

A generic method for bijections between blossoming trees and planar maps

Marie Albenque* Dominique Poulalhon†

Submitted: May 16, 2013; Accepted: May 14, 2015; Published: Jun 3, 2015
Mathematics Subject Classifications: 05A19; 05A15; 05C05; 05C10; 05C38; 05C40; 06B99

Abstract

This article presents a unified bijective scheme between planar maps and blossoming trees, where a blossoming tree is defined as a spanning tree of the map decorated with some dangling half-edges that enable to reconstruct its faces. Our method generalizes a previous construction of Bernardi by loosening its conditions of application so as to include *annular maps*, that is maps embedded in the plane with a root face different from the outer face.

The bijective construction presented here relies deeply on the theory of α -orientations introduced by Felsner, and in particular on the existence of minimal and accessible orientations. Since most of the families of maps can be characterized by such orientations, our generic bijective method is proved to capture as special cases many previously known bijections involving blossoming trees: for example Eulerian maps, m -Eulerian maps, non-separable maps and simple triangulations and quadrangulations of a k -gon. Moreover, it also permits to obtain new bijective constructions for bipolar orientations and d -angulations of girth d of a k -gon.

As for applications, each specialization of the construction translates into enumerative by-products, either via a closed formula or via a recursive computational scheme. Besides, for every family of maps described in the paper, the construction can be implemented in linear time. It yields thus an effective way to encode or sample planar maps.

In a recent work, Bernardi and Fusy introduced another unified bijective scheme; we adopt here a different strategy which allows us to capture different bijections. These two approaches should be seen as two complementary ways of unifying bijections between planar maps and decorated trees.

Keywords: Planar maps; bijection; trees; graph orientation; minimal orientation

*LIX – CNRS, École Polytechnique, France – albenque@lix.polytechnique.fr

†LIAFA – CNRS, Université Paris Diderot, France – poulalhon@liafa.univ-paris-diderot.fr

Introduction

The enumeration of planar maps was initiated in the 60's with the pioneering work of Tutte [38]. To obtain enumeration formulas for planar maps, Tutte uses some properties about their decomposition to write equations satisfied by their generating series. The equations thus obtained are quite complicated, in particular some additional parameters (known as *catalytic variables*) have usually to be introduced to write them. The work of Tutte is a computational *tour de force*, since he managed, in most cases, to solve these equations and to obtain closed (and particularly simple) formulas for numerous families of maps.

This method turned out to be extremely versatile and can be applied to many different models with only slight modifications. Furthermore, the structure of the equations and of their solutions is now better understood and some standard methods such as the “quadratic method” [23, sec.2.9] and its extensions [7] are available to compute them in standard cases. Tutte's approach, however, gives little insight about the combinatorial structure of maps and in particular produces no explanation of the reason why the formulas obtained should be so simple.

Since then, successful ideas have been used to rederive and generalize the results of Tutte, including work based on matrix integrals [37, 12], algebraic approach [25] and bijective constructions [15, 35]. The latter yield an elegant way to rederive the formulas of Tutte but they also provide tools for a better understanding of the combinatorial structure of maps: they produce for example an efficient way to encode them by words of algebraic languages and hence to sample them efficiently [35, 33]. This led to establish new conjectures about the asymptotic behaviour of random maps, which gave birth to a new field of research that has been extremely active since (see for instance [14, 28, 26, 30, 27]). It is noteworthy that a key ingredient in all these works about the convergence of random planar maps is Schaeffer's bijection [35] between quadrangulations and well-labelled trees, where the distance in the quadrangulation between a vertex and the root is encoded by the label of the corresponding vertex in the tree, or one of its generalizations [10, 13, 29].

Let us now focus on those bijective proofs. Formulas obtained by Tutte and its successors suggest that maps could be interpreted as trees with some decorations. After initial work in that direction by Cori and Vauquelin [15], Schaeffer [35] drew new attention on this field by obtaining numerous bijective constructions. This founding work was followed by a series of papers dealing with various families of maps: a non-exhaustive list includes maps with prescribed degree sequence [34, 9, 10], maps endowed with a physical model [8, 11] or with connectivity constraints [33]. Each of these bijections appears as an ad-hoc explanation of the known enumeration formula, but they present strong similarities, which calls for a unified bijective theory. An important step in that direction has been achieved in [4, 5], where a “master bijection” is introduced in order to see many constructions as special cases of a common construction. The main purpose of the present paper is to present a different attempt in unifying the bijective constructions, in particular so as to include some bijections that are not captured by the work of Bernardi and Fusy.

Allow us to be slightly more precise. The first bijections obtained rely on the existence

of a canonical spanning tree of the map [34, 35, 9] or of its quadrangulation [15, 35, 10]. As emphasized by Bernardi [3], a map endowed with a spanning tree can also be viewed as a map endowed with an orientation of its edges with specific properties. The latter point of view appears to be more suitable to unify and generalize the constructions. In particular, the master bijection defined in [4] and [5] is based on a generalized notion of orientations. This construction includes as special cases many previously known bijections, but unfortunately not all of them and in particular not the case of simple triangulations [33] and quadrangulations [20, ch.3]. We would like to emphasize that in both these cases, the tree obtained is simply a spanning tree of the original map. The master bijection, in contrast, produces a spanning tree of the superimposition of a planar map, its dual and their (common) quadrangulation.

The ground result of our paper is to present a new bijective scheme which relies on an orientation of the map and yields a spanning tree of the map with some decorations that allow to reconstruct facial cycles. It generalizes the result of [3] by loosening the rooting conditions. In particular it enables to deal with *annular maps* (that is, rooted planar maps with a marked face) such as triangulations of a p -gon [33]. Notably all the previous bijective constructions that involve a spanning tree of the map are captured by our generic scheme and, moreover, we obtain new bijections for plane bipolar orientations and d -angulations of a p -gon with girth d . Besides, the first bijective proof of a well-known theorem by Hurwitz on products of transpositions in the symmetric group has been obtained recently by Duchi, Schaeffer and the second author [17] using this generic scheme.

Bijective proofs appear often as an a posteriori enlightening explanation of a simple enumerative formula. In fact, the formula is used as a guide to construct the “simplest” objects that it enumerates: the right objects to consider can be seen as its combinatorial translation. Here, remarkably, the satisfying orientations are often natural enough so that they can be guessed even if a formula is not available.

Another important feature of our generic scheme is its constructive character: given a blossoming tree, the corresponding map can be computed in linear time. Reciprocally given a map endowed with the appropriate orientation, the corresponding blossoming tree can be computed in quadratic time by a generic algorithm. In fact, for all the families of maps considered in this paper, ad-hoc algorithms can be designed to compute the blossoming tree in linear time. This was known to be true for rooted maps [3], and one of our main contributions is to design a linear-time algorithm that computes the blossoming tree of a d -angulation of a p -gon.

To conclude, let us mention three perspectives to continue this work. The bijective method we develop relies deeply on orientations and, algorithmically speaking, takes as input a map endowed with a specific orientation. An algorithm by Felsner [18] (see also [20, p.56]) ensures that, for a fixed map with n vertices and a prescribed sequence α of outdegrees, an α -orientation (if any exists) can be computed by a generic algorithm of complexity $\Theta(n^{3/2})$, if the map has bounded maximal outdegree. For various families of maps, linear time algorithms do exist, but it is still an open problem to design such an al-

gorithm for d -angulations when $d \geq 5$ (unlike simple triangulations and quadrangulations, see for instance [20, ch.2]).

Secondly, almost all models of maps appear now as special cases either of our generic scheme or of the master bijections of Bernardi and Fusy. Nevertheless, a few models are still not captured; this in particular the case of models with “matter” such as the Ising model, for which some bijections with blossoming trees are known [8]. Some additional work is needed to either generalize one of those schemes or to come up with an alternative approach.

Lastly, as mentioned above, the scaling limit of random plane maps has been a very active area of research in the last years. So far, it has been proved that for $p = 3$ or p even, the limit of (properly rescaled) p -angulations is the so-called “Brownian map” [30, 27]. It is widely believed that all the reasonable families of maps – which includes for instance p -angulations for p odd or maps with constraints on their connectivity – belong to the same universality class, or in other words, should converge to the same limit object. A first result in this direction about simple triangulations and quadrangulations has been obtained very recently by Addario-Berry and the first author [1]. The proof of their result relies on the bijections of [33] and [20] and on a way to interpret the distances in the map on the corresponding blossoming tree. It would be a major breakthrough to generalize their result to all the maps captured by our scheme.

Outline. In Section 1, we gather definitions about maps and their orientations and recall the fundamental result of Felsner about uniqueness of minimal α -orientations (Theorem 1.1). We introduce *blossoming maps* in Section 2 so as to describe and prove our bijective scheme along with some remarks about its complexity.

In Section 3, previous bijections obtained for Eulerian maps and general maps (Subsection 3.1), m -Eulerian maps (Subsection 3.2) and non-separable maps (Subsection 3.3) are rederived via our bijective technique. A new bijection between bipolar orientations and some triples of paths is obtained in Section 4.

Sections 5 and 6 are devoted to d -angulations of a p -gon. More precisely, Section 5 describes $\frac{d}{d-2}$ -orientations and the bijection between p -gonal d -fractional forests and p -gonal d -angulations as well as enumerative consequences, while Section 6 focuses on the description and the proof of the linear time opening algorithm in that setting.

1 Maps and orientations

1.1 Planar maps

A *planar map* is a proper embedding of a connected graph in the sphere, where *proper* means that edges are smooth simple arcs which meet only at their endpoints. Two planar maps are identified if they can be mapped one onto the other by a homeomorphism that preserves the orientation of the sphere. *Edges* and *vertices* of a map are the natural counterparts of edges and vertices of the underlying graph. The *faces* of a map are the connected components of the complementary of the embedded graph. The embedding fixes the cyclical order of edges around each vertex, which defines readily a *corner* as a

couple of consecutive edges around a vertex (or, equivalently, around a face). Corners may also be viewed as incidences between vertices and faces. The *degree* of a vertex or a face is defined as the number of its corners. In other words, it counts incident edges with multiplicity 2 for each loop (in the case of vertex degree) or for each bridge (in the case of face degree).

In these definitions, vertices and faces play a similar role. It is often useful to exchange them and to consider the *dual* M^* of a given map M , whose vertices correspond to faces of M and faces to vertices of M . Edges are somehow unchanged: each edge e of M corresponds to an edge of M^* that is incident to the same vertices and faces as e , see Fig.1(a).




Figure 1: Example of classical constructions on a map. The original map is drawn in dashed grey lines with white vertices. Its dual map, quadrangulation and radial map are drawn in plain colored lines.

A planar map is said to be *d-regular* if all its vertices have degree d . Dually, a planar map is called a *d-angulation* if all its faces have degree d ; the terms *triangulation*, *quadrangulation* and *pentagulation* correspond respectively to the cases where $d = 3, 4, 5$. The following classical and useful construction associates a quadrangulation to each planar map M . Let us say that vertices of M are white; add a black vertex in each face of M , and an edge in each corner of M between the corresponding white and black vertices. This produces a triangulation with bicolored vertices. Keeping only the additional edges leads to a quadrangulation \mathcal{Q}_M which is called *the* quadrangulation of M , see Fig.1(b). Its dual map \mathcal{R}_M is called the *radial map* of M , see Fig.1(c).

A *plane map* is a proper embedding of a connected graph in the plane; its unique unbounded face is called the *outer face*, and all the other faces are called *inner* faces. Vertices and edges are called *outer* or *inner* depending on whether they are incident to the outer face or not. Observe that a plane map is in fact a planar map with a distinguished marked face.

A planar or plane map is said to be (*corner*)-*rooted* if one of its corners is distinguished. The corresponding vertex and face are called *root vertex* and *root face*. In the figures, the root corner is indicated by a double arrow and the root vertex by a square vertex, see Fig.2 for instance. The *root edge* is defined as the edge that follows the root corner

in clockwise order around the root vertex. The usual convention is to associate to each rooted planar map the rooted plane map in which the root and the outer faces coincide. However in this work, plane maps are allowed to have one root face different from the outer face (in the literature, planar maps for which the root face is different from the outer face are sometimes called *annular maps*). Some weaker rootings will sometimes be also considered by only pointing either a root vertex, a root edge or a root face. In the latter case, vertices or edges incident to the root face are called *root vertices* or *root edges*.

A *plane tree* is a planar map with a single face; its vertices are called *nodes* and *leaves* depending on their degree being greater than or equal to one. A *planted tree* is a plane tree rooted at a leaf. Observe that usual “planar trees” or “ordered trees” are obtained from planted trees by deleting their root leaf.

1.2 Orientations

This section gathers definitions and fundamental results about orientations of planar maps. The terminology and convention are not completely standard and we emphasize the differences when needed.

An *orientation* of a map is the choice of a direction for each of its edges. The *indegree* or *outdegree* of a vertex v , denoted $\text{in}(v)$ or $\text{out}(v)$, is the number of edges oriented inwards or outwards v . Let M be a planar map, V the set of its vertices, and let $\alpha : V \rightarrow \mathbb{N}$ be an application which associates a natural number to each vertex of the map. An α -*orientation* – as introduced by Felsner in [18] – is an orientation of M such that for each vertex v in V , $\text{out}(v) = \alpha(v)$. If such an orientation exists, α is said to be *feasible*. An orientation of a corner-, vertex- or face-rooted map is said to be *accessible* if for any vertex v , there exists an oriented path (see below) from v to the root vertex (or to one of the root vertices, in the case of a face rooting).

An *oriented path* is an alternating sequence $(v_0, e_1, v_1, \dots, v_{\ell-1}, e_\ell, v_\ell)$ of incident vertices and edges in which each edge e_i is oriented from v_{i-1} to v_i . An *oriented cycle* (also called a *circuit*) is defined accordingly. The (canonical) embedding of plane maps enables to define *clockwise* and *counterclockwise cycles* as simple oriented cycles with the outer face respectively on their left and on their right. Observe that the orientation obtained after reverting all the edges of a given oriented cycle is still an α -orientation. In fact, all the α -orientations of a map M can be obtained by a sequence of such *flips*, see [18]. In particular, it implies that either all or none α -orientations of M are accessible. In the former case, α is said *accessibly feasible*. Moreover:

Theorem 1.1 (Felsner [18]). *Let M be a plane map and α be a feasible function on its vertices. Then, there exists a unique α -orientation without counterclockwise cycles.*

The relevance of this theorem for our purpose is to associate *canonically* to any given feasible (M, α) one specific α -orientation. From now on, we call *minimal* any orientation without counterclockwise cycles.

Remark 1.2. Let us mention that the result of Felsner is in fact much stronger. He proves indeed that the set of α -orientations of M can be endowed with a lattice structure, where

the cover relation corresponds essentially to the flip of an oriented cycle. The (unique) minimum element of this lattice is the minimal α -orientation, hence its name.

