Enumeration of planar constellations

Mireille Bousquet-Mélou
CNRS, LaBRI, Université Bordeaux 1
351 cours de la Libération
33405 Talence Cedex, France
bousquet@labri.u-bordeaux.fr

Gilles Schaeffer*
CNRS, Loria, Campus Scientifique
615, rue du Jardin Botanique – B.P. 101
54602 Villers–lès–Nancy Cedex, France
Gilles.Schaeffer@loria.fr

Abstract

The enumeration of transitive ordered factorizations of a given permutation is a combinatorial problem related to singularity theory. Let $n \geqslant 1$, and let σ_0 be a permutation of $\mathfrak{S}_n$ having d_i cycles of length i, for $i \geqslant 1$. Let $m \geqslant 2$. We prove that the number of m-tuples $(\sigma_1, \ldots, \sigma_m)$ of permutations of $\mathfrak{S}_n$ such that:

- $\sigma_1 \sigma_2 \cdots \sigma_m = \sigma_0$,
- the group generated by $\sigma_1, \ldots, \sigma_m$ acts transitively on $\{1, 2, \ldots, n\}$,
- $\sum_{i=0}^{m} c(\sigma_i) = n(m-1) + 2$, where $c(\sigma_i)$ denotes the number of cycles of σ_i,

is

$$m \, \frac{[(m-1)n-1]!}{[(m-1)n - c(\sigma_0) + 2]!} \prod_{i \geqslant 1} \left[i \binom{mi-1}{i} \right]^{d_i}.$$

A one-to-one correspondence relates these m-tuples to some rooted planar maps, which we call constellations and enumerate via a bijection with some bicolored trees. For $m = 2$, we recover a formula of Tutte for the number of Eulerian maps. The proof relies on the idea that maps are conjugacy classes of trees and extends the method previously applied to Eulerian maps by the second author.

Our result might remind the reader of an old theorem of Hurwitz, giving the number of m-tuples of *transpositions* satisfying the above conditions. Indeed, we show that our result implies Hurwitz' theorem. We also briefly discuss its implications for the enumeration of nonequivalent coverings of the sphere.

1 Introduction

Let σ_0 be a permutation in the symmetric group $\mathfrak{S}_n$. An *ordered factorization* of σ_0 is an m-tuple $(\sigma_1, \ldots, \sigma_m)$ of permutations of $\mathfrak{S}_n$ such that $\sigma_1 \sigma_2 \cdots \sigma_m = \sigma_0$.

The enumeration of ordered factorizations of a fixed permutation is a widely studied problem. Its numerous different motivations make it very versatile, and give rise to different kinds of conditions that can be imposed on the factors. Here are some conditions often met in the literature.

- **The cyclic type of the factors.** One can decide that each factor σ_i must be taken inside a prescribed conjugacy class of $\mathfrak{S}_n$: in this case, one is merely trying to compute the *connection coefficients* of the symmetric group. A very general formula can be given in terms of characters [24, p.68]. The *rank* $n - c(\sigma)$ of a permutation σ having $c(\sigma)$ cycles gives the length of the shortest ordered factorization of σ into transpositions. The rank being clearly sub-additive, we observe that the connection coefficient is zero unless

$$\sum_{i=1}^{m} [n - c(\sigma_i)] \geqslant n - c(\sigma_0),$$

where $c(\sigma_i)$ denotes the number of cycles of σ_i (which only depends on its conjugacy class). Equivalently,

$$\sum_{i=1}^{m} c(\sigma_i) \leqslant n(m-1) + c(\sigma_0). \tag{1}$$

This condition is necessary, but not sufficient, for the corresponding connection coefficient to be non zero.

*This work was achieved while the author was a member of the LaBRI.

$$\sum_{i=1}^{m} c(\sigma_i) = n(m-1) + c(\sigma_0), \tag{2}$$

which is minimal in terms of the ranks of the factors. This problems amounts to computing the *top connection coefficients* of the symmetric group [10]. The most celebrated result in this field corresponds to the case where all factors are transpositions and σ_0 is an n-cycle. The extremality condition (2) becomes $m = n-1$, and the number of such factorizations is n^{n-2}, the number of Cayley trees [4, 11, 21].

- **The transitivity condition** requires that the group generated by $\sigma_1, \ldots, \sigma_m$ acts transitively on $\{1, 2, \ldots, n\}$. This condition finds its origin in the link between ordered factorizations and branched coverings of Riemann surfaces: roughly speaking, the transitivity condition is implied by the connectivity of the surfaces. This condition is widely considered, and will also be adopted in this paper: our factorizations will correspond to branched coverings of the two-dimensional sphere by itself (see Section 3).

