices of \mathbf{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, flipping a vertex change its left corners into right corners and *vice versa*). 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 that vertex 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 [3] for precubic maps on orientable surfaces.

Definition 3. *Let \mathbf{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 \mathbf{m} . The vertex v is called an intertwined node if c_3 appears before c_2 during the tour of \mathbf{m} .*

*Moreover, we say that the vertex 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. The intertwined of some unicellular maps on the Klein bottle are indicated in Figure 8. We will now show that intertwined nodes are exactly the ones whose deletion decreases the type of the map without disconnecting it nor increasing its number of faces.

The *opening* of an intertwined node of a map \mathbf{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 \mathbf{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 \mathbf{m} .




Figure 3: Opening of an intertwined node (of a precubic map).




Figure 4: Opening for non-precubic unicellular maps (dominant case).




Figure 5: The tours of \mathbf{m} and \mathbf{n} , in the case of flavor **A**, and in the case of flavor **B**.

Proposition 4. *Let n be a positive integer and let h be in $\{1, 3/2, 2, 5/2, \dots\}$. For each flavor \mathbf{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 \mathbf{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 \mathbf{F} .*

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

The opening of intertwined nodes of type **A** and **B** are represented in Figure 5.

Proof. We first show that the opening of an intertwined vertex produces a unicellular map (and decreases the type by 1). Let \mathbf{m} be a precubic unicellular map, and let v be an intertwined node. Let c_1, c_2, c_3 be the three corners incident to v in counterclockwise order, with the convention that c_1 is the first of these corners to appear during the tour of \mathbf{m} . Since v is intertwined, the sequence of corners appearing during the tour of \mathbf{m} has the form

$$w(\mathbf{m}) = w_1c_1w_2c_3w_3c_2w_4,$$

where w_1, w_2, w_3, w_4 are sequences of corners. Let \mathbf{n} be the embedded graph with marked vertices v_1, v_2, v_3 obtained by opening \mathbf{m} . We identify the corners of \mathbf{m} distinct from c_1, c_2, c_3 with the corners of \mathbf{n} distinct from the corners d_1, d_2, d_3 incident to v_1, v_2, v_3 . By following the edges of \mathbf{n} starting from the root corner along the distinguished side of the root half-edge, one gets a sequence of corners $w(\mathbf{n})$. If v has flavor **A**, this sequence of corners is

$$w(\mathbf{n}) = w_1d_1w_3d_2w_2d_3w_4,$$

as can be seen from Figure 5. Similarly, if v has flavor **B** (resp. **C**, **D**) then the sequence of corner is

$$w(\mathbf{n}) = w_1 d_1 \bar{w}_3 d_2 w_2 d_3 w_4, \quad (\text{resp. } w(\mathbf{n}) = w_1 d_1 w_3 d_2 \bar{w}_2 d_3 w_4, \quad w(\mathbf{n}) = w_1 d_1 \bar{w}_2 d_2 \bar{w}_3 d_3 w_4),$$

where \bar{w}_i is the *mirror* of the sequence w_i obtained by reading w_i backward. In each case, the sequence $w(\mathbf{n})$ contains all the corners of \mathbf{n} , implying that \mathbf{n} is a unicellular map. Moreover, \mathbf{n} has two more vertices than \mathbf{m} , so by Euler formula, its type is $h(\mathbf{n}) = h(\mathbf{m}) - 1$.

We now prove that \mathbf{m} is orientable if and only if v has flavor **A** and \mathbf{n} is orientable. Suppose first that \mathbf{m} is orientable. In this case, \mathbf{m} has no right corner, so that v has flavor **A**. Hence, the relation between $w(\mathbf{m})$ and $w(\mathbf{n})$ indicated above shows that \mathbf{n} has no right corner (in the orientation convention inherited from \mathbf{m}). Hence \mathbf{n} is orientable. Suppose conversely that v has flavor **A** and \mathbf{n} is orientable. Since v has flavor **A** the orientation convention of \mathbf{n} inherited from \mathbf{m} is the canonical one (indeed the corners of \mathbf{n} are right if and only if they are right in \mathbf{m} , and \mathbf{m} has a minority of right corner at each vertex). Since \mathbf{n} is orientable, this orientation convention has no right corner. Therefore, the map \mathbf{m} has no right corner either. Thus, \mathbf{m} is orientable.




Figure 6: The gluing of flavor **A**, **B**, **C** and **D**.

We now define the *gluing* operation (of flavor **A**, **B**, **C** or **D**) which we shall prove to be the inverse of the opening operation (on node of flavor **A**, **B**, **C** or **D**). Let us treat in details the gluing of flavor **B**; the other flavors being similar. Let \mathbf{n} be a precubic unicellular map with three distinguished leaves v_1, v_2, v_3 encountered in this order during the tour of \mathbf{n} . For $i = 1, 2, 3$ we denote by e_i and d_i respectively the edge and corner incident to v_i (see Figure 6). We consider the canonical orientation convention of \mathbf{n} . Clearly, e_1, e_2, e_3 are two-way edges, hence they are not twists for this convention (by Lemma 2). The *gluing of flavor B* on the map \mathbf{n} gives a map \mathbf{m} defined as follows: the graph of \mathbf{m} is the graph of \mathbf{n} after identification of the three leaves v_1, v_2, v_3 into a single vertex v , the rotation system of \mathbf{m} is the same as the rotation system of \mathbf{n} at any vertex distinct from v , and the rotation system at v is (e_1, e_2, e_3) in counterclockwise order, lastly the set of twists of \mathbf{m} is the set of twists of \mathbf{n} together with the two edges e_1 and e_2 . We now prove that the map \mathbf{m} is unicellular. Let us denote by