The k -expanded version of a plane map M is defined as the plane map where each edge of M has been replaced by k copies. A k -fractional orientation of M is defined in [4] as an orientation of the k -expanded map of M , with the additional property that two copies of the same edge cannot create a counterclockwise cycle. It is conveniently considered as an orientation of M in which edges can be partially oriented in both directions and the in- or out-degree of a vertex v (that can now be fractional) is defined as the in- or out-degree of v in the k -expanded map, divided by k . In this setting, a *saturated edge* is an edge which is totally oriented in the same direction, an *oriented path* is a path in which each edge is at least partially oriented in the considered direction. The notions of clockwise or counterclockwise cycles, of minimality and of accessibility follow. Figure 2 shows the two possible representations of a given 2-fractional orientation, with two counterclockwise cycles, and the corresponding minimal 2-fractional orientation.




Figure 2: A rooted plane map endowed with accessible 2-fractional orientations.

2 A generic bijective scheme for maps endowed with a minimal orientation

2.1 Blossoming maps and closure

Definition 2.1. A blossoming map is a plane map, in which each outer corner can carry a sequence of opening or closing stems (in the literature, opening and closing stems are sometimes referred to as buds and leaves).

The cyclic contour word of a blossoming map is the word on $\{e, b, \bar{b}\}$, which encodes the cyclic clockwise order of edges and stems along the border of the outer face with e coding for an edge and b and \bar{b} for opening and closing stems, see Fig.3(a).

A local closure of a blossoming map is a substitution of a factor $be^*\bar{b}$ by the letter e in its contour word, where e^* denotes any sequence of e (possibly empty), see Fig.3(b). In terms of maps, it corresponds to the creation of a new edge (and hence a new face)




Figure 3: Closure of a blossoming map. Opening stems are represented by plain green arrows, and closing stems by reverse blue arrows. Factors to be substituted by e are underlined in the contour word.

by merging an opening stem with the following closing stem (provided that there is no other stem in between) in clockwise order around the border of the outer face. The new edge is canonically oriented from the opening vertex to the closing vertex, with the new bounded face on its right. If several local closure operations are possible on a blossoming map, performing all of them in either order yields the same result, hence iterating such local closures produces eventually a unique object:

Definition 2.2. *The closure \bar{M} of a blossoming map M is the (possibly blossoming) map obtained after iterating as many local closure operations as possible, see Fig.3(c). When only a subset of local closures is performed, the map obtained is called a partial closure of M . The edges created during local closures operations are called closure edges.*

In particular, the closure of a blossoming map with an equal number of opening and closing stems is a (non-blossoming) map. Since closure edges are canonically oriented, if a blossoming map is endowed with an orientation (possibly k -fractional), so is its closure. Moreover, considering opening and closing stems respectively as outgoing and incoming (half-)edges, in- and out-degrees are preserved. Since all the closures are performed in clockwise direction around the map, no counterclockwise cycle can be created during a local closure operation. Consequently if the initial orientation is minimal, then so is its closure. Accessibility is preserved as well.

The most interesting special case is the one of a rooted plane tree, endowed with an accessible orientation, which, in the classical non-fractional setting, implies that edges are oriented towards the root vertex. Examples of a blossoming tree and of its closure are given in Figs. 4(c) and 4(a). See also Fig. 7(c) and 7(a) for a 2-fractional example.

The aim of the next section is to provide an inverse construction of the closure.

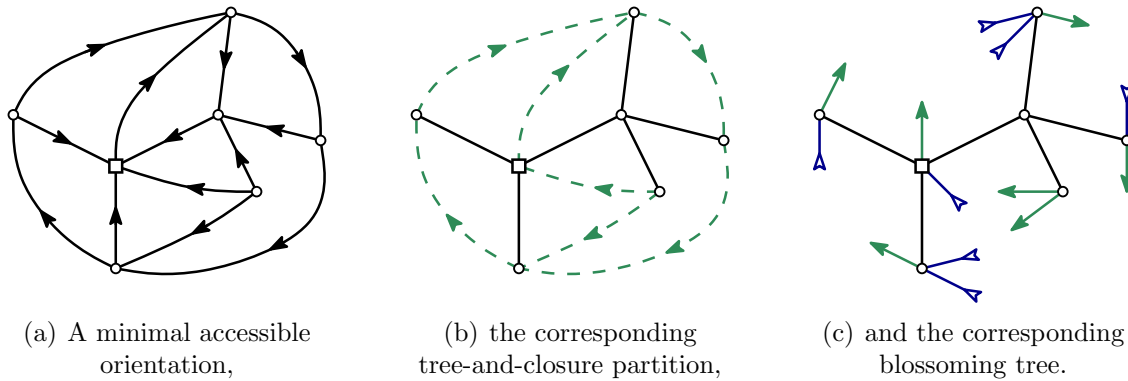


Figure 4: From a minimal accessible orientation to a blossoming tree.

2.2 Orientations and opening

The following theorem generalizes a result on tree orientations that can be explicitly found *e.g.* in [3], and which is at the heart of all bijections between map orientations and blossoming trees.

Theorem 2.3. *Let M be a plane map vertex-rooted at r , and suppose that M is endowed with a minimal accessible orientation O . Then M admits a unique edge-partition $(\mathcal{T}_M, \mathcal{C}_M)$ such that:*

- *edges in \mathcal{T}_M (called tree edges) form a spanning tree of M , rooted at r , on which the restriction of O is accessible;*
- *any edge in \mathcal{C}_M (called a closure edge) is a saturated clockwise edge in the unique cycle it forms with edges in \mathcal{T}_M .*

Let us call such a partition a tree-and-closure partition .

Before proving this theorem, we would like to emphasize in which way it differs from the result obtained in [3]. In the latter work, the outer face of a rooted map is required to be the root face. In this case and also in the particular case of triangulations treated in [33], a contour algorithm, starting at the root edge, enables to identify the edges of \mathcal{C}_M . The proof that this algorithm is correct relies deeply on the fact that *both* the accessibility and the minimality of the orientation are defined according to the root face. We show here that this hypothesis is unnecessary.

Proof. We prove this result by induction on the number of faces of M . If M has only one face, it is an accessible tree, hence the property is satisfied.

Let now $n \geq 2$, and suppose that the property is satisfied for any minimally oriented plane map with less than n faces. Let M be a vertex-rooted plane map with n faces endowed with a minimal accessible orientation.




Figure 5: Existence of a unique tree-and-closure partition.

We shall first prove that one of its outer edges e may be removed to obtain a vertex-rooted map $M_{\setminus e}$ endowed with a minimal accessible orientation. As outer edges do not form a counterclockwise cycle, at least one of them has the outer face strictly on its left, meaning that it is saturated and is not a bridge. Let (u, v) be such an edge and consider the map $M_{\setminus (u, v)}$. If it is accessible, then choose $e = (u, v)$. Otherwise, let C be the accessible component of the root vertex r in $M_{\setminus (u, v)}$. Then clearly v belongs to C while u does not. Moreover, u is accessible from all vertices not in C , see Fig. 5(a). Let D denote the complement of C . Then, the cut between C and D , made up of saturated edges oriented from C to D , is incident twice to the outer face of $M_{\setminus (u, v)}$. Let e be the edge of the cut with the outer face on its left. Since e is not a bridge in M , the map $M_{\setminus e}$ has $n - 1$ faces and the orientation induced by that of M is minimal and accessible.

Hence, by induction, $M_{\setminus e}$ admits a (unique) tree-and-closure partition $(\mathcal{T}_M, \mathcal{C}_M)$, and $(\mathcal{T}_M, \mathcal{C}_M \cup \{e\})$ is a tree-and-closure partition for M .

Let us now prove that M does not admit any other tree-and-closure partition, that is, does not admit any tree-and-closure partition with e in the tree. Suppose by contradiction that $(\mathcal{T}'_M, \mathcal{C}'_M)$ is a tree-and-closure partition for M with $e \in \mathcal{T}'_M$. Let us denote $e = (x, y)$, oriented from x to y , and consider the simple path γ from x to the root vertex in \mathcal{T}_M . At least one edge of γ does not belong to \mathcal{T}'_M (otherwise this would create a cycle in the tree), hence γ contains saturated edges. Let $e_1 = (x_1, y_1)$ be the first saturated edge in γ . Since x_1 lies in the subtree $\mathcal{T}'_M(x)$ of \mathcal{T}'_M rooted at x , e_1 belongs to \mathcal{C}'_M . Hence y_1 is explored after x_1 in the clockwise contour of \mathcal{T}'_M . Moreover, because (x, y) has the outer face on its left, the path γ_1 from y_1 to the root vertex of \mathcal{T}'_M cannot wrap (x, y) . Moreover, because (x, y) has the outer face on its left, it cannot be wrapped by a closure edge. Hence y_1 belongs to $\mathcal{T}'_M(x)$, see Fig. 5(b). In particular, there exists another (saturated) edge e_2 of γ that does not belong to γ_1 nor \mathcal{T}'_M for which the same reasoning applies. This implies the existence of an infinite sequence of edges of γ not belonging to \mathcal{T}'_M , a contradiction. \square

As an immediate corollary of Theorem 2.3, we obtain:

Corollary 2.4. *Let M be a vertex-rooted plane map endowed with a minimal accessible orientation O . Then there exists a unique vertex-rooted blossoming tree, endowed with an accessible orientation, the closure of which is M oriented with O .*

This blossoming tree is denoted \mathcal{B}_M .

Proof. Let M be a vertex-rooted plane map endowed with a minimal accessible orientation O . Consider the tree-and-closure partition $(\mathcal{T}_M, \mathcal{C}_M)$ of M . The edges of the blossoming tree \mathcal{B}_M are the edges of \mathcal{T}_M and each edge of \mathcal{C}_M is cut in two to produce a pair of opening and closing stems (see Fig.4). \square

Hence, in any particular case where a family of plane maps may be canonically endowed with a family of specific minimal and accessible orientations, Corollary 2.4 gives a bijection between that family of maps and a family of blossoming trees with the same distribution of in- and out- degrees. Whenever these trees are easily described and enumerated, this leads to a bijective proof of enumerative results.

2.3 Effective opening and closure

Let us point out some facts about the complexity of computing effectively the closure of a blossoming tree and the opening of a map.

To close a blossoming tree into a map, it is enough to perform a contour process and to match iteratively each opening stem with its corresponding closing stem along the way. Each time a new opening stem is explored, it can be stored in a stack structure (Last-In-First-Out), out of which one stem will be popped each time a closing stem is explored. This process goes around the outer face at most twice, hence the total time complexity is linear in the number of edges of the final map.

Unfortunately, things are not so smooth when it comes to opening an oriented plane map into its blossoming tree. Since the proof of Theorem 2.3 is essentially constructive, it yields an algorithm which identifies a closure edge at each step. Each of these steps consists in computing an accessible component, which can be done in linear time, resulting in a total quadratic complexity.

However, in the case where the map is corner-rooted in the outer face, the opening operation can be realized in linear time by an adapted depth-first search process. This construction has been introduced in a series of papers (see for example [35, 33]) in some particular cases and formally stated in [3] (where it appears in a slightly different form since the convention for tree edges orientation is opposite to ours).

Proposition 2.5 ([35, 33, 3]). *Let M be a corner-rooted plane map in which the outer and root faces coincide, and assume that M is endowed with a minimal orientation.*

Then, the tree-and-closure partition of M can be computed in linear time: initialize \mathcal{T}_M and \mathcal{C}_M as empty sets, and v and e to be respectively the root vertex and the root edge; then repeat the following steps until all edges belong either to \mathcal{C}_M or to \mathcal{T}_M :

- *if e does not belong to \mathcal{C}_M nor \mathcal{T}_M yet, add it to \mathcal{C}_M if it is oriented outwards v , and to \mathcal{T}_M otherwise;*

- if e belongs to \mathcal{T}_M , switch v to the other extremity of e ;
- update e to the next edge around v in clockwise order.

The proof of this proposition can be found in [3], in a different setting: it deals with rooted planar maps endowed with a distinguished spanning tree, for which an orientation is canonically defined by orienting tree edges towards the root, and any other clockwise in the unique cycle it forms with the tree. The set of all these *tree orientations* is hence equal to the set of minimal α -orientations for all accessibly feasible α .

Observe that as opposed to the purpose of [3], our work aims at defining one canonical orientation for each map. This requires to seek for an appropriate family of functions α for each family of maps we want to enumerate, as already mentioned at the end of Section 2.2. In the next section, we demonstrate that numerous previously known bijections can be easily retrieved as soon as we exhibit the adequate α .

3 Recovering previous bijections

Previous bijections between planar maps and blossoming trees obtained after Schaeffer [34] can all be seen as applications of Corollary 2.4, and more precisely of Proposition 2.5. In this section, we only consider corner-rooted planar maps, that is, corner-rooted plane maps with the root corner in the outer face. Hence the blossoming trees involved are *balanced*, meaning that no local closure may wrap the root corner. More formally, let us define the (non-cyclic) *contour word* of a rooted blossoming tree as the natural counterpart of the cyclic contour word starting at the root corner, see Definition 2.1. Then a rooted blossoming tree is said *balanced* if the restriction of its contour word on $\{b, \bar{b}\}$ is a Dyck word.

The intuition behind these bijections was originally relying on the interpretation of the enumerative formulas. We want to emphasize here that most of the time, a natural choice for a function α leads to the same construction. We only sketch how to retrieve the proofs that can be found in the original papers.

3.1 Maps with prescribed vertex degree sequence

Eulerian maps. The first bijection obtained by Schaeffer in [34] concerns planar *Eulerian maps* with prescribed vertex degrees, and in particular 4-regular planar maps with n vertices, that correspond bijectively to planar maps with n edges.

This bijection can be recovered in the following way. First recall that a map is said *Eulerian* if its vertices have even degrees. It is a classical result that Eulerian maps may be endowed with orientations with equal in- and out-degrees for each vertex. Besides these orientations are accessible. In particular, the minimal *Eulerian orientation* of a given plane Eulerian map can be obtained recursively by orienting clockwise the outer cycle and erasing it, see Fig. 6(a).

The generic opening of a planar rooted Eulerian map M with n_i vertices of degree $2i$ for any $i \in \llbracket 1, k \rrbracket$ (and hence $n = \sum_i in_i$ edges) endowed with its minimal Eulerian




Figure 6: Blossoming trees for an Eulerian map.

orientation leads to a balanced rooted blossoming tree with the same distribution of in- and out-vertex degrees, and both $\ell = 1 + \sum_i (i-1)n_i$ opening and closing stems. Observe that each non-root vertex has exactly one outgoing edge that belongs to the blossoming tree. Moreover, since the tree is balanced, the root corner is necessarily followed by an opening stem.