- **The transitive minimality condition.** Most importantly, the upper bound on $\sum c(\sigma_i)$ given by (1) is no longer sharp under the transitivity condition. For instance, all transitive factorizations of a permutation σ_0 into m transpositions satisfy the following inequality [8, 25]:

$$m \geqslant n + c(\sigma_0) - 2 \tag{3}$$

which is stronger than the inequality $m \geqslant n - c(\sigma_0)$ provided by (1). From (3), we easily derive the following inequality, valid for all transitive factorisations of σ_0:

$$\sum_{i=1}^{m} c(\sigma_i) \leqslant n(m-1) - c(\sigma_0) + 2,$$

which is stronger than (1). It can be understood in terms of the genus of the underlying Riemann surfaces.

We shall focus on *extremal transitive* factorisations. The case where all factors are transpositions was solved long time ago by Hurwitz [14] (see also [6, 8, 25]).

Theorem 1.1 (Hurwitz) *Let $n \geqslant 1$. Let σ_0 be a permutation of $\mathfrak{S}_n$ having d_i cycles of length i, for $i \geqslant 1$. Then the number of m-tuples $(\tau_1, \ldots, \tau_m)$ of transpositions of $\mathfrak{S}_n$ such that:*

- $\tau_1 \tau_2 \cdots \tau_m = \sigma_0$,
- *the group generated by $\tau_1, \ldots, \tau_m$ acts transitively on $\{1, 2, \ldots, n\}$,*
- $m = n + c(\sigma_0) - 2$, *where $c(\sigma_0)$ denotes the number of cycles of σ_0,*

is

$$H_{\sigma_0} = n^{c(\sigma_0)-3} (n + c(\sigma_0) - 2)! \prod_{i \geqslant 1} \left[\frac{i^i}{(i-1)!} \right]^{d_i}.$$

In this paper, we count extremal transitive factorizations *regardless of the cyclic type of the factors.* Our main theorem follows. We shall see that it implies Hurwitz' theorem.

Theorem 1.2 *Let $n \geqslant 1$. Let σ_0 be a permutation of $\mathfrak{S}_n$ having d_i cycles of length i, for $i \geqslant 1$. For $m \geqslant 0$, let $G_{\sigma_0}(m)$ denote the number of m-tuples $(\sigma_1, \ldots, \sigma_m)$ of permutations of $\mathfrak{S}_n$ such that:*

- $\sigma_1 \sigma_2 \cdots \sigma_m = \sigma_0$,
- *the group generated by $\sigma_1, \ldots, \sigma_m$ acts transitively on $\{1, 2, \ldots, n\}$,*
- $\sum_{i=0}^{m} c(\sigma_i) = n(m-1) + 2$, *where $c(\sigma_i)$ denotes the number of cycles of σ_i.*

Then for $m \geqslant 2$,

$$G_{\sigma_0}(m) = m \, \frac{[(m-1)n - 1]!}{[(m-1)n - c(\sigma_0) + 2]!} \prod_{i \geqslant 1} \left[i \binom{mi-1}{i} \right]^{d_i}.$$

that one of the factors is a transposition, and the other is a 3-cycle. We find $G_{\sigma_0}(2) = 4$. Indeed, the four minimal transitive factorizations are $\sigma_0 = (13)(123) = (23)(132) = (132)(13) = (123)(23)$. (We multiply permutations from right to left, as we compose functions.) $\square$

Let us call an ordered factorization *proper* if none of its factors is the identity. Note that if we remove a trivial factor from a minimal transitive m-factorization, we obtain a minimal transitive $(m-1)$-factorization. This allows us to use the inclusion-exclusion principle to express the number $F_{\sigma_0}(m)$ of proper minimal transitive m-factorizations of σ_0, for $m \geqslant 0$:

$$F_{\sigma_0}(m) = \sum_{k=0}^{m} (-1)^{m-k} \binom{m}{k} G_{\sigma_0}(k). \tag{4}$$

Observe that a proper factorization $(\sigma_1, \ldots, \sigma_m)$ satisfies $\sum_{i=0}^{m} c(\sigma_i) \leqslant c(\sigma_0) + m(n-1)$. If it is also transitive and minimal, then $\sum_{i=0}^{m} c(\sigma_i) = n(m-1)+2$ and thus, $m \leqslant n + c(\sigma_0) - 2$. Moreover, the choice $m = n + c(\sigma_0) - 2$ forces $c(\sigma_i)$ to be $n-1$, for $1 \leqslant i \leqslant m$, so that each factor is a transposition. This shows that the number of minimal transitive factorizations into transpositions, evaluated by Hurwitz, is

$$H_{\sigma_0} = F_{\sigma_0}(d)$$

where $d = n + c(\sigma_0) - 2$.