$$w(\mathbf{n}) = w_1 d_1 w_2 d_2 w_3 d_3 w_4$$

the sequence of corners encountered during the tour of \mathbf{n} (where the w_i are sequences of corners distinct from d_1, d_2, d_3). Let us denote by c_1, c_2, c_3 the corners of the new vertex v incident to $(e_2, e_1), (e_1, e_3), (e_3, e_2)$ respectively. By following the edges of \mathbf{m} starting from the root corner along the distinguished side of the root half-edge, one gets the sequence of corners

$$w(\mathbf{m}) = w_1 c_1 w_3 c_3 \bar{w}_2 c_2 w_4,$$

as can be seen from Figure 6. This sequence contains all corners of \mathbf{m} showing that \mathbf{m} is unicellular. We will now show that v is an intertwined vertex of flavor \mathbf{B} . Observe first that the corners in the sequence w_1 are followed in the same direction during the tour of \mathbf{m} and \mathbf{n} , so that the corners in w_1 are left corners in the map \mathbf{n} (for its canonical orientation convention) if and only if they are left corners in the map \mathbf{m} (for its non-canonical orientation convention inherited from \mathbf{n}). In particular, the corner preceding c_1 during the tour of \mathbf{m} (the last corner in the sequence w_1) is a left corner since it is a left corner in \mathbf{n} (indeed, e_1 is a two-way edge in \mathbf{n} incident only to left corners by Lemma 2). Since e_1 is a twist of \mathbf{m} , this implies that c_1 is a right corner of \mathbf{m} (for its non-canonical convention). A similar reasoning shows that c_2 and c_3 are left corners of \mathbf{m} (for its non-canonical convention). Since v is incident to a majority of left corners, the orientation convention at v is the canonical one. Hence c_1, c_2, c_3 are in counterclockwise order around v for the canonical orientation convention of \mathbf{m} , which together with the expression of $w(\mathbf{m})$ shows that v is an intertwined node of flavor \mathbf{B} .

It only remains to prove that the opening of a node of flavor \mathbf{B} and the gluing of flavor \mathbf{B} are inverse operations. The reader might already be convinced of this fact by reasoning in terms of ribbon graphs. Otherwise, the proof (which must deal with some orientation conventions) runs as follows.

We first prove that opening a glued map gives the original map. Let \mathbf{n} be a map with marked leaves v_1, v_2, v_3 , let \mathbf{m} be the map with new vertex v obtained by the corresponding gluing of flavor \mathbf{B} , and let \mathbf{n}' be the map obtained by opening \mathbf{m} at v . It is clear that the graph G underlying \mathbf{n} and \mathbf{n}' is the same and we want to prove that $\mathbf{n} = \mathbf{n}'$ (that is, there exists a set of vertices U' such that flipping U' changes the system of rotation and set of twists of \mathbf{n} to those of \mathbf{n}'). The map \mathbf{m} inherits an orientation convention from \mathbf{n} which might differ from its canonical convention. These two conventions on \mathbf{m} differ by the flipping of a certain subset of vertices U , and gives two different systems of rotations and two sets of twists for \mathbf{m} . By definition, the map \mathbf{n} and \mathbf{n}' have graph G and rotation system and set of twists given by the non-canonical and canonical convention for \mathbf{m} . Observe now that one can get the rotation system and set of twists of \mathbf{n} to those of \mathbf{n}' by flipping the set of vertices U' , where $U' = U$ if $v \notin U$ and $U' = U \setminus \{v\} \cup \{v_1, v_2, v_3\}$ otherwise. Hence, $\mathbf{n} = \mathbf{n}'$.

We now consider a map \mathbf{m} with intertwined vertex v of flavor \mathbf{B} , the map \mathbf{n} with marked leaves obtained from the opening of \mathbf{m} at v , and the map \mathbf{m}' obtained by the gluing of flavor \mathbf{B} . It is clear that the graph G underlying \mathbf{m} and \mathbf{m}' is the same and we want to prove that $\mathbf{m} = \mathbf{m}'$ (that is, there exists a set of vertices U such that flipping U changes the system of rotation and set of twists of \mathbf{m} to those of \mathbf{m}'). The map \mathbf{n} inherits an orientation convention from \mathbf{m} which might differ from its canonical convention. Let v_1, v_2, v_3 be the marked leaves of \mathbf{n} appearing in this order during the tour of \mathbf{n} and let e_1, e_2, e_3 be the incident edges. In the orientation convention C of \mathbf{n} inherited from \mathbf{m} , the corners incident to v_1 and v_2 are right corners, while v_3 is a left corner (see Figure 5). In

the orientation convention C' of \mathbf{n} obtained from the canonical convention by flipping v_1 and v_2 , the corners incident to v_1 and v_2 are right corners, while v_3 is a left corner. This implies that one goes from the convention C to the convention C' by flipping a subset of vertices U not containing v_1, v_2, v_3 . By definition, the maps \mathbf{m} and \mathbf{m}' have graph G and system of rotation and twists inherited respectively from the conventions C and C' on \mathbf{m} . Since the orientation convention of \mathbf{m} and \mathbf{m}' coincide at v (the edges e_1, e_3, e_2 appear in this counterclockwise order around v) and $v_1, v_2, v_3 \notin U$, one gets from the system of rotation and twists of \mathbf{m} to those of \mathbf{m}' by flipping the set of vertices U . Hence, $\mathbf{m} = \mathbf{m}'$.