To enumerate balanced blossoming trees, a general strategy is to consider a larger family of planted blossoming trees that is stable by rerooting and in which the proportion of balanced ones can be evaluated. Let us sketch this strategy in the case of Eulerian maps, following [34]. Let us first consider planted trees with n_i nodes of degree $i+1$ (*i.e.* arity i) for any $i > 0$ and hence $\ell + 1 = 2 + \sum_i (i-1)n_i$ leaves *including the root one*; they are enumerated by ([24]):

$$T_{n_1, \dots, n_k} = \frac{1}{n} \binom{n}{\ell, n_1, \dots, n_k} = \frac{(n-1)!}{\ell!} \prod_{i=1}^k \frac{1}{n_i!}.$$

Start from one such tree and add $(i-1)$ opening stems on each node of arity i . The total number of trees that can be obtained in this way is then:

$$B_{n_1, \dots, n_k} = \prod_{i=1}^k \binom{2i-1}{i}^{n_i} T_{n_1, \dots, n_k} = \frac{(n-1)!}{\ell!} \prod_{i=1}^k \binom{2i-1}{i}^{n_i} \frac{1}{n_i!}.$$

Consider now each leaf (including the root one) as a closing stem, it yields a blossoming tree with $\ell - 1$ opening stems and $\ell + 1$ closing stems, whose closure gives a map with two unmatched closing stems.

Among all the planted blossoming trees that give the same map, a proportion $2/(\ell+1)$ of them are rooted on one of the unmatched stems. For those ones, changing their root (closing) stem into an opening one leads to a balanced blossoming tree. Hence:

Proposition 3.1. (Eulerian planar maps with prescribed vertex degrees) *The number of rooted planar Eulerian maps with n_i vertices of degree $2i$ for any $i \in \llbracket 1, k \rrbracket$ is*

given by:

$$\frac{2 \cdot (n-1)!}{(\ell+1)!} \prod_{i=1}^k \binom{2i-1}{i}^{n_i} \frac{1}{n_i!}.$$

An interesting particular case concerns rooted planar 4-regular maps with n vertices, that correspond bijectively to rooted planar maps with n edges. Indeed, let M be a planar map and consider its radial map \mathcal{R}_M (see Fig.1(c)). By convention, \mathcal{R}_M is rooted with the same root face as M , and its root vertex corresponds to the root edge of M .

It is clear from its definition that \mathcal{R}_M is a 4-regular map and that reciprocally every rooted 4-regular map with n vertices corresponds to a unique rooted planar map with n edges. According to our generic bijective scheme, rooted 4-regular maps are in bijection with balanced planted blossoming trees with n nodes, of in- and out-degrees 2, that is obtained from a planted binary tree by adding one opening stem to each node. Hence:

Corollary 3.2. (Planar maps with prescribed number of edges) *The number of rooted planar maps with n edges is:*

$$\frac{2 \cdot 3^n}{(n+2)(n+1)} \binom{2n}{n}.$$

General maps. In [9], the bijection for Eulerian maps is generalized into a bijection between planar maps with prescribed vertex degree sequence and some blossoming trees. We sketch in this paragraph how this construction can be derived from our generic scheme.

Any maps may be endowed with a *quasi-Eulerian* orientation, that is, a partial orientation (or a 2-fractional orientation) with equal in- and out-degrees for each vertex. In particular, as minimal Eulerian orientations, minimal quasi-Eulerian orientations can be obtained recursively; orient clockwise the outer cycle (with the convention that an edge adjacent twice to the outer face is partially oriented in both directions) and iterate after erasing outer edges (see Fig. 7(a) and 8(a)).

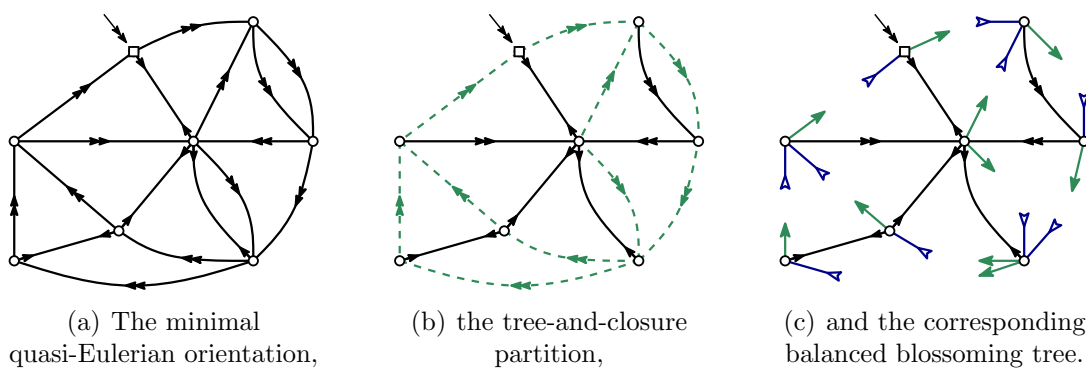


Figure 7: Blossoming trees for general maps: the generic R-case.

Let M be a rooted map endowed with its minimal quasi-Eulerian orientation. Opening M gives a rooted balanced blossoming tree endowed with an accessible 2-fractional orientation such that the in- and out-degrees of each vertex are equal. Now, to characterize blossoming trees that admit such a 2-fractional orientation, it is convenient to follow [9] and introduce the notion of *charge* of a subtree as the difference between the numbers of its closing stems and opening stems. It is then easily seen that a blossoming tree with total charge 0 admits an orientation with equal in- and out-degrees at each vertex if and only if its proper subtrees all have charge 0 or 1. More precisely, subtrees planted at a saturated edge have charge 1, while those planted at a bi-oriented edge have charge 0. This corresponds respectively to the R- and S-trees in [9].

In particular, in the generic case where the root edge e of M is not an isthmus, it is oriented clockwise and thus belongs to the closure, therefore \mathcal{B}_M consists of an R-tree and an opening stem situated right after the root corner (see Fig. 7). In the special case where e is an isthmus, it is partially oriented and belongs to the blossoming tree \mathcal{B}_M . Therefore the right subtree of \mathcal{B}_M is an S-tree carried by e and the other subtrees also form an S-tree (see Fig. 8).




Figure 8: Blossoming trees for general maps: the isthmus S-case.

Let us mention that the enumeration of these trees is not straightforward (and is carried out in Section 3 of [9]). The main difficulty comes from the fact that R-trees are not stable by rerooting.

3.2 m -Eulerian maps with prescribed degree sequence

In [8], the authors define m -Eulerian maps as bipartite maps such that black vertices all have degree m and each white vertex has degree multiple of m . We give in this section a much shorter proof that these maps are in bijection with the so-called m -Eulerian trees. Considering black vertices as hyper-edges, m -Eulerian maps form in fact a subclass of m -regular hypermaps, which boils down to Eulerian maps if $m = 2$. Let us root m -Eulerian maps in a white corner, and call the extremities of the root edge respectively the *white* and the *black root vertex*.

An important feature of these maps is that they can be canonically labelled on edges with integers in $\llbracket 1, m \rrbracket$ in such a way that:

- the root edge has label 1;
- around each black vertex, edges are labelled from 1 to m in clockwise order;
- around each white vertex, edges are cyclically labelled in counterclockwise order; in particular, if v is a white vertex with degree km , it is incident to exactly k edges of label i for any i in $\llbracket 1, m \rrbracket$.

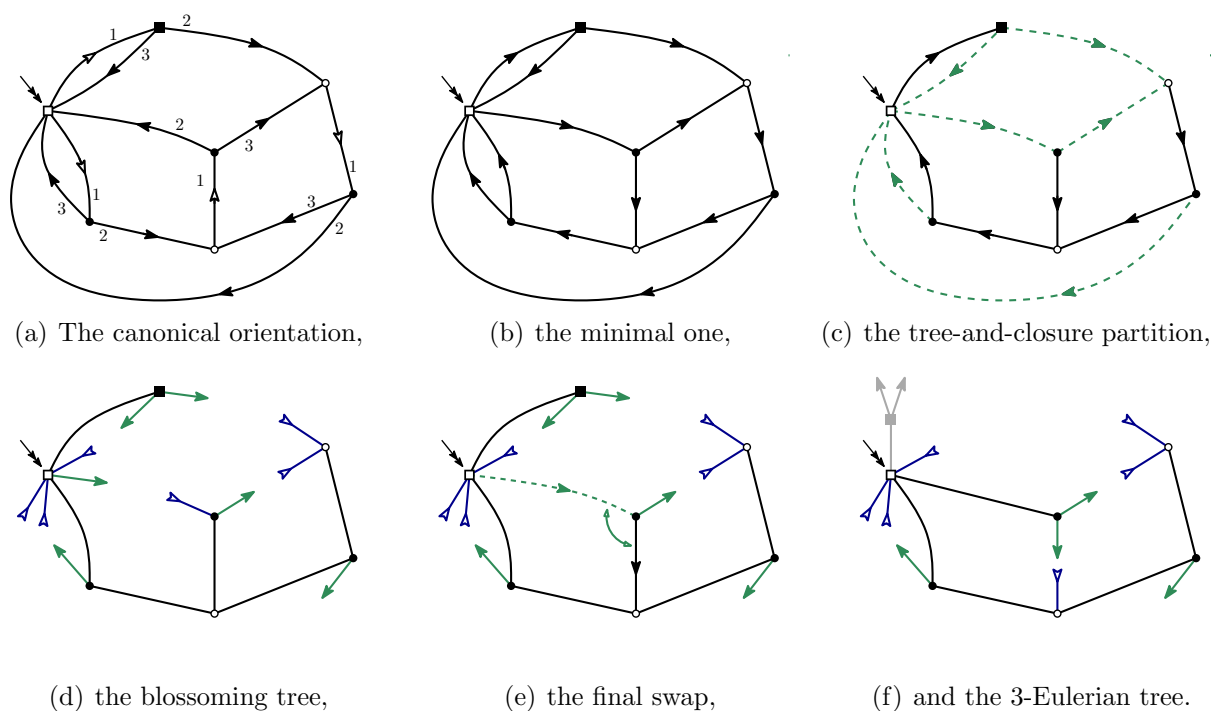


Figure 9: Blossoming trees for m -Eulerian maps: example of a 3-Eulerian map with two white vertices of degree 3 and one with degree 6, rooted on the edge with square ends.

This implies that any m -Eulerian map can be endowed with a *canonical* orientation: orient all edges but the 1-labelled ones from their black end to their white end, see Fig. 9(a). We will now prove that this canonical orientation is accessible. First observe that, for any face f , labels of edges incident to f are alternatively equal to i and $i + 1 \pmod m$ for a given i . Hence if $i = 1$ or m , f is an oriented cycle and all its vertices lie in the same strong connected component. For other values of i , its edges are all oriented from black to white, the i -labelled ones counterclockwise and the others clockwise. Hence, if its vertices were not in one and the same component, there would exist a cocycle which contains both an i -labelled edge e_0 and an $i + 1$ -labelled edge e_1 incident to f . Iterating this argument along the cocycle produces a cyclical sequence of edges (e_j) such that the

label of e_j is equal to $i + j$. Since no 1- or m -labelled edge can belong to a cocycle, we obtain a contradiction.

Observe that in the canonical orientation of any m -Eulerian map, each black vertex has outdegree equal to $m - 1$ (and indegree equal to 1), and each white vertex of degree km has outdegree equal to k . Consequently, the function α defined by $\alpha(v) = m - 1$ if v is black, and $\alpha(v) = k$ if v is white with degree km , is accessibly feasible.

Now any rooted m -Eulerian map M can be canonically embedded in the plane with the root face as outer face, and endowed with its minimal α -orientation. The generic opening according to the *black* root vertex leads to a rooted blossoming plane bipartite tree T such that (see Fig. 9(d)):

- the (black) root vertex has only one child – the white root vertex, and carries $m - 1$ opening stems;
- any white vertex with total degree km carries $k - 1$ opening stems, and $k(m - 1)$ black children or closing stems;
- black non-root vertices carry $m - 2$ opening stems and either a white child or a closing stem.

T is not exactly a *m-Eulerian tree* as defined in [8], but the transition between the two families is easy. It is enough to modify the tree locally in a way such that black and white vertices respectively carry only opening and closing stems. To do so, observe that opening stems carried by white vertices are matched with closing stems carried by black ones (since the underlying map is bipartite). Now suppose that such a couple is carried by a white vertex u and a black vertex v . It can be replaced by its closure edge (u, v) , creating a cycle that is broken by opening the edge connecting v to its father in T so as to create a new couple of opening and closing stems (see Fig. 9(e)-(f)). This *swap* leads to a blossoming tree with only opening stems on black vertices and closing stems on white vertices; removing the black root vertex gives precisely an m -Eulerian tree. As shown in [8]:

Proposition 3.3. *Let $m \geq 2$, the number of edge-rooted m -Eulerian maps with d_i white vertices of degree mi for each $i \geq 1$ is:*

$$m(m - 1)^{v_\circ - 1} \frac{[(m - 1)v_\bullet]!}{[(m - 1)v_\bullet - v_\circ + 2]!} \prod_{i \geq 1} \frac{1}{d_i!} \binom{mi - 1}{i - 1}^{d_i},$$

where $v_\bullet = \sum id_i$ and $v_\circ = \sum d_i$ denote respectively the number of black and white vertices.

3.3 Non-separable maps with prescribed number of edges

A *cut vertex* of a map is a vertex that is incident twice to the same face; a map is said to be *non-separable* if it has no cut vertex. Observe that this definition is stable by duality.

It includes the two maps with only one edge, but in the following, we will consider only maps with at least two edges. In [35], bijections are described for general and cubic non-separable planar maps, and we show here that they are indeed special cases of our generic bijective scheme.

A *bipolar orientation* of a map (or more generally a graph) is an acyclic orientation of its edges with a single *source* (vertex without any incoming edge) and a single *sink* (vertex without any outgoing edge), which are called the *poles* of the orientation. Non-separable maps (or graphs) are characterized by the following property (see *e.g.* [16]):

Proposition 3.4. *A rooted map is non-separable if and only if it can be endowed with a bipolar orientation with the two ends of the root edge as poles.*




Figure 10: Properties of bipolar orientations; the non-separable map is drawn in plain lines, its quadrangulation in dashed green ones.

In the planar case, rooted non-separable maps endowed with a bipolar orientation (with, say, the root vertex as the sink) have some interesting properties, illustrated in Fig. 10:

1. each face is itself bipolar, hence its corners may be classified into *lateral* ones (left or right) and two *polar* ones (source and sink);
2. each vertex but the two poles has exactly one bundle of incoming edges and one bundle of outgoing ones, hence its corners may be classified into some polar ones (source or sink) and two lateral ones (left and right);
3. the quadrangulation Q_M of the map can be endowed with an α -orientation so that each extremity of the root edge has indegree 0 and every other vertex has 2 incoming edges (one for each special corner of the corresponding face or vertex of M , see Fig.10(a) and 10(b));
4. reciprocally to each such orientation of Q_M corresponds a bipolar orientation of M , hence the set of bipolar orientations of a plane non-separable map M is endowed with a lattice structure inherited from the lattice structure of the α -orientations of its quadrangulation;

5. the minimal element of this lattice is the unique bipolar orientation of M such that, for any distinct vertices u and v both incident to distinct faces f and g , the following configuration is forbidden: u sink of f and on the left of g , v on the right of f and source of g , see Fig. 10(c). This pattern corresponds to a counterclockwise 4-cycle in the quadrangulation.