It is not difficult to check that, for each $\sigma_0 \in \mathfrak{S}_n$, the expression of $G_{\sigma_0}(m)$ given in Theorem 1.2 is a *polynomial* $P_{\sigma_0}(m)$ in m, of degree $d = n + c(\sigma_0) - 2$, and that the identity $G_{\sigma_0}(m) = P_{\sigma_0}(m)$ actually holds for $m \geqslant 0$. Defining the difference operator Δ by $\Delta P(x) = P(x+1) - P(x)$, we can rewrite (4) as follows:

$$F_{\sigma_0}(m) = \Delta^m P_{\sigma_0}(0).$$

Observe that $\Delta^d(x^k) = 0$ if $k < d$ and $\Delta^d(x^d) = d!$. This implies that H_{σ_0} is, up to a factorial, the leading coefficient of $P_{\sigma_0}(x)$:

$$\begin{aligned} H_{\sigma_0} &= F_{\sigma_0}(d) \\ &= \Delta^d P_{\sigma_0}(0) \\ &= d![x^d]P_{\sigma_0}(x) \\ &= (n + c(\sigma_0) - 2)!\; n^{c(\sigma_0)-3} \prod_{i \geqslant 1} \left[\frac{i^i}{(i-1)!} \right]^{d_i}. \end{aligned}$$

This is exactly Hurwitz' theorem.

Many of the enumeration problems mentioned above have an alternative description in terms of trees, maps, or hypermaps. Our theorem is not an exception to this rule: in Section 2, we describe a family of maps, called *constellations*, which are in one-to-one correspondence with minimal transitive factorizations. In Section 3, we give more details on the connection with branched covers of the sphere. The rest of the paper focuses on constellations: we first define and enumerate a family of trees (Section 4), describe a correspondence between these trees and constellations (Section 5) and prove this correspondence is one-to-one (Section 6). We thus obtain the number of constellations, and hence, of minimal transitive factorizations.

2 Constellations and their relatives

A *planar map* is a 2-cell decomposition of the oriented sphere into vertices (0-cells), edges (1-cells), and faces (2-cells). Loops and multiple edges are allowed. The degree of a vertex (or a face) is the number of incidences of edges to this vertex (or face). Two maps are *isomorphic* if there exists an orientation preserving homeomorphism of the sphere that maps cells of one of the maps onto cells of the same type of the other map and preserves incidences. We shall consider maps up to isomorphism.

Definition 2.1 *Let $m \geqslant 2$. An m-constellation is a planar map whose faces are coloured black and white in such a way that*

- *all faces adjacent to a given white face are black, and vice-versa,*
- *the degree of any black face is m,*

A constellation is rooted *if one of its edges, called the root edge, is distinguished.*

The black faces of a constellation will often be called its *polygons* or its *m-gons*. In what follows, we will mainly consider *rooted* constellations, and the word "rooted" will often be omitted. The polygon containing the root edge will be called the *root polygon*. Observe that it is possible to label the vertices of an m-constellation with $1, 2, \ldots, m$ in such a way the vertices of any m-gon are labelled $1, 2, \ldots, m$ in counterclockwise order. We adopt the convention that the ends of the root edge are labelled 1 and 2: this determines the *canonical labelling* of the constellation (Fig. 1).

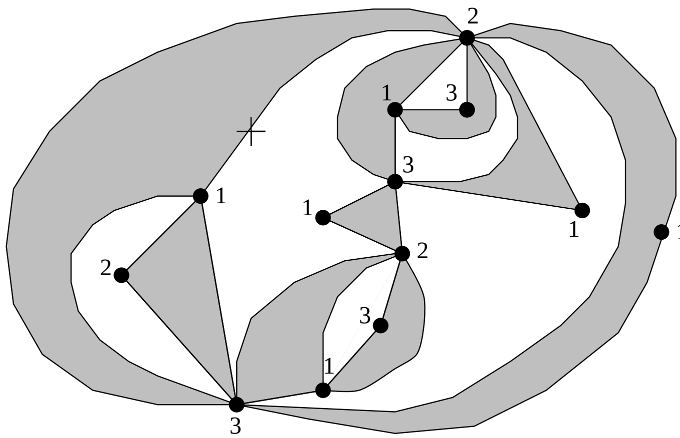

Figure 1: A rooted 3-constellation and its canonical labelling.