This concludes the proof that the opening and gluing of flavor **B** are inverse operations. The proofs for the other flavors is similar (see Figure 6 for the definition of the gluings of type **A**, **C**, and **D**). \square

3 Main results.

3.1 The number of precubic unicellular maps.

In this section, we present our main results, which rely on two ingredients. The first one is Proposition 4, which enables us to express the number of precubic unicellular maps of type h with a marked intertwined node in terms of the number of unicellular maps of a smaller type. The second ingredient is the fact (to be discussed in Section 4) that, among maps of type h and fixed size, the average number of intertwined nodes in a map is $2h - 1$.

In order to use Proposition 4, we first need to determine the number of way of choosing non-root leaves in precubic maps.

Lemma 5. *Let $h \in \frac{1}{2}\mathbb{N}$ and let \mathbf{m} be a precubic unicellular map of type h . Then, the number of edges of \mathbf{m} is at least $6h - 1$ and is odd if the type h is an integer and even otherwise. Moreover, if \mathbf{m} has $2m + \mathbb{1}_{h \in \mathbb{N}}$ edges, then it has $m + 1 - 3h - \frac{1}{2}\mathbb{1}_{h \notin \mathbb{N}}$ non-root leaves.*

Proof. Let n_1 and n_3 be the number of vertices of degree 1 and 3 in \mathbf{m} , respectively. One has $n_1 + n_3 = v(\mathbf{m})$ and $n_1 + 3n_3 = 2e(\mathbf{m})$. Moreover, Euler formula gives $v(\mathbf{m}) = e(\mathbf{m}) + 1 - 2h$. Solving this system of equations gives $n_1 = e(\mathbf{m})/2 + 3/2 - 3h$. Since $n_1 \geq 1$ (because the root vertex of a precubic maps is a leaf) this implies the stated conditions on $e(\mathbf{m})$. Moreover the number of non root leaves is $n_1 - 1 = m + 1 - 3h - \frac{1}{2}\mathbb{1}_{h \notin \mathbb{N}}$. \square

Let $h \geq 1$ be an element of $\frac{1}{2}\mathbb{N}$, and let $m \geq 1$ be an integer. We denote by $\mathcal{O}_h(m)$ and $\mathcal{N}_h(m)$ respectively the sets of orientable and non-orientable precubic unicellular maps of type h with $2m + \mathbb{1}_{h \in \mathbb{N}}$ edges, and we denote by $\xi_h(m)$ and $\eta_h(m)$ their cardinalities. From Lemma 5 and Proposition 4, the number $\eta_h^\bullet(m)$ of *non-orientable* unicellular precubic maps of type h with $2m + \mathbb{1}_{h \in \mathbb{N}}$ edges and a marked intertwined node is given by:

$$\eta_h^\bullet(m) = 4 \binom{\ell}{3} \eta_{h-1}(m) + 3 \binom{\ell}{3} \xi_{h-1}(m), \quad (2)$$

where $\ell = m + 4 - 3h - \frac{1}{2}\mathbb{1}_{h \notin \mathbb{N}}$ is the number of non-root leaves in precubic unicellular maps of type $h - 1$ having $2m + \mathbb{1}_{h \in \mathbb{N}}$ edges. 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 **A**, **B**, **C** or **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 **B**, **C** or **D** in order to destroy the orientability).

The keystone of this paper, to be proved in Section 4 is the following result:

Proposition 6. *There exists an involution Φ of $\mathcal{N}_h(m)$ such that for all maps $\mathbf{m} \in \mathcal{N}_h(m)$, the total number of intertwined nodes in the maps \mathbf{m} and $\Phi(\mathbf{m})$ is $4h - 2$. In particular, the average number of intertwined nodes of elements of $\mathcal{N}_h(m)$ is $(2h - 1)$, and one has $\eta_h^\bullet(m) = (2h - 1)\eta_h(m)$.*

It is interesting to compare Equation (2) and Proposition 6 with the analogous results given in [3] for the orientable case. First, the number $\xi_h(m)$ of unicellular orientable maps satisfies

$$\xi_h^\bullet(m) = \binom{\ell}{3} \xi_{h-1}(m),$$

because only the gluings of type **A** are to be considered. Moreover, *each* map of genus h has exactly $2h$ intertwined nodes (yielding $\xi_h^\bullet(m) = 2h\xi_h(m)$), whereas here the quantity $(2h - 1)$ is only an *average value*. For example, Figure 8 shows two maps on the Klein bottle ($h = 1$) which are related by the involution Φ : they have respectively 2 and 0 intertwined nodes. It should also be noted that (contrarily to the situation for the orientable case) certain unicellular maps have no intertwined nodes, hence cannot be obtained by gluing a map of smaller type.

As a direct corollary of Proposition 6 and Equation (2), we can state our main result:

Theorem 7. *The numbers $\eta_h(m)$ of non-orientable precubic unicellular maps of type h with $2m + 1_{h \in \mathbb{N}}$ edges obey the following recursion:*

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

where $\ell = m + 4 - 3h - \frac{1}{2}1_{h \notin \mathbb{N}}$, and where $\xi_h(m)$ is the number of orientable precubic unicellular maps of genus h with $2m + 1_{h \in \mathbb{N}}$ edges, which is 0 if $h \notin \mathbb{N}$, and is given by the following formula otherwise [2]:

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

Explicit formulas for the numbers $\eta_h(m)$ can now be obtained by iterating the recursion given in Theorem 7.

Corollary 8 (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 9 (the case $h \notin \mathbb{N}$). *Let $h \in \{\frac{1}{2}, \frac{3}{2}, \frac{5}{2}, \dots\}$ 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:*

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

From Corollaries 8 and 9 one gets $\eta_h(m) \sim_{m \rightarrow \infty} k_h 4^m m^{3h-3/2}$, where $k_h = \frac{c_h}{6^h \sqrt{\pi}}$ if $h \in \mathbb{N}$, and $k_h = \frac{4^{h-3/2}}{6^{h-1/2}(2h-1)!!}$ otherwise.