The constructions for general and cubic non-separable planar maps are based on this minimal bipolar orientation. In the following, we denote by s and t respectively the source and the sink of the considered bipolar orientations, and say that a face is *generic* if it is *not* incident to the root edge. We use the usual convention that bipolar oriented maps are drawn in the plane with oriented paths going upwards. Moreover s and t are outer vertices, s is at the bottom and t at the top of the figure, and with this convention we choose the embedding in which the root edge is the rightmost one.

3.3.1 Non-separable cubic maps

The case of cubic maps (treated in [35], extended in [32]) is very constrained: since each vertex has degree 3, in any bipolar orientation, each non-polar vertex has either one incoming edge and two outgoing ones, or the opposite. Hence each vertex (including the two poles) is either the source or the sink of exactly one generic face. Let M be such a rooted non-separable cubic map endowed with its minimal bipolar orientation, and let us add an extra *bipolar edge* in each generic face of M (that is, an edge between its two poles). These extra edges realize a perfect matching of the vertices, hence the resulting planar map \overline{M} is 4-regular. As such, \overline{M} can be endowed with its minimal Eulerian orientation, and opened accordingly into a blossoming tree rooted at the sink t . We say that \overline{M} is the *completion* of M .

Lemma 3.5. *All bipolar edges belong to the resulting blossoming tree \mathcal{B}_M , and each vertex carries exactly one opening stem, immediately before the bipolar edge in clockwise order.*

Proof. This is proved inductively on the number of non-polar vertices of M . The lemma is true in the smallest case (three parallel edges between s and t). Let now M have at least two inner vertices. Let f be the (only) generic face incident to s , g the bounded non-generic face, and u the sink of f . The vertex u may be incident to the outer face (and even be equal to t); this specific case is illustrated in Fig. 12, and the generic case is illustrated in Fig. 11. We show that the bipolar edge (s, u) satisfies the lemma, and in each possible case we build smaller non-separable cubic maps such that \mathcal{B}_M is obtained from their respective blossoming trees by grafting them together with the small subtree made of s , u and their incident edges and stems.

Let us first prove that the bipolar edge (s, u) satisfies the lemma.

Properties of the minimal Eulerian orientation of \overline{M} imply that the border of the outer face is a clockwise cycle, and that the edges incident to g are oriented clockwise or counterclockwise around g depending on whether they are incident to the outer face or not. Observe also that the two outer edges incident to s are oriented clockwise and since

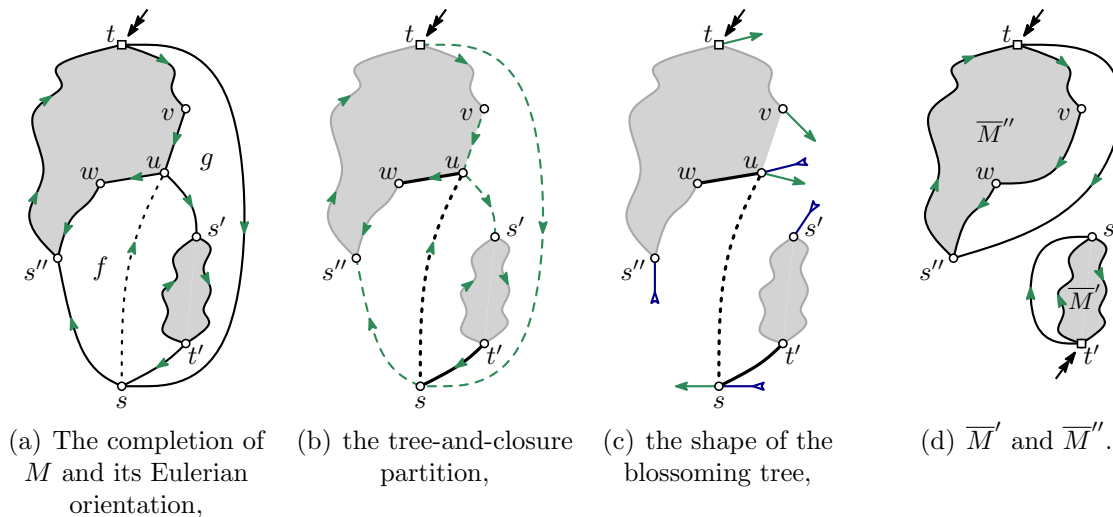


Figure 11: Generic case of non-separable cubic maps.

they are unnecessary to the accessibility of t , they belong to the closure. Therefore the bipolar edge (s, u) , oriented from s to u , belongs to the tree.

Let us now prove by contradiction that the face on the right of u is the non-generic bounded face g , that is, its source and sink are s and t . Otherwise, let us define sequences of faces (f_i) and (g_i) and of vertices (s_i) , (u_i) and (v_i) as follows. Let f_0 be equal to f , and for any $i \geq 0$, let s_i and u_i be respectively the source and the sink of f_i , g_i the face on the right of u_i , v_i its source, and finally f_{i+1} the face on the left of v_i . By assumption, g_0 is different from g (and clearly also from f), hence v_0 is not equal to s , and f_1 is well defined. More generally, it can be checked that v_i cannot be equal to s for any $i \geq 1$, hence the whole sequence (f_i) is well defined.

Let λ_i be the leftmost path from s to s_i , and ρ_i be the rightmost path from s to v_i . Denote now by R_i the region delimited by λ_i and the left border of f_i on the left, and ρ_i and the left border of g_i on the right. Then u_{i+1} belongs to R_i , hence (u_i) is decreasing for the partial order induced by the acyclic orientation. It is therefore stationary after some i_0 . By definition, u_{i_0} , v_{i_0} , f_{i_0} and g_{i_0} define exactly the forbidden pattern 10(c), hence the underlying bipolar orientation is not minimal, a contradiction.

Since the face on the right of u is g , the edge following the bipolar edge (s, u) in counterclockwise order around u is incident to both f and g . Hence it is necessarily outgoing, and belongs to the closure, which proves the lemma locally for s and u .

Let us decompose M into smaller non-separable cubic maps.

Now let s' and t' be the respective neighbours of u and s that are both incident to f and g – possibly equal to s and u , see Fig. 11-12(a). If $s' \neq s$ and $t' \neq u$, deleting (u, s') and (t', s) disconnects M into two submaps; as these edges are both oriented clockwise around f , (u, s') belongs to the closure and (t', s) belongs to the tree, see Fig. 11-12(b). Adding a root edge between s' and t' with root vertex t' in the corresponding submap leads to a smaller non-separable cubic map M' , see Fig. 11-12(d). Observe that the restriction




Figure 12: Specific *series parallel* case for non-separable cubic maps.

to M' of the minimal bipolar orientation of M is precisely the opposite of the minimal bipolar orientation of M' . Hence these two orientations define the same matching between the sources and the sinks of generic faces. Thus the completion of M' is a submap of \overline{M} (except for its root edge), and its minimal Eulerian orientation is the restriction of the one of \overline{M} .

We resume to s and u being possibly equal to s' and t' . To end the construction in the generic case, let s'' be the last neighbour of s , and let v and w be the two last neighbours of u , with v incident to g and w incident to f , see Fig. 11. Then (u, w) is a tree edge, while (s, s'') and (v, u) are closure edges. Let M'' be obtained from M by removing s and u and their incident edges, and adding to the component of t and s'' successively in the outer face a (closure) edge between v and w , and a (closure) root edge between t and s'' . Then M'' is a smaller non-separable cubic map, whose minimal bipolar orientation is the restriction of the minimal bipolar orientation of M . Hence \mathcal{B}_M is obtained from $\mathcal{B}_{M''}$ by grafting the subtree made from u , s and $\mathcal{B}_{M'}$ instead of the suitable stem of w .

Now, in the case where u is incident to the outer face, the situation between the face f and the outer face is similar to the one between the faces f and g , as illustrated in Fig. 12. Let s'' and t'' be the neighbours of s and u between these two faces, then (s, s'') belongs to the closure and, as soon as the two edges are distinct, (t'', t) belongs to the tree. Adding a root edge between s'' and t'' in the corresponding submap leads to a smaller non-separable cubic map M'' .

If moreover $u \neq t$, let s''' be its fourth neighbour; the edge (u, s''') (oriented towards s''') is incident to the two non-generic faces, hence deleting the root edge and (u, s''') disconnects M , which implies in particular that (u, s''') belongs to the tree. In this case, let M''' denote the submap containing s''' and t , with an additional root edge between s''' and t : M is somehow a *series parallel* compound of three submaps M' , M'' and M''' , each possibly empty in degenerate cases, see Fig. 12. \square




Figure 13: Complete example of a non-separable cubic map M .

Following [35, 32], let us define a blossoming *twin ternary tree* as a tree obtained by fairly splitting each node of a ternary tree into two *twin* nodes linked by a special *twinning edge*, with an additional opening stem on each node right before its twinning edge in clockwise order (see Fig.13(d)).

Observe that the opening of the completion of a non-separable cubic map leads to a (balanced) blossoming twin ternary tree, where twinning edges are bipolar edges. Indeed, non-root opening stems always follow bipolar edges in counterclockwise order around vertices and the contraction of bipolar edges yields a planted ternary tree (the leaves of which are the closing stems).

Conversely, we prove now that the closure of any balanced blossoming twin ternary tree is the completion of a non-separable cubic map. For any subtree (a node, its stems and its descendants), let us call *free* the stems that are matched with stems not belonging to the subtree. An immediate counting shows that each subtree has one more closing stem than opening ones, hence one more free closing stem than free opening ones. Let us denote by s the node carrying the closing stem c corresponding to the (opening) root stem, and let u be its twin node; let us prove by contradiction that u is its parent. If u were the right child of s , its opening stem would be matched with the (only) closing stem of s . If u were its left child, the subtree of s would have at least two free closing stems since the opening stem of u would indeed be free, which prevents c from being matched with the root stem. Hence u is necessarily the parent of s , and as the opening stem of u is matched to a closing stem in the right subtree of s , c is necessarily just before the opening stem of s in clockwise order around s . Hence s and u are exactly in the configuration of Fig.11(a) or Fig.12(a). We can conclude by a similar induction argument as in the proof of Lemma 3.5.

Hence rooted non-separable cubic planar maps with $2n$ vertices are in one-to-one correspondence with balanced blossoming twin ternary trees with $2n$ nodes. The number of planted ternary trees with n nodes is $\frac{1}{2n+1} \binom{3n}{n}$, each one leading to 2^n distinct planted

blossoming twin ternary trees. A fraction $\frac{2}{2n+2}$ of them is balanced, leading to:

Corollary 3.6. *The number of rooted planar non-separable cubic maps with $2n$ vertices and $3n$ edges is equal to:*

$$\frac{2^n}{(n+1)(2n+1)} \binom{3n}{n}.$$

3.3.2 General non-separable maps

In the case of general planar non-separable maps, the generic scheme is applied on (an extension of) their radial maps, as in the proof of Corollary 3.2. Let M be a rooted non-separable map, and \mathcal{R}_M its radial map, rooted with the same root face as M , and root vertex r corresponding to the root edge of M . The orientation of the quadrangulation of M defined below Proposition 3.4 yields naturally an orientation of \mathcal{R}_M such that each generic face of \mathcal{R}_M has exactly two clockwise edges (corresponding to the two special corners of the corresponding face or vertex of M). For edges incident to r , we adopt the convention that the two outer ones are oriented clockwise, and the last two ones are outgoing for the root vertex. This orientation is said *minimal* if the corresponding bipolar orientation of M is itself minimal, see Fig. 14(a).

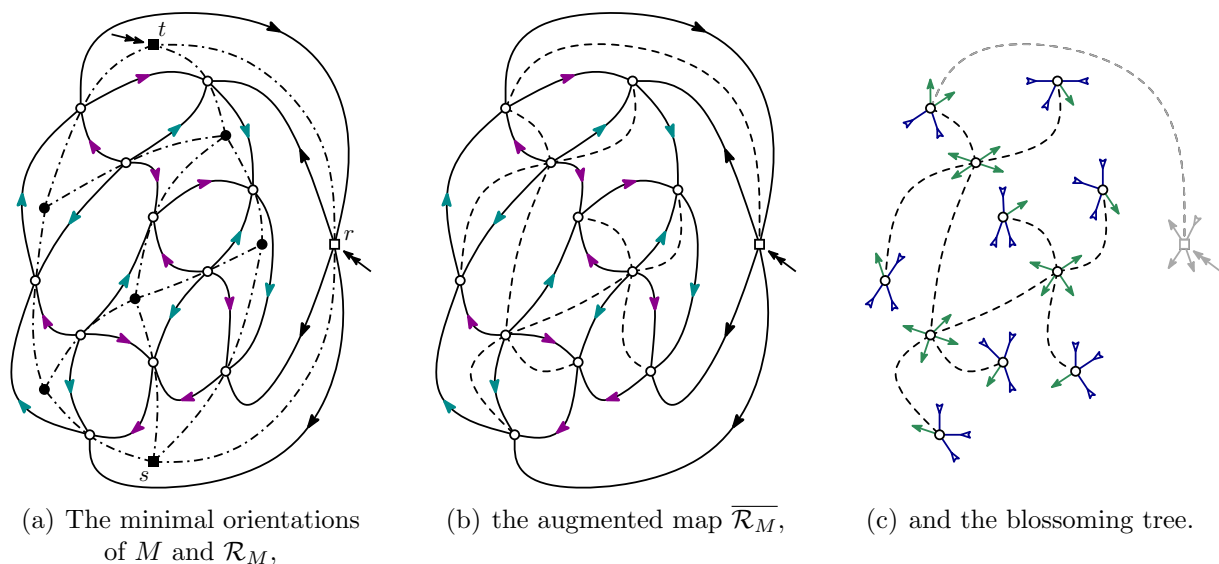


Figure 14: Complete example of a non-separable planar map M .

Given such an orientation of \mathcal{R}_M , each generic face has two special corners (the origins of the two clockwise edges), and so does the face that corresponds to t . Let us add an extra edge in each such face between these two corners, and denote respectively $\overline{\mathcal{R}_M}$ the resulting map and \mathcal{T}_M the map made exactly of these extra edges and their incident vertices, see Fig. 14(b). Then:

Lemma 3.7 ([35]). *\mathcal{T}_M is a spanning tree of $\overline{\mathcal{R}_M}$ if and only if the underlying orientation of \mathcal{R}_M is minimal. Moreover in this case, the edges of \mathcal{R}_M go clockwise around \mathcal{T}_M .*

Hence in the minimal case, \mathcal{R}_M is actually a valid set of closure edges around \mathcal{T}_M , whatever accessible orientation is chosen for edges in \mathcal{T}_M . For instance, we may orient all edges towards r , or simply leave them unoriented (that is, oriented both ways in a 2-fractional orientation). Then $(\mathcal{T}_M, \mathcal{R}_M)$ is the tree-and-closure partition of $\overline{\mathcal{R}_M}$.