One can object that our maps do not look very much like real constellations. The terminology[1], which is due to Alexander Zvonkin, becomes more transparent if we replace each m-gon by an m-star (Fig. 2): we thus obtain a connected set of stars, which is undoubtly a constellation [12].

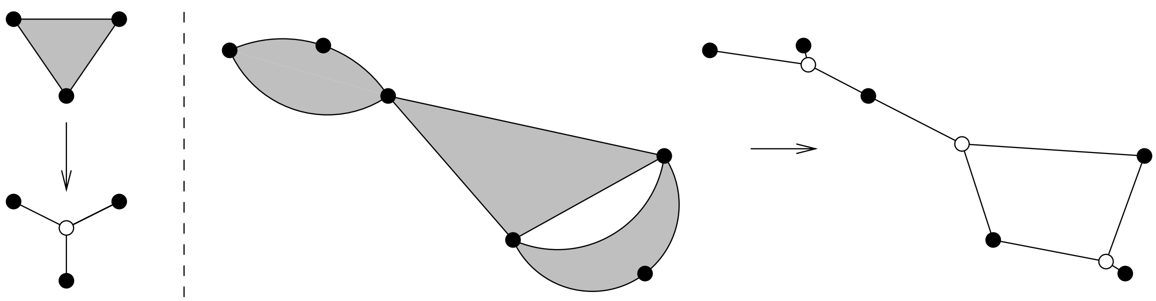

Figure 2: How constellations appear.

Proposition 2.2 *Let $n \geqslant 1$ and $m \geqslant 2$. There exists a one-to-one correspondence between m-tuples $(\sigma_1, \ldots, \sigma_m)$ of permutations of $\mathfrak{S}_n$ such that:*

- *the group generated by $\sigma_1, \ldots, \sigma_m$ acts transitively on $\{1, 2, \ldots, n\}$,*
- *$\sum_{i=0}^{m} c(\sigma_i) = n(m-1)+2$, where $\sigma_0 = \sigma_1 \sigma_2 \cdots \sigma_m$,*

and rooted m-constellations formed of n polygons, labelled from 1 to n in such a way the root polygon has label 1. Moreover, if the constellation has d_i white faces of degree mi, then σ_0 has d_i cycles of length i.

Proof. Let C be a rooted m-constellation formed of n polygons labelled from 1 to n. Recall there is a canonical labelling (by $1, 2, \ldots, m$) of the vertices of C. For $1 \leqslant i \leqslant m$, each m-gon is adjacent to exactly

[1]Note that the word "constellation" was formerly used by Jacques with the meaning of "map" [15].

polygons, denoted σ_i, which we identify with a permutation of S_n.

As the constellation is connected, the group generated by $\sigma_1, \ldots, \sigma_m$ acts transitively on $\{1, 2, \ldots, n\}$.

Moreover, let W be a white face of degree mi: it has exactly i vertices of label m. Let $B_1, B_2, \ldots, B_i$ denote the i black faces adjacent to W by an edge labelled $(1, m)$, arranged in counterclockwise order around W. Then the permutation $\sigma_0 = \sigma_1 \sigma_2 \cdots \sigma_m$ maps B_j onto B_{j+1} for $1 \leqslant j \leqslant i$ (with $B_{i+1} = B_1$). Hence, each cycle of σ_0 corresponds to a white face of C, and the cycle type of σ_0 is given by the degrees of the white faces.

Finally, the number of vertices of C is $v = \sum_{i=1}^{m} c(\sigma_i)$, the number of its faces is $f = n + c(\sigma_0)$ and the number of its edges is $e = nm$. The constellation C is drawn on the sphere, so that Euler's characteristic formula $v + f = e + 2$ reads $\sum_{i=0}^{m} c(\sigma_i) = n(m-1) + 2$.