Proof of Corollary 8. It follows by induction on h and Equations (3) and (4) that the statement of Equation (5) holds, with the constant c_h defined by the recurrence $c_0 = 0$ (since there is no non-orientable map of type 0) and $c_h = \lambda_{h-1}c_{h-1} + a_{h-1}$, with $\lambda_{h-1} = \frac{4}{2h-1}$ and $a_{h-1} = \frac{3}{(2h-1)(2h-2)!!} = \frac{3}{2^{h-1}(h-1)!(2h-1)}$. The solution of this recurrence is $c_h = \sum_{l=0}^{h-1} a_l \lambda_{l+1} \lambda_{l+2} \dots \lambda_{h-1}$. Now, by definition, the product $a_l \lambda_{l+1} \lambda_{l+2} \dots \lambda_{h-1}$ equals $\frac{3 \cdot 4^{h-1-l}}{2^l!(2l+1)(2l+3)(2l+5)\dots(2h-1)}$. Using the expression $\frac{1}{(2l+1)(2l+3)\dots(2h-1)} = \frac{2^h h!}{2^l h!} \cdot \frac{(2l)!}{2^l l!}$ and reporting it in the sum gives the expression of c_h given in Corollary 8. \square

Proof of Corollary 9. Since for non-integer h we have $\xi_{h-1}(m) = 0$, the first equality is a direct consequence of an iteration of Theorem 7. Therefore the only thing to prove is that the number $\eta_{1/2}(m)$ of precubic maps in the projective plane is 4^{m-1} . This can be done by induction on m via an adaptation of Rémy's bijection [11]. For $m = 1$, we have $\eta_{1/2}(m) = 1$ (the only map with two edges is made of an edge joining the root-leaf to a node and of a twisted loop incident to this node). 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 $2m$ edges has $m-1$ non-root leaves and $4m$ 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

Remark. Before closing this subsection we point out that the bijection *à la Rémy* presented in the proof of Corollary 9 can be adapted to any surface. Such a bijection implies that the number $\eta_h(m)$ of maps on a surface of integer type h satisfies $(m+1-3h)\eta_h(m) = 2(2m-1)\eta_h(m-1)$. This recursion on m imposes the general form of the numbers $\eta_h(m)$:

$$\eta_h(m) = \frac{2m(2m-1)}{m(m-3h+1)} \eta_h(m-1) = \dots = K_h \frac{(2m)!}{m!(m-3h+1)!}$$

for the constant $K_h = \frac{m_{\min}!}{(2m_{\min})!} \eta_h(m_{\min})$, where $m_{\min} = 3h-1$. Similarly, for non-integer type h , one gets $(m-1-3\lfloor h \rfloor)\eta_h(m) = 2(2m-2)\eta_h(m-1)$, hence

$$\eta_h(m) = \frac{4(m-1)}{(m-1-3\lfloor h \rfloor)} \eta_h(m-1) = \dots = K_h \frac{4^m(m-1)!}{(m-1-3\lfloor h \rfloor)!},$$

for the constant $K_h = \frac{1}{4^{3[h]+1} m_{\min}!} \eta_h(m_{\min})$, where $m_{\min} = 3[h] + 1$. Observe however that the approach *à la Rémy* is not sufficient to determine explicitly the value of $\eta_h(m_{\min})$.

3.2 The asymptotic number of general unicellular maps.

In this subsection we derive the asymptotic number of *arbitrary* (i.e., non-necessarily pre-cubic) unicellular maps of given type (the type is fixed, the number of edges goes to infinity). The expression of these asymptotic numbers was already given in [1] in terms of the number of unicellular *cubic maps* (maps with vertices of degree 3) of the same type (see also [4] for the orientable case). Moreover, the number of unicellular cubic maps of type h is easily seen to be $\eta_h(m_{\min})$, where $m_{\min} = 3h - 1 + \frac{1}{2} \mathbb{1}_{h \notin \mathbb{N}}$, hence can be obtained from Corollaries 8 and 9. The goal of this subsection is rather to explain how to adapt the *opening bijections* described in Subsection 2.3 to (almost all) arbitrary unicellular maps.

Let \mathfrak{m} be a unicellular map of type h . The *core* of \mathfrak{m} is the map obtained by deleting recursively all the leaves of \mathfrak{m} (until every vertex has degree at least 2). Clearly, the core is a unicellular map of type h formed by paths of vertices of degree 2 joining vertices of degree at least 3. The *scheme* of \mathfrak{m} is the map obtained by replacing each of these paths by an edge, so that vertices in the scheme have degree at least 3. We say that a unicellular map is *dominant* if the scheme is *cubic* (every vertex has degree 3).

Proposition 10 ([4, 1]). *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} (there is only a finite number of schemes), by observing that these maps are obtained by substituting each edge of the scheme with a path of trees. A generating function approach then easily shows that the schemes with maximum number of edges are the only one contributing to the asymptotic number of unicellular maps. These schemes are precisely the cubic ones.

The opening bijection of Subsection 2.3 can be adapted to dominant unicellular maps as follows. Given a dominant map \mathfrak{m} of type h and scheme \mathfrak{s} , and v an intertwined node of \mathfrak{s} , we can define the opening operation of \mathfrak{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 4): one obtains a dominant map \mathfrak{n} of type $h - 1$ with three distinguished vertices. These vertices are not *any* three vertices: they have to be in *general position* in \mathfrak{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 triples 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 [3]), but rather we state now our asymptotic theorem:

Theorem 11. *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. Therefore,

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

where the constant c_h is defined in Corollary 8.