The resulting (balanced) blossoming tree is such that each non-root vertex is incident to 4 stems, and to as many edges as opening stems – hence it has in-degree equal to 4. Considering closing stems as leaves, and after some surgery to remove the root vertex, we get a planted ternary tree with one extra opening stem at each corner before an inner edge (clockwise around each vertex), see Fig. 14(c). Reciprocally, as shown in [35], the closure edges of such a balanced blossoming tree T form a 4-regular map R endowed with an orientation with 2 clockwise edges per face, which by Lemma 3.7 ensures that R is the radial map of a non-separable map, endowed with its minimal orientation. Since any planted blossoming ternary tree with n nodes has $2n + 2$ leaves (including the root stem) and $2n - 2$ opening stems, a fraction $\frac{4}{2n+2}$ among them is balanced, hence:

Corollary 3.8. *For $n \geq 1$, the number of rooted planar non-separable maps with $n + 1$ edges is equal to:*

$$\frac{2}{(n + 1)(2n + 1)} \binom{3n}{n}.$$

3.4 Simple triangulations and quadrangulations

Bijections between simple triangulations or quadrangulations and blossoming trees, as described in [33, 20], are special cases of the general bijection for d -angulations of girth d as explained in Sections 5 and 6, with a special emphasis in Subsection 6.3. In particular, the uniqueness part in Theorem 2.3 gives a more direct proof that the closure construction of [33, 20] for simple triangulations and quadrangulations of a p -gon is injective, while the existence part proves surjectivity without requiring a cardinality argument.

Besides, it is noteworthy that in such cases where the degree of faces are prescribed, closing stems are redundant; since the underlying orientation is regular, blossoming trees are trees with a fixed number of opening stems per vertex.

4 Plane bipolar orientations

Recall that a bipolar orientation of a map is an acyclic orientation of its edges with a single source (vertex without incoming edge) and a single sink (vertex without outgoing edge). A *plane bipolar orientation* is a (non-separable) corner-rooted map (with at least two edges) endowed with a bipolar orientation such that the root vertex of the map is the sink of the orientation and its source is the other extremity of the root edge, see Fig. 15(a). We emphasize that this section is devoted to the study of *all* plane bipolar orientations, as opposed to Section 3.3 which focuses only on maps endowed with their minimal bipolar orientation as a tool to enumerate non-separable maps. Plane bipolar orientations have a nice enumerative formula:


Theorem 4.1 (Baxter [2]). *For all non-negative integers i and j , the number Θ_{ij} of plane bipolar orientations with i non-pole vertices and j generic faces is equal to:*

$$\Theta_{ij} = \frac{2(i+j)!(i+j+1)!(i+j+2)!}{i!(i+1)!(i+2)!j!(j+1)!(j+2)!}. \quad (1)$$

The first proof of this formula was given by Baxter [2]. His proof involves quite technical computation and relies on a “guess and check” approach. Since then, some bijective proofs of this result have been obtained in [21], [6] and [19]. Our generic scheme provides a new bijective proof.

4.1 Bijection with triple of paths

Let us consider paths on \mathbb{Z}^2 made of right-steps $(0, 1)$ and up-steps $(1, 0)$. A configuration of such paths is called *non-intersecting* if each vertex of \mathbb{Z}^2 belongs to at most one path. For $i, j \in \mathbb{N}$, define the set $\mathcal{P}_{i,j}$ of non-intersecting triple of paths $(p^{(1)}, p^{(2)}, p^{(3)})$, each made of exactly i right- and j up-steps and starting respectively at $(-1, 1)$, $(0, 0)$ and $(1, -1)$ (and hence ending at $(i-1, j+1)$, (i, j) and $(i+1, j-1)$), see Fig. 15(c).



(a) A plane bipolar orientation, (b) its blossoming tree, (c) the resulting triple of paths.

Figure 15: Example of the bijection for a bipolar orientation M .

The contour word w of its blossoming tree is $ebbb\bar{e}e\bar{e}bb\bar{e}\bar{e}b\bar{e}bb\bar{e}\bar{e}b\bar{e}bb\bar{e}\bar{e}b\bar{e}bb\bar{e}\bar{e}$.
 Then $w^{(1)} = eee\bar{e}b\bar{e}b\bar{e}b\bar{e}bb\bar{e}bb$, $w^{(2)} = \bar{e}b\bar{e}bb\bar{e}\bar{e}\bar{e}b\bar{e}bb\bar{e}$ and $w^{(3)} = bbb\bar{e}b\bar{e}b\bar{e}b\bar{e}b\bar{e}$.

The rest of this section is devoted to the proof of the following theorem, from which a direct application of Lindström-Gessel-Viennot Lemma [22] yields the enumerative result of Baxter:

Theorem 4.2. *For all positive integers i and j , there exists a one-to-one constructive correspondence between the set of plane bipolar orientations with i generic faces and j non-pole vertices and the set $\mathcal{P}_{i,j}$.*

Other bijections between plane bipolar orientations and the set \mathcal{P} already appear in [21, 19]. It must nevertheless be emphasized that the bijection we obtained is different, thus providing a first example where our general scheme yields a new bijective construction. To prove the theorem, we start by applying the generic scheme to open a bipolar orientation into a blossoming tree, which is then encoded by a non-intersecting triple of paths.

4.2 From bipolar orientations to configurations of paths

Since any bipolar orientation is acyclic and its sink vertex t is accessible from any vertex, Corollary 2.4 can be applied to open it into a blossoming tree, rooted at the former outer corner of t . Let \mathcal{T} (resp. $\mathcal{T}_{i,j}$) be the set of balanced blossoming trees obtained when opening a plane bipolar orientation (resp. with i generic faces and j non-pole vertices).

Let T be a rooted blossoming tree, we consider its contour word w on the alphabet $\{e, \bar{e}, b, \bar{b}\}$, where e and \bar{e} encode respectively the first and second exploration of an edge and b and \bar{b} encode respectively the exploration of an opening or a closing stem. For instance, the contour word of the tree of Fig.15(b) is $ebbb\bar{e}ee\bar{b}\bar{b}\bar{e}\bar{e}\bar{b}\bar{b}\bar{e}\bar{b}\bar{e}\bar{b}\bar{e}\bar{b}\bar{e}\bar{b}\bar{e}\bar{b}\bar{e}\bar{b}\bar{e}$. Let furthermore define $w^{(1)}$, $w^{(2)}$ and $w^{(3)}$ as the subwords of w obtained respectively by keeping only the letters e and \bar{b} , \bar{e} and \bar{b} , \bar{e} and b :

$$w^{(1)} = w_{|e,\bar{b}}, \quad w^{(2)} = w_{|\bar{e},\bar{b}} \quad \text{and} \quad w^{(3)} = w_{|\bar{e},b}.$$

Claim 4.3. *For T a blossoming tree in $\mathcal{T}_{i,j}$, the words $w^{(1)}$, $w^{(2)}$ and $w^{(3)}$ have length $\ell = i + j + 2$ and furthermore:*

$$w_1^{(1)} = e \text{ and } w_\ell^{(1)} = \bar{b}, \quad w_1^{(2)} = \bar{e} \text{ and } w_2^{(2)} = \bar{b}, \quad w_1^{(3)} = b \text{ and } w_\ell^{(3)} = \bar{e},$$

where w_k denotes the k -th letter of a word w .

This claim is a consequence of Lemma 4.5 below, we omit its proof. For $T \in \mathcal{T}$, Claim 4.3 enables to define $(\tilde{w}^{(1)}, \tilde{w}^{(2)}, \tilde{w}^{(3)})$ as:

$$w^{(1)} = e\tilde{w}^{(1)}\bar{b}, \quad w^{(2)} = \bar{e}\bar{b}\tilde{w}^{(2)} \quad \text{and} \quad w^{(3)} = b\tilde{w}^{(3)}\bar{e}. \quad (2)$$

Triple of words $(\tilde{w}^{(1)}, \tilde{w}^{(2)}, \tilde{w}^{(3)})$ can be naturally represented by triple of up-right paths $(\tilde{p}^{(1)}, \tilde{p}^{(2)}, \tilde{p}^{(3)})$, with initial points $(-1, 1)$, $(0, 0)$ and $(1, -1)$, by replacing letters e or \bar{e} by up-steps and letters b and \bar{b} by right-steps. Let Φ be the application that associates to each tree of \mathcal{T} the corresponding triple of paths, see Fig. 15. Observe that if M has i generic faces and j non-pole vertices, then the corresponding blossoming tree T has $i + 1$ pairs of opening-closing stems and $j + 1$ edges, therefore each of the paths $\tilde{p}^{(1)}$, $\tilde{p}^{(2)}$ and $\tilde{p}^{(3)}$ have exactly i right-steps and j up-steps.

Proposition 4.4. *Let T be an element of $\mathcal{T}_{i,j}$. Then its image by Φ is non-intersecting, in other words it belongs to $\mathcal{P}_{i,j}$.*

This follows from:

Lemma 4.5. *A word w on the alphabet $\{e, \bar{e}, b, \bar{b}\}$ is the contour word of an element of \mathcal{T} if and only if the five following conditions hold:*

- (1) $w_1 = e$ and $w_2 = b$; (2) $w_{|e, \bar{e}}$ is a Dyck word; (3) $w_{|b, \bar{b}}$ is a Dyck word;
- (4) for all $i < j$, if $w_i = b$ and $w_j = \bar{b}$, there exists $k \in [i, j]$ such that $w_k = \bar{e}$;
- (5) for all $1 < i < j$, if $w_i = e$ and $w_j = \bar{e}$, there exists $k \in [i, j]$ such that $w_k = \bar{b}$.

Proof. Let T be an element of \mathcal{T} and let w be its contour word. For the first condition, observe that the root edge, oriented from s to t , belongs to the blossoming tree and is the first edge encountered in the contour of T , thus giving a first letter e . Moreover, s is a leaf of T that carries only opening stems and at least one, so $w_2 = b$. Condition (2) reflects the fact that T is a tree and Condition (3) that it is balanced.

Conditions (4) and (5) are proved by contradiction. If Condition (4) does not hold, there exists a factor of w of the form $be \cdots e\bar{b}$ or $b\bar{b}$. It implies that a vertex is matched with one of its descendants (possibly itself) in the closure, producing an oriented cycle, a contradiction. Similarly, if Condition (5) does not hold, there exists a leaf of T (different from s) which does not carry a closing stem. It contradicts the uniqueness of the source.

Reciprocally, Conditions (2) and (3) imply that w is the contour word of a balanced blossoming tree T . Moreover, Conditions (1), (4) and (5) imply that the first subtree of the root is reduced to one edge (t, s) , where s carries at least one opening stem. To ensure that T belongs to \mathcal{T} , it is enough to prove (a) that each vertex of T different from s or t has at least one ingoing and one outgoing edges and (b) that the orientation of the closure is acyclic.

Each node of T has at least one ingoing edge per child and Condition (4) ensures that each leaf of T carries at least one closing stem. Moreover each vertex but t has one outgoing edge (oriented towards its parent in T), hence (a) is satisfied.

If the opening stem b is matched with the closing stem \bar{b} , then the corner incident to b is explored before the one incident to \bar{b} in the contour process of T . It implies that if there exists an oriented cycle in the closure of T , at least one closure edge links a node to one of its descendants in T . Let us prove that this cannot happen, since in the contour word of T , each occurrence of b is followed by $b^*\bar{e}$. Consider an occurrence of b in w and the first occurrence of \bar{b} after it. Consider the first occurrence of \bar{e} after b (which precedes \bar{b} by Condition (4)). The corresponding occurrence of e necessarily precedes b by Condition (5), hence there is no occurrences of e between b and \bar{e} . Hence (b) is satisfied. \square

Proof of Proposition 4.4. For any $\ell \geq 0$, denote $w_{\leq \ell}$ the prefix of length ℓ of a word w , and $|w|_x$ the number of occurrences of x in w . Let w be the contour word of T , Condition (2) implies that:

$$\forall \ell \geq 0, \quad |w_{\leq \ell}^{(1)}|_{\bar{b}} \leq |w_{\leq \ell}^{(2)}|_{\bar{b}} \quad \text{and} \quad |w_{\leq \ell}^{(1)}|_e \geq |w_{\leq \ell}^{(2)}|_{\bar{e}}.$$

Consequently, the corresponding paths $(p^{(1)}, p^{(2)})$ starting at $(0, 0)$ are such that $p^{(1)}$ lies above and on the left of $p^{(2)}$ with possible common vertices or edges. However the two

paths share no vertical edge but the first one by Condition (5). Hence, after translating $p^{(1)}$ by an up-step, $p^{(1)}$ and $p^{(2)}$ are non-intersecting and so are $\tilde{p}^{(1)}$ and $\tilde{p}^{(2)}$.

We can prove similarly that $(\tilde{p}^{(2)}, \tilde{p}^{(3)})$ is non-intersecting by observing first that $p^{(2)}$ lies above and on the left of $p^{(3)}$ thanks to Condition (3) of Lemma 4.5 and that $p^{(2)}$ and $p^{(3)}$ cannot share horizontal edge by Condition (4). Translating $p^{(3)}$ by a right-step and deleting the appropriate steps of $p^{(2)}$ and $p^{(3)}$ yield the desired result. \square

4.3 From configurations of paths to blossoming trees

Let $p = (\tilde{p}^{(1)}, \tilde{p}^{(2)}, \tilde{p}^{(3)})$ be a configuration of paths in $\mathcal{P}_{i,j}$ and $(w^{(1)}, w^{(2)}, w^{(3)})$ the corresponding triple of words. Let us decompose $w^{(1)}$ and $w^{(3)}$ as a sequence of factors according respectively to the occurrences of \bar{b} and \bar{e} :

$$w^{(1)} = e \cdot w_{[1]}^{(1)} \cdot w_{[2]}^{(1)} \cdot \dots \cdot w_{[j+1]}^{(1)} \quad \text{and} \quad w^{(3)} = b \cdot w_{[1]}^{(3)} \cdot w_{[2]}^{(3)} \cdot \dots \cdot w_{[i+1]}^{(3)},$$

where each factor $w_{[k]}^{(1)}$ (resp. $w_{[k]}^{(3)}$) is of the form $e^* \bar{b}$ (resp. $b^* \bar{e}$). These factors are “bricks” used to reconstruct a *compatible* word w , that is a word such that:

$$w^{(1)} = w_{|e, \bar{b}}, \quad w^{(2)} = w_{|\bar{e}, \bar{b}} \quad \text{and} \quad w^{(3)} = w_{|\bar{e}, b}.$$

The order in which those bricks are added is driven by $w^{(2)}$: let \bar{w} be obtained from $w^{(2)}$ by replacing its k -th occurrence of \bar{b} by $w_{[k]}^{(1)}$ and its k -th occurrence of \bar{e} by $w_{[k]}^{(3)}$, and define finally $w = e\bar{w}$.