Conversely, let $(\sigma_1, \ldots, \sigma_m)$ be an m-tuple of permutations as described in the proposition. We consider elementary black m-gons with vertices labelled from 1 to m in counterclockwise order, and white polygons of degree mi for $i \geqslant 1$, the vertices of which are labelled $1, 2, \ldots, m, 1, 2, \ldots, m$, *etc.* in clockwise order. We take n black m-gons, labelled from 1 to n, and $c(\sigma_0)$ white polygons, d_i of which have degree mi. The m-tuple $(\sigma_1, \ldots, \sigma_m)$ describes an incidence relation on these $n + c(\sigma_0)$ polygons. Following this relation, we glue polygons together by identifying edges. According to general topology theory [19, chap.1], this yields a unique 2-cell decomposition of a compact connected surface without boundary. The condition $\sum_{i=0}^{m} c(\sigma_i) = (m-1)n + 2$ ensures, via Euler's characteristic formula, that this surface is the sphere, and hence that the map we have obtained is a planar constellation.

$\blacksquare$

Example. For the labelled rooted 3-constellation C of Fig. 3, we find $\sigma_1 = (1)(2,3)$, $\sigma_2 = (1,2,3)$ and $\sigma_3 = (1,3)(2)$. We compute $\sigma_0 = \sigma_1\sigma_2\sigma_3 = (1)(2)(3)$ which fits with the fact that C has three white faces, each of degree 3. $\square$

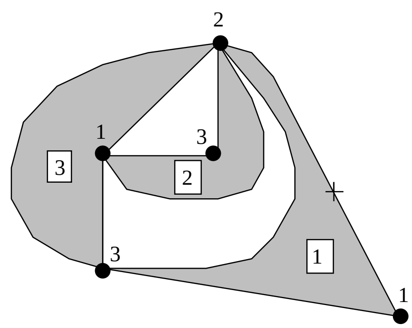

Figure 3: A 3-constellation with labelled 3-gons.

As the m-gons of a rooted constellation formed of n polygons can be labelled in $(n-1)!$ different ways and $n!/\prod_{i\geqslant 1}\left[i^{d_i} d_i!\right]$ permutations have exactly d_i cycles of length i, Proposition 2.2 implies the equivalence between Theorem 1.2 and Theorem 2.3 below, on which we shall focus from now on.

Theorem 2.3 *Let $m \geqslant 2$. The number of rooted m-constellations C having d_i white faces of degree mi, for $i \geqslant 1$, is*

$$m(m-1)^{f-1} \frac{[(m-1)n]!}{[(m-1)n - f + 2]!} \prod_{i\geqslant 1} \frac{1}{d_i!} \binom{mi-1}{i-1}^{d_i},$$

where $n = \sum i d_i$ is the number of m-gons, and $f = \sum d_i$ the number of white faces of C.

We can derive right now two interesting corollaries.

Corollary 2.4 *Let $n \geqslant 1$ and $m \geqslant 2$. The number of rooted m-constellations formed of n polygons is*

$$C_m(n) = \frac{(m+1)m^{n-1}}{[(m-1)n+2][(m-1)n+1]} \binom{mn}{n}.$$

Proof. There is a simple one-to-one correspondence, which preserves the number of polygons, between m-constellations and $(m+1)$-constellations whose white faces have degree $m+1$. Our result will thus follow from Theorem 2.3, by replacing m by $m+1$ and setting $d_1 = n$, $d_i = 0$ for $i \geqslant 2$.

the center of each white face a new vertex labelled $m+1$, and pull the center of each edge $(m,1)$ of the face so that it coincides with the new vertex (Fig. 4). We obtain an $(m+1)$-constellation whose white faces have degree $m+1$ (as each of them contains exactly one vertex labelled $m+1$), and the construction is clearly reversible.

■

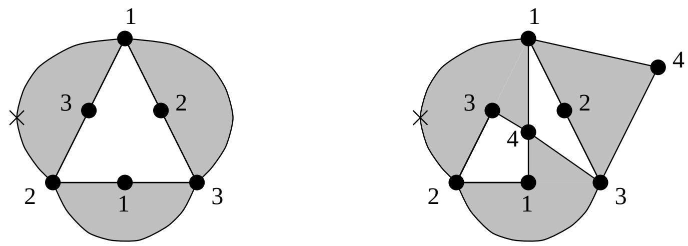

Figure 4: From a 3-constellation to a 4-constellation with all faces of degree 4.