We recall for comparison that the number of orientable rooted unicellular maps of type h with n edges satisfies $\epsilon_h(n) \sim_{n \rightarrow \infty} \frac{n^{3h - 3/2} 4^n}{12^h h! \sqrt{\pi}}$.

4 The average number of intertwined nodes

In this section we prove Proposition 6 stating that the average number of intertwined nodes among precubic unicellular maps of type h and size m is exactly $(2h - 1)$:

$$\eta_h^\bullet(m) = (2h - 1)\eta_h(m). \tag{6}$$

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 \mathcal{N}_1 , three maps have 2 intertwined nodes, and three maps have none; see Figure 7. As stated in Proposition 6, our strategy to prove Equation (6) is to exhibit an involution Φ from the set $\mathcal{N}_h(m)$ to itself, such that for any given map \mathbf{m} , the total number of intertwined nodes in the maps \mathbf{m} and $\Phi(\mathbf{m})$ is $4h(\mathbf{m}) - 2$. Observe from Figure 7 that the involution Φ cannot be a simple re-rooting of the map \mathbf{m} .




Figure 7: The precubic unicellular maps with 5 edges on the Klein bottle (the root is incident to the unique leaf). Intertwined nodes are indicated as white vertices.

Before defining the mapping Φ , we relate the number of intertwined nodes of a map to certain properties of its twists. Let \mathbf{m} be a (canonically oriented) precubic map, and let e be an edge of \mathbf{m} which is a twist. Let c be the corner incident to e which appears first in the tour of \mathbf{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 8). 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).

Lemma 12. *Let \mathbf{m} be a precubic unicellular map of type $h(\mathbf{m})$, considered with its canonical orientation convention. Then, its numbers $\tau(\mathbf{m})$ of intertwined nodes, $T_{LR}(\mathbf{m})$ of left-to-right twists, and $T_{RL}(\mathbf{m})$ of right-to-left twists are related by:*

$$\tau(\mathbf{m}) = 2h(\mathbf{m}) + T_{RL}(\mathbf{m}) - T_{LR}(\mathbf{m}). \tag{7}$$




Figure 8: Two maps on the Klein Bottle \mathcal{N}_1 and their intertwined nodes (white vertices). The number of twists are (a) $T_{LR}(\mathbf{m}) = 1, T_{RL}(\mathbf{m}) = 1$; (b) $T_{LR}(\mathbf{m}) = 2, T_{RL}(\mathbf{m}) = 0$.

Proof. We first define the *label* of a corner of \mathbf{m} as the element of $\{1, \dots, 2e(\mathbf{m})\}$ indicating the position of appearance of this corner during the tour of the map: the root corner has label 1, the corner that follows it in the tour has label 2, etc... We say that a corner of \mathbf{m} is a *descent* if it is followed, counterclockwise around its vertex, by a corner of smaller or equal canonical label, and that it is an *ascent* otherwise. We let $dsc(\mathbf{m})$ and $asc(\mathbf{m})$ be the total numbers of descents and ascents in \mathbf{m} , respectively. We will now compute the difference $dsc(\mathbf{m}) - asc(\mathbf{m})$ in two different ways: one by summing over edges, the other by summing over vertices (extending the ideas used in [2] for the orientable case).

To each edge $e = (v_1, v_2)$ of \mathbf{m} we associate the two (distinct) corners c_1, c_2 incident to the vertices v_1 and v_2 respectively and following e clockwise around their vertex. Clearly this, creates a partition of the set of corners of \mathbf{m} and we can compute $dsc(\mathbf{m}) - asc(\mathbf{m})$ by adding the contribution of each edge. Let $e = (v_1, v_2)$ be an edge. For $i = 1, 2$ we denote by l_i the label of c_i and by l'_i the label of the other corner incident to v_i and e ($l'_i = l_i$ if v_i is a leaf). Up to exchanging v_1 and v_2 , we can assume that $l_1 < l_2$. We now examine five cases:

- e is the root edge of \mathbf{m} ; Figure 9(a). Since v_1 is a leaf, the edge e is two-ways hence not a twist and incident to left-corners only. We get $l_1 = l'_1 = 1, l_2 = 2e(\mathbf{m})$, and $l'_2 = 2$. Hence both c_1 and c_2 are descents.
- e is not the root edge, is not a twist, and is two-ways; Figure 9(b). In this case, we know by Lemma 2 that all the corners are *left*, from which $(l'_1, l'_2) = (l_2 + 1, l_1 + 1)$. Therefore c_1 is an ascent and c_2 is a descent.
- e is not the root edge, is not a twist, and is one-way; Figure 9(c). In this case, we have $(l'_1, l'_2) = (l_2 - \epsilon, l_1 + \epsilon)$, where $\epsilon = 1$ or -1 according to whether c_1 is a left or a right corner. In both cases, c_1 is an ascent and c_2 is a descent (observe that v_1 cannot be a leaf by Lemma 2, hence $l_1 \neq l'_1$).
- e is a right-to-left twist; Figure 9(d). In this case, $l_1 < l_2$ implies that e is followed from v_1 to v_2 . By definition of a right-to-left twist, $l'_1 < l_1$ hence c_1 is an descent. Moreover, $l_2 = l_1 + 1$ and $l'_2 = l'_1 + 1$, therefore c_2 is also a descent.
- e is a left-to-right twist; Figure 9(e). This case is similar to the previous one. The corner c_1 is necessarily ascent, and since $l_2 = l_1 + 1$ and $l'_2 = l'_1 + 1$, the corner c_2 is also an ascent.




Figure 9: The five cases of the proof of Lemma 12.