Proposition 4.6. *Let p be an element of $\mathcal{P}_{i,j}$ and let w be the corresponding word as defined above. The word w is the unique word compatible with p that satisfies the five conditions of Lemma 4.5. In other words it is the contour word of the unique blossoming tree of \mathcal{T} compatible with p .*

Proof. First observe that w is compatible with p and that no other such word may satisfy the conditions of Lemma 4.5: Conditions (4) and (5) imply that the factors $w_{[k]}^{(1)}$ and $w_{[k]}^{(3)}$ have to be factors of w , and the order in which they appear is completely determined by $w^{(2)}$.

Let us now prove that w encodes indeed an element of \mathcal{T} , by applying Lemma 4.5. Condition (1) is clearly satisfied. Conditions (4) and (5) follow also easily from the definition of the decomposition in factors of $w^{(1)}$ and $w^{(3)}$: observe for instance, for Condition (4), that any occurrence of b in w (but the first one) comes from a factor $w_{[k]}^{(3)}$ of the form $b^* \bar{e}$. The first occurrence of b does not either raise a problem since $w_1^{(2)} = \bar{e}$ and is hence replaced by $w_{[1]}^{(3)}$.

It remains to prove Conditions (2) and (3), namely that $w_{|e, \bar{e}}$ and $w_{|b, \bar{b}}$ are Dyck words. We only give the proof for $w_{|e, \bar{e}}$, since both proofs work along the same lines. From the construction of w , the number of occurrences of e and of \bar{e} in w are both equal to $i + 1$, hence it is enough to prove that $|w_{\leq k}|_e \geq |w_{\leq k}|_{\bar{e}}$, for all k . We consider the following decomposition of w into product of factors:

$$w = eb \cdot w_{[1]} \cdot w_{[2]} \cdot \dots \cdot w_{[i+j+2]},$$

where each of the $w_{[k]}$ is equal to the corresponding factor of $w^{(1)}$ or of $w^{(3)}$. It is then enough to check that for each $1 \leq k \leq i + j + 2$:

$$\left| eb \prod_{i=1}^k w_{[i]} \right|_e \geq \left| eb \prod_{i=1}^k w_{[i]} \right|_{\bar{e}},$$

which can be rewritten as:

$$1 + \sum_{i=1}^{k_1} |w_{[i]}^{(1)}|_e \geq \sum_{i=1}^k |w_i^{(2)}|_{\bar{e}}, \quad \text{where } k_1 = |w_{\leq k}^{(2)}|_{\bar{e}}.$$

Let (x_2, y_2) be the point of $p^{(2)}$ reached after k steps and let (x_2, y_1) be the point of $p^{(1)}$ of abscissa x_2 with minimal ordinate. By construction of w , the value of y_1 is equal to the left-hand side of the above inequality, while y_2 is equal to the right hand-side. Since $p^{(1)}$ lies above $p^{(2)}$, we obtain the desired result. \square


5 Blossoming trees for d -angulations

The aim of this section is to generalize bijections previously obtained for simple triangulations [33] and simple quadrangulations [20], that is triangulations and quadrangulations without loops nor multiple edges. In other words, triangulations and quadrangulations in which the contours of the faces are shortest cycles. More generally, the *girth* of a map is defined as the minimal length of its cycles. Obviously a d -angulation has girth at most d (except if it is a tree), hence simple triangulations and simple quadrangulations are exactly triangulations and quadrangulations with maximal girth. In the remaining sections, we aim at applying the general scheme to d -angulations of girth d , for any $d \geq 3$, and also to their following generalization. For any integers $p \geq d \geq 3$, define a *d -angulation of a p -gon* or a *p -gonal d -angulation* as a face-rooted plane map such that the contour of the root face is a simple cycle of length p and all non-root faces have degree d , see Fig. 17(a). We denote respectively \mathbf{M}_d and $\mathbf{M}_{d,p}$ the set of d -angulations and p -gonal d -angulations of girth d , with distinct root and outer faces.


We do not use here the canonical plane embedding of face-rooted maps with the root face as the outer face. On the contrary, from now on, we consider only *face-rooted plane maps in which the outer face and root face are different*. This convention yields obviously equivalent enumerative byproducts but proves to fit better.

5.1 Orientations for p -gonal d -angulations

For any $j, k \geq 0$, a j/k -orientation of a face-rooted map is defined as a k -fractional orientation such that for each root vertex v , $\text{out}(v) = k$, and $\text{out}(v) = j$ otherwise (see Fig. 16(b)). Bernardi and Fusy show in [4] that the existence of $\frac{d}{d-2}$ -orientations characterizes d -angulations of girth d , generalizing previous results obtained for triangulations [36] and quadrangulations [31]:



(a) A simple quadrangulation endowed with its minimal 2-orientation.



(b) A pentagulation of girth 5 endowed with its minimal 5/3-orientation.

Figure 16: Examples of $\frac{d}{d-2}$ orientations for d -angulations of girth d .

Theorem 5.1 (Schnyder [36], Ossona de Mendez [31], Bernardi and Fusy [4]). *Let $d \geq 3$ and M be a face-rooted d -angulation; then M admits a $\frac{d}{d-2}$ -orientation if and only if it has girth d . Moreover any such orientation is accessible.*

Besides if d is even, all the flows in this orientation are even.

Remark 5.2. For d even, the parity of the flows implies that maps of \mathbb{M}_d can be endowed with $\frac{d/2}{d/2-1}$ -orientations: for instance quadrangulations are naturally endowed with 2-orientations. For sake of conciseness we work here mainly with $\frac{d}{d-2}$ -orientations and distinguish odd and even cases only when needed.

For $p > d$, a simple application of Euler formula proves that a p -gonal d -angulation cannot admit a $\frac{d}{d-2}$ -orientation. The appropriate generalization is to define a (*pseudo*-) $\frac{d}{d-2}$ -orientation as a $(d-2)$ -fractional orientation in which the contour of the root face is a circuit of saturated edges and $\text{out}(v) = d$ for any non-root vertex v . Observe that for $p = d$, minimal pseudo- $\frac{d}{d-2}$ -orientations are minimal $\frac{d}{d-2}$ -orientations. Pseudo- $\frac{d}{d-2}$ -orientations characterize p -gonal d -angulations of girth d :

Proposition 5.3 ([4, Lemma 18]). *Let $p \geq d \geq 3$ be integers. A p -gonal d -angulation M admits a pseudo- $\frac{d}{d-2}$ -orientation if and only if it has girth d . In this case every pseudo- $\frac{d}{d-2}$ -orientation is accessible, and the sum of the outdegrees of the root vertices is equal to $(d-2)p + (p-d)$.*

Besides, there exists a unique minimal such pseudo- $\frac{d}{d-2}$ -orientation.

Remark 5.4. Since we allow accessibility and minimality to be defined relatively to two different faces (the root face and the outer face), this minimal pseudo- $\frac{d}{d-2}$ -orientation fits

in our scheme. In [4] however, accessibility and minimality have to be defined relatively to the same face. It implies [4, Proposition 19] that the existence of a minimal suitable orientation is conditioned on the map being non-separated (*i.e.* there cannot exist a cycle of girth length that separates the root face and the outer face).

5.2 Bijection for p -gonal d -angulations

To adapt Theorem 2.3 to our context, we need to introduce the following family of planar maps akin to forests of trees. A p -cyclic forest is a face-rooted plane map with two faces, the root one and the outer one, such that the border of the root face is a simple cycle of length p . Observe that a cyclic forest is nothing else but a cycle of sequences of planted trees.

Corollary 5.5. *Let $M = (V, E)$ be a plane face-rooted map with distinct root and outer faces and endowed with a minimal accessible orientation O . Then M admits a unique edge-partition $(\mathcal{T}_M, \mathcal{C}_M)$ such that:*

- *edges in \mathcal{T}_M form a spanning cyclic forest of M with same root-face as M , on which the restriction of O is accessible;*
- *edges in \mathcal{C}_M are saturated edges, and any of them turns clockwise around the unique cycle it forms with edges in \mathcal{T}_M .*

Proof. The map M admits such a unique partition of its edges if and only if \tilde{M} does, where \tilde{M} is the vertex-rooted map constructed from M by contracting its root face. We conclude by applying Theorem 2.3 to \tilde{M} . □

The two following definitions describe the corresponding blossoming objects in the setting of p -gonal d -angulations (see Fig. 17(a)).

For any integers $d \geq 3$ and $0 \leq i < d - 2$, a d -fractional tree of excess i is a planted blossoming tree endowed with an accessible $(d - 2)$ -fractional orientation such that the root leaf has outdegree i and each non-root vertex has outdegree $\text{out}(u) = d$, where each opening stem contributes $d - 2$. The set of d -fractional trees of excess i (resp. with n vertices) is denoted by $\mathsf{T}^{(i)}$ (resp. $\mathsf{T}_n^{(i)}$).

A p -gonal d -fractional forest is a p -cyclic forest, the planted trees of which are d -fractional trees. The sum of their excesses is moreover required to be equal to $p - d$. Observe that such a forest is naturally endowed with a pseudo- $\frac{d}{d-2}$ -orientation. The set of p -gonal d -fractional forests (resp. with n vertices) is denoted by $\tilde{\mathsf{F}}_d^p$ (resp. $\tilde{\mathsf{F}}_d^p(n)$). The closure of a p -gonal d -fractional forest is the natural counterpart of the closure of a blossoming map, in which the local closure of an opening stem creates a face of degree d .

The main theorem of this section is the following application of Corollary 2.4 to p -gonal d -angulations.

Theorem 5.6. *There exists a one-to-one constructive correspondence between p -gonal d -fractional forests with n vertices and p -gonal d -angulations of girth d with n non-root vertices.*




Figure 17: Example of the correspondence between p -gonal d -angulations of girth d and p -gonal d -fractional forests.

Proof. Corollary 5.5 and 2.4 entail that a p -gonal d -angulation endowed with its minimal $\frac{d}{d-2}$ -orientation can uniquely be opened into a p -cyclic d -fractional forest with additional closing stems. Deleting them yields a p -cyclic d -fractional forest, from which they can be retrieved using the condition that non-root faces have degree d .

The only point to check is that the closure of a p -gonal d -fractional forest is indeed a p -gonal d -angulation, *i.e.* that the degree of the outer face of the full closure is d . A d -fractional tree of excess i with k vertices has $\frac{2k+i}{d-2}$ opening stems. So a p -gonal d -fractional forest with $n + p$ vertices has $\frac{2n+p-d}{d-2}$ opening stems and its outer face has degree $2n + p$. Since each local closure reduces this degree by $d - 2$, it is exactly equal to d at the end of the closing process.

As mentioned in Section 2.3, the complexity of the opening procedure in the general case is quadratic. For p -gonal d -angulations, we were able to find a linear constructive algorithm. Its description is postponed to Section 6 for sake of clarity. \square

5.3 Enumerative consequences

By Theorem 5.6, the enumeration of d -angulations of girth d boils down to the enumeration of d -fractional forests and the latter reduces to the counting of d -fractional trees of excess i for $i \in \{0, \dots, d - 3\}$. Let T_i be the generating series of $\mathbb{T}^{(i)}$ according to the number of opening stems. The only way to generate and enumerate the elements of $\mathbb{T}_n^{(i)}$ seems to require a recursive approach, described in this section. It comes as no surprise that the recursive scheme we obtain is essentially the same as the one that counts the d -regular mobiles of [4]. We found it nevertheless interesting to note that the same

enumerative results can be obtained by our approach.

Following [4], for any positive integer j , let h_j be the polynomial in the variables t_1, t_2, \dots defined by:

$$h_j(t_1, t_2, \dots) := [x^j] \frac{1}{1 - \sum_{i>0} x^i t_i} = \sum_{r>0} \sum_{\substack{i_1, \dots, i_r > 0 \\ i_1 + \dots + i_r = j}} t_{i_1} \dots t_{i_r}.$$

In other terms, h_j is the generating function of compositions of j where the variable t_i keeps track of the number of parts of size i . Now, the generating series of the fractional trees of excess i can be obtained by a recursive decomposition. Cut the child of the root leaf to obtain a forest of trees (where one of the tree can be reduced to an opening stem) such that the sum of their excesses is equal to $i + 2$. In other words it can be written as a sequence $(s_0, t_1, s_1, \dots, t_l, s_l)$, where s_0, \dots, s_l are some sequences of trees with excess 0 (possibly reduced to a single vertex) and each of the t_i 's is a tree with positive excess. It yields the following system of equations:

$$T_i(x) = \frac{1}{1 - T_0} \cdot h_{i+2} \left(\frac{T_1}{1 - T_0}, \dots, \frac{T_{d-3}}{1 - T_0}, \frac{x}{1 - T_0} \right), \quad (3)$$

for $0 \leq i \leq d - 3$ and $T_i = 0$ otherwise. This set of equations characterizes T_0, T_1, \dots, T_{d-3} as formal power series. The constant coefficient of all these series is clearly equal to zero and the other coefficients can be computed recursively.

Proposition 5.7 (Bernardi and Fusy [4]). *For $p \geq d \geq 3$, the generating function $M_{d,p}(x)$ of corner-rooted p -gonal d -angulations of girth d with a marked outer face and counted according to the number of non-root faces is equal to:*

$$M_{d,p}(x) = x \left(\frac{1}{1 - T_0} \right)^p \cdot h_{p-d}^{(p)} \left(\frac{T_1}{1 - T_0}, \dots, \frac{T_{d-3}}{1 - T_0}, \frac{x}{1 - T_0} \right),$$

where $h_j^{(p)}$ is defined in (4) and the series T_0, \dots, T_{d-3} are characterized by (3).

For $p = d$, this equation reduces to:

$$M_{d,d}(x) = x \left(\frac{1}{1 - T_0} \right)^d$$

Proof. By Theorem 5.6, we have:

$$M_{d,p}(x) = x F_{d,p}(x),$$

where $F_{d,p}$ is the generating function of corner-rooted p -gonal d -fractional cyclic forests counted according to their number of opening stems. Erasing the edges of the root face of such a forest produces a p -tuple $P := (P_1, \dots, P_p)$ of sequences of fractional trees such that the total sum of their excesses is equal to $p - d$.

For each $1 \leq i \leq p$, the sequence P_i can be written as $P_i = (s_{i_0}, t_{i_1}, s_{i_1}, \dots, t_{i_l}, s_{i_l})$, such that each of the s_j is a sequence made of trees of excess 0 (possibly empty) and each t_j is a tree of positive excess. This decomposition gives the following formula for $F_{d,p}(x)$:

$$F_{d,p}(x) = \left(\frac{1}{1 - T_0} \right)^p \cdot h_{p-d}^{(p)} \left(\frac{T_1}{1 - T_0}, \dots, \frac{T_{d-3}}{1 - T_0}, \frac{x}{1 - T_0} \right),$$

where for $j \geq 0$, the polynomial $h_j^{(p)}$ is the generating function of p -tuples of compositions of integers, the sum of which is equal to j . For $p = 0$, $h_j^{(0)} = 1$, for $p = 1$ we retrieve the quantity h_j and more generally:

$$h_j^{(p)}(t_1, t_2, \dots) := [x^j] \frac{1}{(1 - \sum_{i>0} x^i t_i)^p}. \quad (4)$$

□

6 Fast opening of p -gonal d -angulations of girth d

6.1 General description of the algorithm

This section is devoted to the description of a linear-time algorithm that associates to each element of M_d its tree-and-closure partition or equivalently its d -fractional forest. We only give the proof in this setting, all steps extend easily to p -gonal d -angulations.