Remark. Using Proposition 2.2, we can translate Corollary 2.4 in terms of permutations: the number of m-tuples $(\sigma_1, \ldots, \sigma_m)$ of permutations of $\mathfrak{S}_n$ such that:

- the group generated by $\sigma_1, \ldots, \sigma_m$ acts transitively on $\{1, 2, \ldots, n\}$,
- $\sum_{i=0}^{m} c(\sigma_i) = n(m-1)+1$, where $\sigma_0 = \sigma_1 \sigma_2 \cdots \sigma_m$,

is

$$\frac{(m+1)m^n(mn-1)!}{((m-1)n+2)!}.$$

Constellations having a unique white face have already appeared in the literature, and are known as *cacti* [1, 7] (see Fig. 11). The case $d_n = 1$, $d_i = 0$ for $i \neq n$ of Theorem 2.3 gives the number of m-cacti having n polygons.

Corollary 2.5 *For $n \geqslant 1$ and $m \geqslant 2$, the number of m-cacti formed of n polygons is*

$$\frac{1}{(m-1)n+1}\binom{mn}{n}.$$

Remarks

1. Cacti correspond to ordered factorizations of an n-cycle of $\mathfrak{S}_n$.

2. Goulden and Jackson gave in [7] a nice closed form expression for the number of ordered factorizations $(\sigma_1, \ldots, \sigma_m)$ of an n-cycle such that each factor σ_i belongs to a prescribed conjugacy class of $\mathfrak{S}_n$. In terms of cacti, this amounts to fixing the degrees of the vertices of each colour. Corollary 2.5 can be derived from their result via a summation, or by recycling their functional equations, only taking into account the number of polygons.

3. We observe that $\frac{1}{(m-1)n+1}\binom{mn}{n}$ is the number of m-ary trees having n vertices. A simple bijection between m-cacti and m-ary trees has been found by Michel Bousquet [2]. The method we use in this paper to prove Theorem 2.3 provides a different bijection between these objects (see Section 5).

Dual maps[2] of constellations will be called *m-Eulerian maps*. The definition of constellations provides the following characterization for m-Eulerian maps (Fig. 5).

Definition 2.6 *A planar map is* m-Eulerian *if it is bipartite (with black and white vertices), and*

- *the degree of any black vertex is m,*
- *the degree of any white vertex is a multiple of m.*

The case $m = 2$ justifies our terminology: if we remove all black vertices from a 2-Eulerian map, we obtain a map having only vertices of even degree; such maps are usually called *Eulerian*.

Of course, counting m-Eulerian maps is equivalent to counting m-constellations. In particular, Theorem 2.3 gives the number of rooted m-Eulerian maps having d_i white vertices of degree mi, for $i \geqslant 1$. When $m = 2$, we recover an old result of Tutte [3, 22, 26].

[2]Recall that the dual map C^* of a map C describes the incidence relation between the faces of C: in particular, the vertices (resp. faces) of C^* are the faces (resp. vertices) of C.

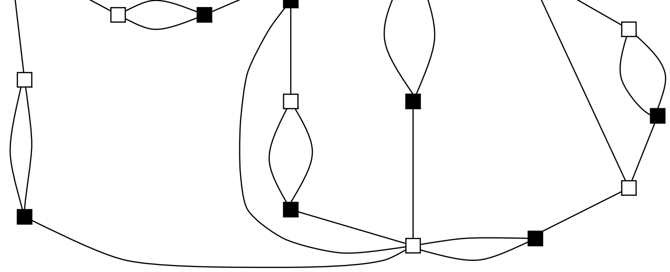

Figure 5: A rooted 3-Eulerian map, dual of the 3-constellation of Fig. 1.

3 Branched coverings and unrooted constellations

This section is independent from the rest of the paper. It contains a brief description of the connection between constellations (or factorizations of permutations) and branched coverings of the sphere, first described by Hurwitz [14], and presented in full generality in [5]. It also addresses the enumeration of *unrooted* constellations having their vertices labelled $1, \ldots, m$ in counterclockwise order around each black face.

Let f be a continuous map from $\mathbb{S}_2$ to $\mathbb{S}_2$, where $\mathbb{S}_2$ is the two-dimensional sphere. Assume every value w on the target sphere has a neighborhood U such that $f^{-1}(U)$ is a finite union of disjoint open sets on which f is topologically equivalent to the complex map z^i around $z=0$, for some positive i (the value 0 corresponding to w). In particular, w has finitely many pre-images. Then f is called a *branched cover* (of the sphere by itself). We say that w has *branching type* $(d_i)_{i\geqslant 1}$ if exactly d_i of its pre-images have a neighborhood of the form z^i for $i \geqslant 1$. It is *regular* if $d_i = 0$ for all $i \geqslant 2$, and *critical* otherwise. It is *simple* if $d_2 = 1$ and $d_i = 0$ for $i \geqslant 3$. The sum $\sum_i id_i$ is constant over all values and called the *degree* of f.