Expressing the difference $\text{dsc}(\mathbf{m}) - \text{asc}(\mathbf{m})$ as a sum over all edges of \mathbf{m} , we obtain from the five cases above:

$$\text{dsc}(\mathbf{m}) - \text{asc}(\mathbf{m}) = 2 + 0 + 0 + 2T_{\text{RL}}(\mathbf{m}) - 2T_{\text{LR}}(\mathbf{m}) = 2\left(1 + T_{\text{RL}}(\mathbf{m}) - T_{\text{LR}}(\mathbf{m})\right).$$

Using the fact that $\text{dsc}(\mathbf{m}) + \text{asc}(\mathbf{m}) = 2e(\mathbf{m})$, we obtain the total number of descents in \mathbf{m} which is $\text{dsc}(\mathbf{m}) = e(\mathbf{m}) + 1 + T_{\text{RL}}(\mathbf{m}) - T_{\text{LR}}(\mathbf{m})$.

Now, there is another way of counting the descents. Indeed, since by definition each non-intertwined node has exactly one descent, and each intertwined node has exactly two of them, one gets: $\text{dsc}(\mathbf{m}) = 2\tau(\mathbf{m}) + (v(\mathbf{m}) - \tau(\mathbf{m}))$. Solving for $\tau(\mathbf{m})$ and using the previous expression for $\text{dsc}(\mathbf{m})$ gives

$$\tau(\mathbf{m}) = \text{dsc}(\mathbf{m}) - v(\mathbf{m}) = e(\mathbf{m}) + 1 - v(\mathbf{m}) + T_{\text{RL}}(\mathbf{m}) - T_{\text{LR}}(\mathbf{m}).$$

The lemma then follows by applying Euler’s formula. □

We now define the promised mapping Φ averaging the number of intertwined nodes. Let \mathbf{m} be a unicellular precubic map on a non-orientable surface. We consider the canonical orientation convention for the map \mathbf{m} , which defines a rotation system and set of twists. The set of twists is non-empty since the map \mathbf{m} lives on a non-orientable surface. By cutting every twist of \mathbf{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 $\widehat{\mathbf{m}}$, can have several connected components and each component (which is a map with buds) can have several faces; see Figure 10. We set a convention for the direction in which one *turns around a face* of $\widehat{\mathbf{m}}$: the edges are followed in such a way that every corner is left (this is possible since $\widehat{\mathbf{m}}$ has no twist). For any bud b of $\widehat{\mathbf{m}}$, we let $\sigma(b)$ be the bud following b when turning around the face of $\widehat{\mathbf{m}}$ containing b . Clearly, the mapping σ is a permutation on the set of buds. We now define $\Phi(\mathbf{m})$ to be the graph with rotation system and twists obtained from $\widehat{\mathbf{m}}$ by gluing together into a twist the buds $\sigma(b)$ and $\sigma(b')$ for every pair of buds b, b' forming a twist of \mathbf{m} . The mapping Φ is represented in Figure 10.




Figure 10: A unicellular map \mathbf{m} and its image by the mapping Φ . The twists are indicated by (partially) dotted lines, while the map $\widehat{\mathbf{m}}$ is represented in solid lines.

Before proving that $\Phi(\mathbf{m})$ is a unicellular map, we set some additional notations. We denote by k the number of twists of \mathbf{m} and we denote by $w(\mathbf{m}) = w_1 w_2 \cdots w_{2k+1}$ the sequence of corners encountered during the tour of \mathbf{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 \mathbf{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 $\widehat{\mathbf{m}}$ 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 $w'_i = \overline{w_i}$ otherwise (where $\overline{w_i}$ is the *mirror* sequence of w_i obtained by reading w_i backwards). We identify the buds of $\widehat{\mathbf{m}}$ (i.e. the half-twists of \mathbf{m} or $\Phi(\mathbf{m})$) with the integers in $\{1, \dots, 2k\}$ by calling i the bud following the sequence of corners w'_i around the faces of $\widehat{\mathbf{m}}$. This labelling is indicated in Figure 10. We will now consider the permutation σ as a permutation on $\{1, \dots, 2k\}$ and we denote $r = \sigma^{-1}(1)$. The map in Figure 10 gives $\sigma = (1, 8, 13, 2, 9, 14, 3, 10)(4, 11, 6, 5)(7, 12)$ and $r = 10$. We first prove a technical lemma.

Lemma 13. *The permutation σ maps odd to even integers. In particular, $r = \sigma^{-1}(1)$ is even.*

Proof. By Lemma 2, all twists of \mathbf{m} are one-way. Hence, every bud of $\widehat{\mathbf{m}}$ is incident both to a left corner and to a right corner of \mathbf{m} . The lemma therefore follows from the fact that left and right corners of \mathbf{m} belong to the sequences w'_i for i odd and i even respectively. \square

We are now ready to prove that $\Phi(\mathbf{m})$ is unicellular and a little more. In the following, we denote by \overline{i} the representative of an integer i modulo $2k$ belonging to $\{1, \dots, 2k\}$.

Lemma 14. *The embedded graph $\Phi(\mathbf{m})$ is a unicellular map. Moreover, the rotation system and set of twists of $\Phi(\mathbf{m})$ inherited from \mathbf{m} correspond to the canonical orientation convention of $\Phi(\mathbf{m})$. Lastly, the sequence of corners encountered during the tour of $\Phi(\mathbf{m})$ reads $v_1 v_2 \dots v_{2k+1}$, where the subsequences v_i separated by twist traversals are given by $v_i = w_{\sigma(\overline{r+1-i})}$ for all $i = 1, \dots, 2k$, and $v_{2k+1} = w_{2k+1}$.*

Proof. We consider, as above, the map \mathbf{m} with its canonical orientation convention and the map $\Phi(\mathbf{m})$ with the orientation convention inherited from \mathbf{m} . We denote by α the (fixed-point free) involution on $\{1, \dots, 2k\}$ corresponding to the twists of \mathbf{m} . That is to say, $\alpha(i) = j$ if the half-edges i, j form a twist of \mathbf{m} . We also denote by $\beta = \sigma \alpha \sigma^{-1}$ the involution corresponding to the twists of $\Phi(\mathbf{m})$.