Let $M \in M_d$; recall that \mathcal{T}_M , \mathcal{C}_M and \mathcal{B}_M denote respectively its sets of tree edges, of closure edges and its d -fractional forest. Since the root face of M is not its outer face, the construction of Proposition 2.5 cannot be applied directly to obtain \mathcal{B}_M . The general idea of our algorithm is however to use a similar construction iteratively to identify and cut subtrees of \mathcal{B}_M .

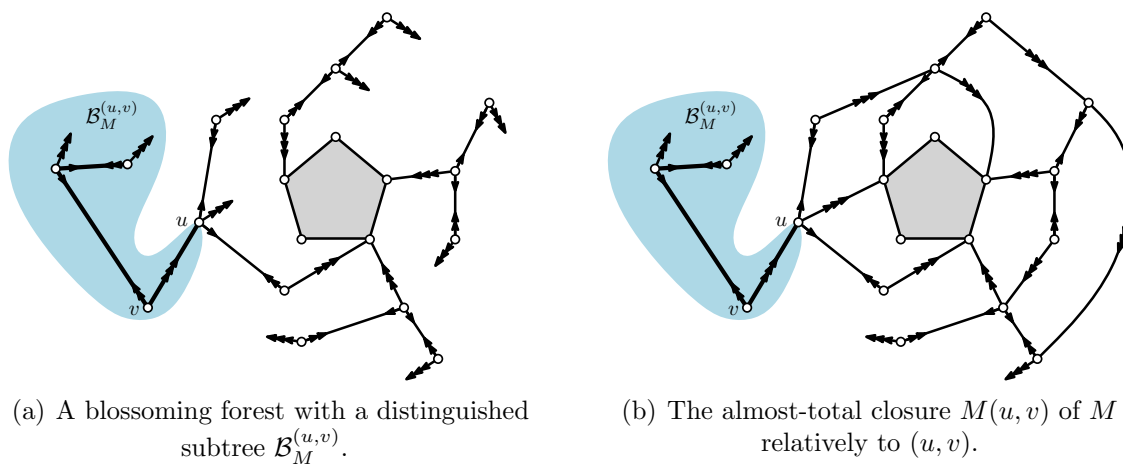


Figure 18: Almost-total closure.

More precisely, for any edge $e \in \mathcal{T}_M$, let us denote by $\mathcal{B}_M^{(e)}$ the (blossoming) subtree of \mathcal{B}_M planted at e . Define the *almost-total closure* $M(e)$ of M relatively to e as the maximal partial closure of \mathcal{B}_M such that the border of the subtree $\mathcal{B}_M^{(e)}$ (including both sides of e) still lies on the outer face of $M(e)$ (see Fig. 18(b)). Deleting $\mathcal{B}_M^{(e)}$ yields a blossoming map with a canonically marked corner, denoted by $M_{/e}$. It is clear that M is the closure of the blossoming map obtained by grafting $\mathcal{B}_M^{(e)}$ on $M_{/e}$ in its marked corner. Then,

Claim 6.1. *For any map $M \in \mathbf{M}_d$ and any edge $e \in \mathcal{T}_M$, \mathcal{B}_M is the forest obtained by grafting $\mathcal{B}_M^{(e)}$ in the marked corner of $\mathcal{B}_{M_{/e}}$.*

Proof. It follows directly from the uniqueness of \mathcal{B}_M (proved in Theorem 5.6). □

Claim 6.1 entails that computing the blossoming forest of M can be achieved in three steps: first, identify $M(e)$ for some tree-edge e , then, compute recursively the blossoming forest of $M_{/e}$, and last, graft $\mathcal{B}_M^{(e)}$ in the marked corner of $\mathcal{B}_{M_{/e}}$. This is precisely described by Algorithms 1 to 3 page 36. The correction of Algorithm 3 relies on the three following propositions, whose proofs are postponed to the next subsections:

Proposition 6.2. *Algorithm 1 returns a lineage path $(v_0, e_1, v_1, \dots, e_\ell, v_\ell)$ in \mathcal{T}_M oriented towards the root polygon and opens a subset C of \mathcal{C}_M such that each e_i is an outer edge of the resulting blossoming map.*

Proposition 6.3. *Algorithm 2 returns the subtree $\mathcal{B}_M^{(e)}$ for some edge e and opens the corresponding closure edges of M so as to compute $M(e)$.*

Proposition 6.4. *There exists a partial order on blossoming trees such that:*

- for any $i \in \llbracket 0, d-2 \rrbracket$, there exists a unique minimum $t_i^{(d)}$ among trees of excess i ,
- the tree T produced by Algorithm 2 does not belong to the family $\left(t_i^{(d)}\right)_{0 \leq i \leq d-2}$.

By Proposition 6.4, the sequence of maps on which Algorithm 3 is recursively performed is decreasing, and hence converges eventually to the trivial map. Following the sequence backwards provides a recursive construction of the tree-and-closure partition of M . A complete run of the algorithm for the pentagulation of Fig.16(b) is shown in Figure 19: Algorithm 1 in Fig.19(a), Algorithm 2 in Fig.19(b), the replacement of the subtree in Fig.19(c), and the recursive call afterwards.

Theorem 6.5. *For any input map $M \in \mathbf{M}_d$, Algorithm 3 computes its blossoming forest in linear time with respect to the size of M .*

Algorithm 1: LineagePath(M)

Input: a map $M \in \mathbf{M}_d$ endowed with its minimal $\frac{d}{d-2}$ -orientation

Output: a lineage path $(v_0, e_1, v_1, \dots, e_\ell, v_\ell)$ in \mathcal{T}_M oriented towards the root face

Side effect: partial opening, such that each e_i is an outer edge

let (u, v) be an outer closure edge; initialize C to $\{(u, v)\}$, i to 0 and set $v_0 = v$;

repeat

 update u to the previous vertex around v in clockwise order;

if (u, v) is saturated towards v **then** add it to C **else** update (u, v) to (v, u) ,

 increment i and set $v_i = v$ and $e_i = (u, v)$

until $i = \lceil d/2 \rceil$ or v is a root vertex ;

replace each edge of C by an opening stem in the appropriate corner;

return $(v_0, e_1, v_1, \dots, e_\ell, v_\ell)$

Algorithm 2: DetachableSubtree(M)

Input: a map $M \in \mathbf{M}_d$, endowed with its minimal $\frac{d}{d-2}$ -orientation

Output: the set T of inner edges of a subtree $\mathcal{B}_M^{(e_\ell)}$

Side effect: M is partially opened into $M(e_\ell)$

let $(v_0, e_1, v_1, \dots, e_\ell, v_\ell) = \text{LINEAGEPATH}(M)$;

initialize T and C to \emptyset , and (u, v) to $(v_\ell, v_{\ell-1})$;

repeat

if $(u, v) \notin C \cup T$ **then**

if (u, v) is saturated towards v **then** add (u, v) to C **else** add (u, v) to T

if $(u, v) \in T$ **then** update (u, v) to (v, u) update v to the next vertex (in clockwise order) around u ;

until $(u, v) = (v_{\ell-1}, v_\ell)$;

replace each edge of C by an opening stem in the appropriate corner;

return T

Algorithm 3: BlossomingForest(M)

Input: a map $M \in \mathbf{M}_d$, endowed with its minimal $\frac{d}{d-2}$ -orientation


Output: its blossoming forest \mathcal{B}_M

let $T = \text{DETACHABLESUBTREE}(M)$ and t be the minimal tree with same excess ;


replace T by t in M and perform the closure;

let $F = \text{BLOSSOMINGFOREST}(M)$ and replace t by T in F ;


return F ;




(a) Identification of a lineage path of length 3,




(b) the subtree $\mathcal{B}_M^{(e_3)}$,




(c) replacing $\mathcal{B}_M^{(e_3)}$ by $t_1^{(5)}$,



(d) the lineage path reaches the root pentagon,



(e) according to Remark 6.10, the DFS identifies \mathcal{B}_{M/e_3} entirely,



(f) which enables to reconstruct \mathcal{B}_M .

Figure 19: Execution of the fast opening algorithm on the pentagulation of Fig.16(b).

6.2 Identification of a subtree of \mathcal{B}_M

In this subsection, we prove Proposition 6.3. First observe that, assuming Proposition 6.2, the proof of Proposition 6.3 is easily derived from the proof of Proposition 2.5, given in [3]: Algorithm 2 is actually nothing but an adapted version of the depth-first search process described in Section 2.3, in a slightly looser case.

Therefore we only need to prove Proposition 6.2. To this purpose, we need the following technical result:

Definition 6.6. Let u and w be two outer vertices of a blossoming map. The clockwise distance between u and w is the number of edge sides in the shortest clockwise path from u to w along the border of its outer face, see Fig.20(a).

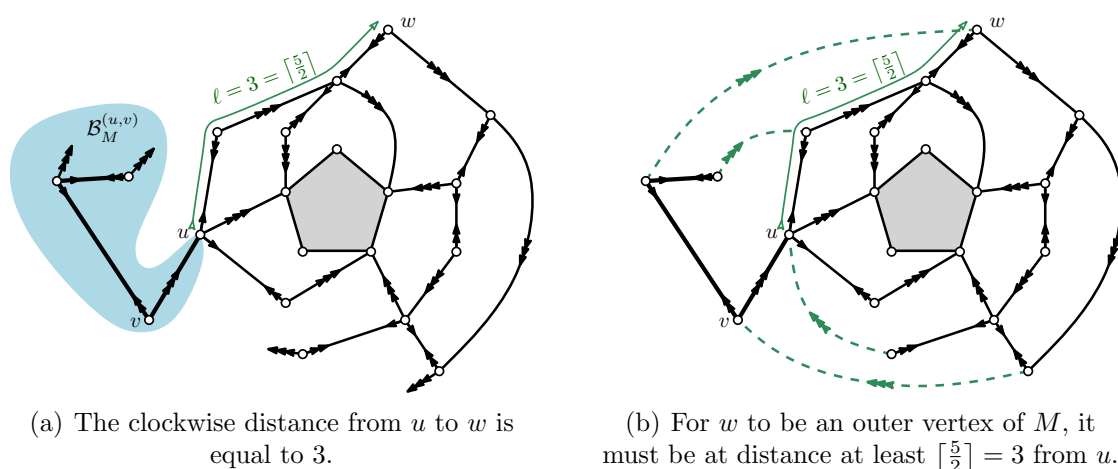


Figure 20: Illustration of Lemma 6.7 for a pentagulation.

Lemma 6.7. Let $M \in \mathcal{M}_d$ and u, v in M such that v is a child of u in \mathcal{B}_M . Let w be an outer vertex of M that does not belong to $\mathcal{B}_M^{(u,v)}$. Then, the clockwise distance between u and w in $M(u, v)$ cannot be smaller than $\lceil d/2 \rceil$, see Fig.20(b).

Proof. Let $T = \mathcal{B}_M^{(u,v)}$, n be its number of edges, and $0 \leq i \leq d - 3$ be the flow from u to v . A simple combinatorial argument yields that the number of stems in T is equal to $(2n + i)/(d - 2)$. It implies that performing all the local closures of stems in T requires $2n + i + 1$ sides of edges (each stem needs $d - 1$ sides of edges, and each new created edge but the last one can be used in a further closure).

Now the number of sides of edges of T is equal to $2n$, but not all of them can be used in local closures involving outgoing stems of T . Indeed the flow between a vertex in T at distance h of u and its parent is at least $d - i - 2h$, hence such a vertex can carry a stem only if $h > \lceil (d - 2 - i)/2 \rceil$. Therefore the first stem explored during a DFS contour of T is discovered after at least $\lceil (d - 2 - i)/2 \rceil$ steps. Hence the number of sides of edges

available for closures in T is at most $2n - \lceil (d-2-i)/2 \rceil$. It follows that closing the stems of T requires at least $\lceil d/2 \rceil$ additional sides of edges; hence outer vertices of $M(u, v)$ at clockwise distance less than $\lceil d/2 \rceil$ from u are separated from the outer face of M by at least one closure edge. \square

This lemma implies that if (u, v) is a tree-edge and u is the parent of v , then the outer face cannot lie on the left of (v, u) . In particular:

Corollary 6.8. *Let $M \in \mathbf{M}_d$. Any outer clockwise saturated edge of M (and there exists at least one of them) is a closure edge, and other outer edges (if any) have the outer face on their right when walking from a child to its parent.*

Proof of Proposition 6.2. This enables to choose an outer closure edge e to initialize Algorithm 1. Let e_1, \dots, e_ℓ be the edges of the path returned by Algorithm 1, and let v_0, \dots, v_ℓ be the corresponding vertices. Let us now prove by induction on the number of iterations that Algorithm 1 correctly separates tree from closure edges. Assume that it is true after the k first steps of the Algorithm 1, and let i_k be the corresponding value of i . Assume that $i_k < l$, then $e = (v_{i_k}, w)$ for some vertex w and v_0 is an outer vertex of M at clockwise distance $i_k < d/2$ from v_{i_k} around $M \setminus C$. By Lemma 6.7, it prevents w from being a child of v_{i_k} . Hence if e is saturated towards v_i , it implies that e is a closure edge; otherwise w is the parent of v_{i_k} . Moreover, from the construction, w is necessarily an outer vertex of $M \setminus C$. This ends the proof. \square

Let us end this section with two comments.

Remark 6.9. The depth-first search process of Algorithm 2 could have been applied on e_1 rather than e_ℓ . But relying on this whole lineage path, we can identify a much bigger subtree. This is actually a key point for proving the termination of recursive calls in Algorithm 3, as explained in Section 6.4.

Remark 6.10. If the lineage path reaches the root face, the depth-first search algorithm can be continued so as to identify \mathcal{B}_M entirely, as in Fig.19(e). Indeed in this case, let us open the closure edges pointing to the lineage path, collapse the root polygon into a single root vertex, and choose as root-corner of this new map the corner of the root vertex incident to e_ℓ and to the outer face. Then we are precisely in the setting of Proposition 2.5.

6.3 Triangulations and quadrangulations

Simple triangulations and simple quadrangulations constitute a much simpler case than general d -angulations of girth d , since all the edges are saturated in their minimal $\frac{d}{d-2}$ -orientation.

Given a d -angulation M of girth d with $d = 3$ or $d = 4$, Algorithm 1 identifies two tree-edges e_1 and e_2 , and Algorithm 2 computes $M(e_2)$. $M(e_2)$ can be decomposed into $\mathcal{B}_M^{(e_2)}$ and $M_{/e_2}$. Since all the edges are saturated, the closure of $M_{/e_2}$ is itself a d -angulation, smaller than M , on which Algorithm 2 can be iteratively applied.