Let $w_0, w_1, \ldots, w_m$ be $m+1$ distinct points in the target sphere. Let $\mathcal{P}$ be a polygon whose vertices are, in counterclockwise order, $w_1, \ldots, w_m$, and that does not contain w_0. If a branched cover f has its critical values in $\{w_0, w_1, \ldots, w_m\}$, then $f^{-1}(\mathcal{P})$ is an unrooted m-constellation with labelled vertices, the elements of $f^{-1}(w_i)$ being labelled i. Conversely, given an m-constellation C with labelled vertices, there exists, up to homeomorphisms of the domain sphere, a unique branched cover f having its critical values in $\{w_0, w_1, \ldots, w_m\}$ such that $f^{-1}(\mathcal{P})$ is isomorphic to C. This one-to-one correspondence between unrooted constellations and branched covers is essentially due to Hurwitz [14], who expressed it in terms of factorizations of permutations. From now on, we shall consider branched covers up to homeomorphisms of the domain sphere *i.e.*, labelled unrooted constellations.

Using this correspondence, Hurwitz translated Theorem 1.1 into the following result: given $m+1$ distinct points $w_0, w_1, \ldots, w_m$ of $\mathbb{S}_2$, the number of branched covers having w_0 as a critical value of branching type $(d_1, d_2, \ldots)$ and $w_1, \ldots, w_m$ as simple critical values is

$$\frac{H_{\sigma_0}}{\prod_{i\geqslant 1} i^{d_i} d_i!},$$

where σ_0 is any permutation with d_i cycles of length i. The proof he gave implicitly relies on the fact that the corresponding unrooted constellations have no non-trivial automorphism, so that their number is directly related to the number of *rooted* constellations. The construction described in the proof of Proposition 2.2 shows that this counting problem is equivalent to the determination of H_{σ_0}.

In general, constellations may have non trivial automorphisms and *unrooting* is more difficult. However, Liskovets' reduction theorem [17, Thm. 3.1] for the unrooting of planar maps can be repeated verbatim to unroot constellations. In this paper, we only apply his argument to m-constellations with n polygons, but it could be done for constellations with given degrees of faces, although the formula becomes more involved.

Theorem 3.1 *The number of unrooted m-constellations with n polygons is*

$$\widetilde{C}_m(n) = \frac{1}{n}\left[C_m(n) + \sum_{\ell \mid n, \ell < n} \phi(n/\ell) \binom{\ell(m-1)+2}{2} C_m(\ell) \right],$$

number of rooted m-constellations with n polygons, given by Theorem 2.1.

Thanks to the one-to-one correspondence described above, we can translate this theorem as follows. Let $w_0, w_1, \ldots, w_m$ be $m+1$ distinct points of $\mathbb{S}_2$. Then $\tilde{C}_m(n)$ is the number of branched covers of degree n, with critical values in the set $\{w_0, w_1, \ldots, w_m\}$. Note that the number of branched covers with prescribed branching types has been given by Mednykh in terms of the characters of the symmetric group [20].

More generally, one can study branched coverings of the sphere by another Riemann surface S. In this context, the covers f are often taken to be holomorphic functions from S to $\mathbb{S}_2$. The case $m = 2$ is the most studied one; it is known as the theory of *dessins d'enfants* (see for instance [16] and references therein).

4 Eulerian trees

In this section, we define and enumerate certain trees that will be proved to be in one-to-one correspondence with constellations.

A *planted tree* is a plane tree with a marked leaf, also called the *root*. In our figures, planted trees hang from their roots. The (total) *degree* of a vertex is the degree in the context of graph theory, *i.e.*, one more than the arity in the functional representation of trees. Vertices of degree 1 are referred to as *leaves*, the others as *inner vertices*. The *inner degree* of a vertex is the number of inner vertices adjacent to it. An *inner edge* connects two inner vertices. The *depth* of a vertex is its distance to the root. The *left-to-right prefix order* (lr-prefix for short) on the vertices of a tree T is obtained recursively by visiting first the root of T, and then its subtrees $T_1, \ldots, T_k$, taken from left to right, in lr-prefix order. The right-to-left prefix (rl-prefix) order is defined symmetrically. The number of planted trees having $n+1$ edges is the famous Catalan number $C_n = \frac{1}{n+1}\binom{2n}{n}$. More generally, the Lagrange inversion formula (see [9] for instance) or encodings by Łukasiewicz words [18, p.221] give the following classical result, first proved by Harary, Prins and Tutte [13].