Fact 1: For $i \in \{1, \dots, 2k\}$, $\alpha(i) = i + 1$ if i is odd (hence, $\alpha(i) = i - 1$ if i is even). Similarly $\beta(i) = \overline{i + 1}$ if i is even (hence $\beta(i) = \overline{i - 1}$ if i is odd).

To prove Fact 1, recall that w_1, \dots, w_n denote the sequences of corners, encountered in that order during the tour of \mathbf{m} . If $i \in \{2, 3, \dots, 2k\}$ is odd (resp. even), then the sequence of left corners $w_i = w'_i$ (resp. right corners $w_i = \overline{w'_i}$) goes from the bud $\sigma^{-1}(i)$ to the bud i (resp. from the bud i to the bud $\sigma^{-1}(i)$) during the tour of \mathbf{m} ; see Figure 11. Hence, if $i \in \{1, 2 \dots 2k\}$ is odd, the twist of \mathbf{m} traversed between w_i and w_{i+1} is made of the half-twists i and $i + 1$, while if i is even it is made of the half-twists $\sigma^{-1}(i)$ and $\sigma^{-1}(\overline{i + 1})$. From the odd case, one gets $\alpha(i) = i + 1$ if i is odd. From the even case, one gets $\alpha(\sigma^{-1}(i)) = \sigma^{-1}(\overline{i + 1})$ if i is even. That is, $\beta(i) \equiv \sigma \alpha \sigma^{-1}(i) = \overline{i + 1}$ if i is even.

We now denote by $v_1 v_2 v_3 \dots v_{\ell+1}$ the sequence of corners encountered by following the edges of $\Phi(\mathbf{m})$ starting and ending at the root corner (tour of the face of $\Phi(\mathbf{m})$ containing the root), where the subsequences v_i and v_{i+1} are separated by a twist traversal. Clearly,




Figure 11: Proof of Fact 1.




Figure 12: Proof of Fact 2.

$v_1 = w_1$, $v_{\ell+1} = w_{2k+1}$ and for $i = 1 \dots \ell + 1$ the corners in v_i are left corners of $\Phi(\mathbf{m})$ if and only if i is odd. For $i = 1 \dots \ell$, we denote $v'_i = v_i$ if i is odd and $v'_i = \bar{v}_i$ otherwise, so that each of the sequences v'_1 belongs to $\{w'_1, \dots, w'_{2k}\}$. For $i = 1 \dots \ell$, we denote by $\phi(i)$ the bud following v'_i around the faces of $\widehat{\mathbf{m}}$. Then, the same reasoning as above (see Figure 12) proves:

Fact 2: For $i = 1, \dots, \ell - 1$, $\beta(\phi(i)) = \phi(i + 1)$ if i is odd, and $\sigma\beta\sigma^{-1}(\phi(i)) = \phi(i + 1)$ if i is even.

Fact 3: If $\phi(i) = \sigma(j)$ for certain integers $1 \leq i < \ell$ and $1 < j \leq 2k$ of different parity, then $\phi(i + 1) = \sigma(\overline{j - 1})$.

The Fact 3 is easily proved by the following case analysis. If i is odd, then

$$\phi(i + 1) = \beta(\phi(i)) = \beta(\sigma(j)) = \sigma\alpha(j) = \sigma(j - 1),$$

where the first and last equalities are given by Fact 2 and Fact 1 respectively. Similarly, if i is even

$$\phi(i + 1) = \sigma\beta\sigma^{-1}(\phi(i)) = \sigma\beta\sigma^{-1}(\sigma(j)) = \sigma\beta(j) = \sigma(\overline{j - 1}).$$

We now consider the relation $\phi(1) = 1 = \sigma(r)$ and recall that 1 and r are of different parity by Lemma 13. Then Fact 3 implies by induction that $\phi(i) = \sigma(\overline{r + 1 - i})$ for $i = 1 \dots \ell$. This proves that $v'_i = w'_{\sigma(\overline{r + 1 - i})}$ for $i = 1 \dots \ell$. Since i and $\sigma(\overline{r + 1 - i})$ have the same parity (by Lemma 13), this also gives $v_i = w_{\sigma(\overline{r + 1 - i})}$ for $i = 1 \dots \ell$. In particular, for $i = \ell$, one gets $v_\ell = w_{\sigma(\overline{r + 1 - \ell})}$. Moreover, by definition $v_{\ell+1} = w_{2k+1}$, hence $v_\ell = w_{\sigma\beta(r)}$. Hence, $\beta(r) = \overline{r + 1 - \ell}$. Since r is even, $\beta(r) = \overline{r + 1}$ (by Fact 1), hence $\ell = 2k$.

The sequence $v_1 v_2 \dots v_{2k+1}$ contains all the corners of $\Phi(\mathbf{m})$. Hence, $\Phi(\mathbf{m})$ is a unicellular map. Moreover, a corner is left for the map \mathbf{m} (resp. the map $\Phi(\mathbf{m})$ considered with its orientation convention inherited from \mathbf{m}) if and only if it belongs to a sequence w_i (resp. $v_i = w_{\sigma(\overline{r + 1 - i})}$) for an odd integer i . Since i and $\sigma(\overline{r + 1 - i})$ have the same parity a corner is left for \mathbf{m} if and only if it is left for $\Phi(\mathbf{m})$. This shows that the orientation convention of $\Phi(\mathbf{m})$ inherited from \mathbf{m} is the canonical convention of $\Phi(\mathbf{m})$. \square

We now make the final strike by considering the action of Φ on the set $\mathcal{N}_h(m)$ of non-orientable maps of type h .