At each iteration the number of edges of the map decreases, the sequence of maps reaching eventually the trivial map reduced to a single cycle. Following the sequence

backwards provides a recursive construction of the tree-and-closure partition of the desired map. Hence:

Proposition 6.11. *For any simple triangulation or simple quadrangulation M , Algorithm 4 computes \mathcal{B}_M in linear time.*

Algorithm 4: BlossomingForest3Or4angulation(M)

Input: a simple triangulation or quadrangulation M

Output: its blossoming forest \mathcal{B}_M

let $T = \text{DETACHABLESUBTREE}(M)$;

delete T in M and perform the closure;

let $F = \text{BLOSSOMINGFOREST3OR4ANGULATION}(M)$;

graft T in F in the appropriate corner;

return F ;

6.4 How to graft small subtrees

This section is dedicated to the proof of the termination of Algorithm 3 in the general case of d -angulations of girth d for $d \geq 5$. Roughly speaking, we want to iterate Algorithm 2 on the closure of $M_{/e_\ell}$, but this map is not necessarily a d -angulation: this can happen if the edge e_ℓ is not saturated, so that its deletion creates a vertex u of outdegree less than d . To circumvent this problem, Algorithm 3 grafts at u a smaller d -fractional subtree of appropriate excess, performs the closure, and recursively calls itself on the resulting (smaller) map in \mathbf{M}_d . The existence of such a smaller subtree is guaranteed by Proposition 6.4, which we shall now prove.

A natural partial order on blossoming trees is the order induced by sizes: a tree t_1 is declared to be smaller than a tree t_2 if t_1 has less edges than t_2 . But this order has to be slightly refined so as to define unambiguously a unique minimal element $t_i^{(d)}$ in each $\mathbb{T}^{(i)}$. Precisely:

- we adopt the convention that the empty tree has excess 0, hence it is equal to $t_0^{(d)}$;
- for d even or i odd, there exists also a unique minimal element $t_i^{(d)}$ in $\mathbb{T}^{(i)}$, the one made of a path of length $(d - 2 - i)/2$ and one stem;
- for d odd and i even, exactly two trees have minimal size, both made of a path of length $d - 2 - i/2$ with two stems; consider their contour words on $\{b, e, \bar{e}\}$: they are respectively equal to $w_1 = e^p b e^q b \bar{e}^{(p+q)}$ and $w_2 = e^{p+q} b \bar{e}^q b \bar{e}^p$, with $p = \frac{1}{2}(d - i - 1)$ and $q = \frac{1}{2}(d - 3)$. We refine the ordering on $\mathbb{T}^{(i)}$ by defining the minimal element $t_i^{(d)}$ as the one with contour word w_1 .

This notion of ordering extends naturally to d -fractional forests: $F_1 < F_2$ if and only if F_1 can be obtained from F_2 by replacing a sequence of subtrees by smaller ones. The maps in \mathbf{M}_d inherit the ordering of their blossoming forests. The unique minimal map of \mathbf{M}_d is the map reduced to a simple cycle of length d .

Proof of Proposition 6.4. The case $i = 0$ is clear, and includes in particular the case where v_ℓ is a root vertex. From the exploration of the lineage path in Algorithm 1, any other tree that Algorithm 2 may return has at least $d/2$ edges before the first stem in its contour word. Since $t_i^{(d)}$ has less than $d/2$ edges for d even or i odd, this proves the result in those cases. For d odd and i even, our particular choice for $t_i^{(d)}$ enables also to conclude. \square

6.5 Proof of Theorem 6.5 – Complexity of the algorithm

Propositions 6.2 and 6.3 prove that the output of the algorithm is the tree-and-closure partition of the input. It remains to prove that the algorithm can be implemented in linear time.

Let (M_k) be the sequence of maps obtained in the execution of the algorithm. Observe that the length of the sequence (M_k) is bounded by $n/2 + 2n/(d-1)$; indeed, for each k , M_{k+1} has at least two edges less than M_k , except in the case where $t_i^{(d)}$ is grafted instead of the other tree in $\mathbf{T}^{(i)}$ of minimal size; the latter case happens at most $2n/(d-1)$ times, since any edge appears at most once in such a subtree.

The cost of step k of the algorithm is clearly linear respectively to the number of edges of (M_k) , since only exploration processes are performed, with a bounded cost per visited edge. Unfortunately, this is not precise enough to evaluate correctly the total complexity of the algorithm, and we need to look more carefully at the number of visited edges of each type at each step. The cost of step k can be decomposed into three parts:

- the contribution $c(k)$ of the handling of closure edges discovered during the exploration of the lineage path, described in Algorithm 1;
- the contribution $t(k)$ of the exploration of the subtree T_k , including the computation of the lineage path, obtained by Algorithm 2;
- the contribution $n(k)$ of the construction of M_{k+1} , including the deletion of $\mathcal{B}_{M_k}^{(e_\ell)}$ and the grafting of the appropriate $t_i^{(d)}$, described in Section 6.4.

The number of edges in T_k is the sum of the size of the adequate $t_i^{(d)}$ – which is bounded by d – and of the size difference between M_{k+1} and M_k . Hence, the sum over k of $t(k)$ is linear in n . Likewise the sum over k of $n(k)$ is linear in n .

To deal with $c(k)$, we need to be slightly more careful. Indeed, the number of closure edges met during the exploration of the lineage path at any given step k can only be bounded by n , resulting in an upper bound of order n^2 for the total complexity. However, let us observe that a bundle of closure edges that end at a same vertex of the lineage

path in M_k also end at a same vertex in M_{k+1} (provided that they still belong to M_{k+1}). Therefore, we assume that M is encoded in such a way that closure edges ending at the same vertex are stored in the same object, called a *bundle*, to be used during the further exploration processes. In particular, the number of bundles encountered during Algorithm 1 is bounded by d .

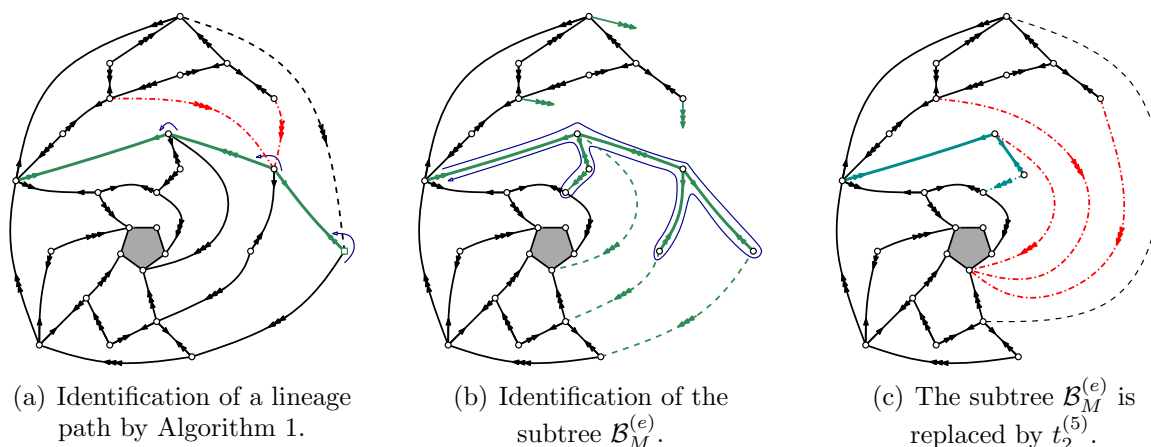


Figure 21: Merging of a bundle (made of the two red dash-dotted edges in Fig.(a)) with one of the closure edges of $t_2^{(5)}$ in Fig.(c).

More precisely, in addition to the classical encoding of a map as a collection of half-edges cyclically ordered around vertices and faces, we introduce *bundles* of half-edges. Each bundle contains a doubly-linked list of consecutive closure half-edges ending at the same vertex. Each half-edge belongs to at most one bundle. At the beginning of the algorithm, there is no bundle. Then, each time an half-edge is identified as the extremity of a closure edge, it is either added to an existing bundle, or used to initialize a new one. If two bundles happen to become consecutive, they are merged. This can occur during Algorithm 1 or after the grafting of the appropriate $t_i^{(d)}$ in Algorithm 3, see Fig.21. The only other operation to be performed on a bundle is the deletion of its half-edges that belong to the subtree $\mathcal{B}_{M_k}^{(e_i)}$ identified by Algorithm 2.

The total number of closure edges that need to be dealt with during Algorithm 3 is linear in the number n of vertices: the number of closure edges of the initial d -angulation is linear in n , and at each step, at most two new closure edges are created. The merging of two bundles and the deletion of consecutive half-edges from one bundle can be both implemented in constant time. Hence the sum over k of $c(k)$ is linear in n .

Acknowledgments

We would like to thank Éric Fusy and Gilles Schaeffer for fruitful discussions and a referee whose careful reading helped us to improve significantly the exposition of our work. We acknowledge the support of the ERC under the agreement “ERC StG 208471 - ExploreMap” and of the ANR under the agreement “ANR 12-JS02-001-01 - Cartaplus”.

References

- [1] L. Addario-Berry and M. Albenque. The Brownian map is the scaling limit of simple triangulations and quadrangulations. *arXiv:1306.5227*, 2013.
- [2] R.J. Baxter. Dichromatic polynomials and potts models summed over rooted maps. *Ann. Comb.*, 5(1):17–36, 2001.
- [3] O. Bernardi. Bijective counting of tree-rooted maps and shuffles of parenthesis systems. *Electron. J. Combin.*, 14:#R9, 2007.
- [4] O. Bernardi and É. Fusy. A bijection for triangulations, quadrangulations, pentagulations, etc. *J. Combin. Theory Ser. A*, 119(1):218–244, 2012.
- [5] O. Bernardi and É. Fusy. Unified bijections for maps with prescribed degrees and girth. *J. Combin. Theory Ser. A*, 119(6):1351–1387, 2012.
- [6] N. Bonichon, M. Bousquet-Mélou, and É. Fusy. Baxter permutations and plane bipolar orientations. *Sém. Lothar. Combin.*, 61A:B61Ah, 2010.
- [7] M. Bousquet-Mélou and A. Jehanne. Planar maps and algebraic series: a generalization of the quadratic method. *J. Combin. Theory Ser. B*, 96:623–672, 2006.
- [8] M. Bousquet-Mélou and G. Schaeffer. Enumeration of planar constellations. *Adv. in Appl. Math.*, 24(4):337–368, Nov 2000.
- [9] J. Bouttier, Ph. di Francesco, and E. Guitter. Census of planar maps: from the one-matrix model solution to a combinatorial proof. *Nuclear Phys. B*, 645(3):477–499, 2002.
- [10] J. Bouttier, Ph. di Francesco, and E. Guitter. Planar maps as labeled mobiles. *Electron. J. Combin.*, 11(1):#R69, 2004.
- [11] J. Bouttier, Ph. Di Francesco, and E. Guitter. Blocked edges on Eulerian maps and mobiles: Application to spanning trees, hard particles and the Ising model. *J. Phys. A*, 40:7411–7440, 2007.
- [12] E. Brézin, C. Itzykson, G. Parisi, and J.-B. Zuber. Planar diagrams. *Comm. Math. Phys.*, 59(1):35–51, 1978.
- [13] G. Chapuy, M. Marcus, and G. Schaeffer. A bijection for rooted maps on orientable surfaces. *SIAM J. Discrete Math.*, 23(3):1587–1611, 2009.
- [14] Ph. Chassaing and G. Schaeffer. Random planar lattices and integrated superbrownian excursion. *Probab. Theory Related Fields*, 128:161–212, 2004.
- [15] R. Cori and B. Vauquelin. Planar maps are well labeled trees. *Canad. J. Math*, 33(5):1023–1042, 1981.
- [16] H. de Fraysseix, P. Ossona de Mendez, and P. Rosenstiehl. Bipolar orientations revisited. *Discrete Appl. Math.*, 56(2-3):157–179, 1995.
- [17] E. Duchi, D. Poulalhon, and G. Schaeffer. Uniform random sampling of simple branched coverings of the sphere by itself. In *Proc. 25th ACM-SIAM Symp. (SODA)*, pages 294–304, 2014.
- [18] S. Felsner. Lattice structures from planar graphs. *Electron. J. Combin.*, 11(1):#R15, 2004.

- [19] S. Felsner, É. Fusy, M. Noy, and D. Orden. Bijections for baxter families and related objects. *J. Combin. Theory, Ser. A*, 118(3):993–1020, 2011.
- [20] É. Fusy. *Combinatoire des cartes planaires et applications algorithmiques*. PhD thesis, École Polytechnique, 2007.
- [21] É. Fusy, D. Poulalhon, and G. Schaeffer. Bijective counting of plane bipolar orientations and Schnyder woods. *European J. Combin.*, 30(7):1646–1658, 2009.
- [22] I.M. Gessel and X. Viennot. Determinants, paths, and plane partitions. 1989.
- [23] I.P. Goulden and D.M. Jackson. *Combinatorial enumeration*, volume 19. Wiley New York, 1983.
- [24] F. Harary, G. Prins, and W.T. Tutte. The number of plane trees. *Indag. Math.*, 26:319–329, 1964.
- [25] D.M. Jackson. Counting cycles in permutations by group characters, with an application to a topological problem. *Trans. Amer. Math. Soc.*, 299(2):785–801, 1987.
- [26] J.-F. Le Gall. The topological structure of scaling limits of large planar maps. *Invent. Math.*, 169(3):621–670, 2007.
- [27] J.-F. Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.*, 41:2880–2960, 2013.
- [28] J.-F. Marckert and A. Mokkadem. Limit of normalized quadrangulations: the Brownian map. *Ann. Probab.*, 34(6):2144–2202, 2006.
- [29] G. Miermont. Tessellations of random maps of arbitrary genus. *Ann. Sci. Éc. Norm. Supér.*, 42(5):725–781, 2009.
- [30] G. Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.*, 210:319 – 401, 2013.
- [31] P. Ossona de Mendez. *Orientations bipolaires*. PhD thesis, École des Hautes Etudes en Sciences Sociales, Paris, 1994.
- [32] D. Poulalhon and G. Schaeffer. A bijection for triangulations of a polygon with interior points and multiple edges. *Theoret. Comput. Sci.*, 307(2):385–401, 2003.
- [33] D. Poulalhon and G. Schaeffer. Optimal coding and sampling of triangulations. *Algorithmica*, 46(3):505–527, 2006.
- [34] G. Schaeffer. Bijective census and random generation of eulerian planar maps with prescribed vertex degrees. *Electron. J. Combin.*, 4(1):#R20, 1997.
- [35] G. Schaeffer. *Conjugaison d’arbres et cartes combinatoires aléatoires*. PhD thesis, Université Bordeaux I, 1998.
- [36] W. Schnyder. Planar graphs and poset dimension. *Order*, 5:323–343, 1989.
- [37] G. t’Hooft. A planar diagram theory for strong interactions. *Nuclear Phys. B*, 72(3):461–473, 1974.
- [38] W.T. Tutte. A census of planar maps. *Canad. J. Math.*, 15:249–271, 1963.