Proposition 15. *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 \mathbf{m} in $\mathcal{N}_h(m)$, the total number of intertwined nodes in the maps \mathbf{m} and $\Phi(\mathbf{m})$ is $4h - 2$.*

Proof. Clearly, the maps \mathbf{m} and $\Phi(\mathbf{m})$ 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 \mathbf{m} , the embedded graphs $\widehat{\mathbf{m}}$ and $\widehat{\Phi(\mathbf{m})}$ are equal; this is because the canonical rotation system and set of twists of \mathbf{m} and $\Phi(\mathbf{m})$ coincide. In particular, the permutation σ on the half-twists of \mathbf{m} can be read from $\Phi(\mathbf{m})$. Hence, the twists of \mathbf{m} are easily recovered from those of $\Phi(\mathbf{m})$: the buds i and j form a twist of \mathbf{m} if $\sigma(i)$ and $\sigma(j)$ form a twist of $\Phi(\mathbf{m})$.

We now proceed to prove that the total number of intertwined nodes in \mathbf{m} and $\Phi(\mathbf{m})$ is $4h - 2$. By Lemma 12, this amounts to proving that $T_{\text{LR}}(\mathbf{m}) - T_{\text{RL}}(\mathbf{m}) + T_{\text{LR}}(\Phi(\mathbf{m})) - T_{\text{RL}}(\Phi(\mathbf{m})) = 2$. Since \mathbf{m} and $\Phi(\mathbf{m})$ both have k twists, $T_{\text{LR}}(\mathbf{m}) - T_{\text{RL}}(\mathbf{m}) + T_{\text{LR}}(\Phi(\mathbf{m})) - T_{\text{RL}}(\Phi(\mathbf{m})) = 2(T_{\text{LR}}(\mathbf{m}) + T_{\text{LR}}(\Phi(\mathbf{m})) - k)$. Hence, we have to prove $T_{\text{LR}}(\mathbf{m}) + T_{\text{LR}}(\Phi(\mathbf{m})) = k + 1$.

Let i be a bud of $\widehat{\mathbf{m}}$, let t be the twist of \mathbf{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 \mathbf{m} is left-to-right if and only if c appears before c' during the tour of \mathbf{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 $i = r$ (and these two possibilities are disjoint since $\sigma(r) = 1$).

Let us now examine under which circumstances the bud $\sigma(i)$ is part of a left-to-right twist of $\Phi(\mathbf{m})$. 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 c' belongs to w_{2k+1}). By Lemma 14, $w_{\sigma(i)} = v_{\overline{r+1-i}}$ for $i = 1 \dots 2k$. Therefore, the twist t' of $\Phi(\mathbf{m})$ containing $\sigma(i)$ is left-to-right (for $\widehat{\mathbf{m}}$) if and only if $\overline{r+1-i} < \overline{r+1-\sigma(i)}$ or $\sigma(i) = r$ (and these two possibilities are disjoint since $\sigma(i) = r$ implies $r+1-\sigma(i) = 1$).

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

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

where $\delta(i) = \mathbb{1}_{i < \sigma(i)} + \mathbb{1}_{\overline{r+1-i} < \overline{r+1-\sigma(i)}}$ 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 $T_{\text{LR}}(\mathbf{m}) + T_{\text{LR}}(\Phi(\mathbf{m})) = k + 1$. \square

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

Proof of Proposition 6. 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

- [1] O. Bernardi and J. Rué. Enumerating simplicial decompositions of surfaces with boundaries. Submitted.
- [2] G. Chapuy. A new combinatorial identity for unicellular maps, via a direct bijective approach. Submitted. Extended abstract in Discrete Math. Theor. Comput. Sci. Proc., FPSAC'09.
- [3] G. Chapuy. The structure of unicellular maps, and a connection between maps of positive genus and planar labelled trees. *Probab. Theory and Related Fields*, 147:415–447, 2010.
- [4] G. Chapuy, M. Marcus, and G. Schaeffer. A bijection for rooted maps on orientable surfaces. *SIAM J. Discrete Math.*, 23:1587–1611, 2009.
- [5] I. Goulden and D. Jackson. Maps in locally orientable surfaces and integrals over real symmetric matrices. *Canad. J. Math*, 48:569–584, 1997.
- [6] I. Goulden and A. Nica. A direct bijection for the Harer-Zagier formula. *J. Combin. Theory Ser. A*, 111:224–238, 2005.
- [7] J. Harer and D. Zagier. The Euler characteristic of the moduli space of curves. *Invent. Math.*, 85:457–485, 1986.
- [8] B. Lass. Démonstration combinatoire de la formule de Harer-Zagier. *C. R. Acad. Sci. Paris*, 333:155–160, 2001.
- [9] M. Ledoux. A recursion formula for the moments of the Gaussian orthogonal ensemble. *Ann. Inst. Henri Poincaré Probab. Stat.*, 45:754–769, 2009.
- [10] B. Mohar and C. Thomassen. *Graphs on surfaces*. Johns Hopkins University Press, Baltimore, 2001.
- [11] J.-L. Rémy. Un procédé itératif de dénombrement d'arbres binaires et son application à leur génération aléatoire. *RAIRO Inform. Théor.*, 19:179–195, 1985.
- [12] T. R. S. Walsh and A. B. Lehman. Counting rooted maps by genus. I. *J. Combin. Theory Ser. B*, 13:192–218, 1972.