Theorem 4.1 *The number of planted plane trees having d_i inner vertices of degree $i+1$ for $i \geqslant 1$, is*

$$\frac{(e-1)!}{(\ell-1)!} \prod_{i \geqslant 1} \frac{1}{d_i!},$$

where $e = 1 + \sum i d_i$ is the number of edges and $\ell = 2 + \sum (i-1) d_i$ the number of leaves of such trees.

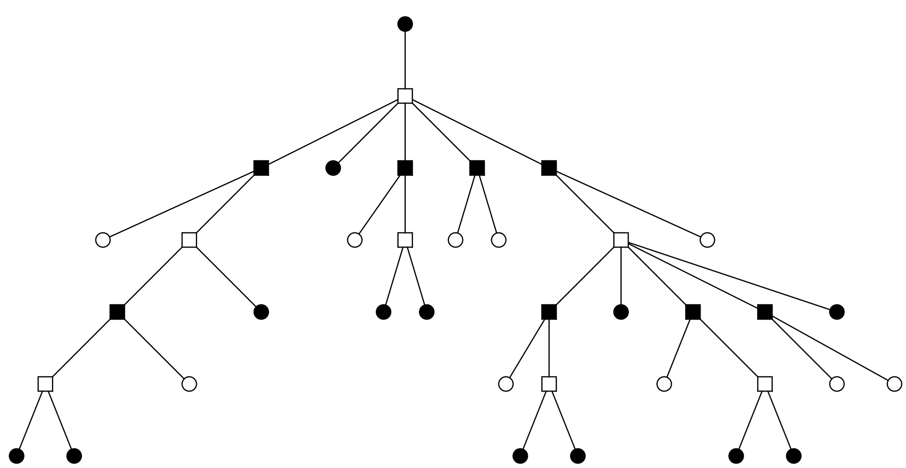

Figure 6: A 3-Eulerian tree (circles represent leaves, squares represent inner vertices).

Definition 4.2 *A bicolored (black and white) tree, planted at a black leaf, is said to be m-*Eulerian *if*

- *all neighbors of a white vertex are black, and vice-versa,*
- *all inner black vertices have total degree m and inner degree 1 or 2,*

of inner degree 1.

Proposition 4.3 *Let $m \geqslant 2$. The number of m-Eulerian trees having d_i white vertices of degree mi, for $i \geqslant 1$, is*

$$(m-1)^{f-1} \frac{[(m-1)n]!}{[(m-1)n-f+1]!} \prod_{i \geqslant 1} \frac{1}{d_i!} \binom{mi-1}{i-1}^{d_i}, \tag{5}$$

where $n = \sum i d_i$ and $f = \sum d_i$. Such trees have exactly:

- *f inner white vertices,*
- *$n-1$ inner black vertices,*
- *$(m-1)n - f - m + 2$ white leaves,*
- *$(m-1)n - f + 2$ black leaves.*

Proof. We can construct all m-Eulerian trees having d_i white vertices of degree mi as follows.
1. We first built the tree T_1 formed of the white inner vertices and black leaves of the final m-Eulerian tree. All the vertices of T_1 have degree 1 modulo $m-1$. More precisely, let d_i be the number of (inner) vertices of degree $(m-1)i+1$, for $i \geqslant 1$. According to Theorem 4.1, the number of such trees is

$$T(d_1, d_2, \ldots) = \frac{[(m-1)n]!}{[(m-1)n-f+1]!} \prod_{i \geqslant 1} \frac{1}{d_i!}.$$

2. In the middle of each inner edge of T_1, add a black vertex of total degree m. This vertex has $m-2$ white leaves, which can be displayed in $m-1$ different ways. As T_1 has $f-1$ inner edges, the number of trees T_2 thus obtained is $(m-1)^{f-1} T(d_1, d_2, \ldots)$.
3. To each of the d_i white vertices of T_2 of degree $(m-1)i+1$, add $i-1$ black children of total degree m. The position of these children can be chosen in $\binom{mi-1}{i-1}$ different ways, and this observation concludes the proof.

■

Let T be an m-Eulerian tree. Let us arrange its leaves cyclically by reading them in lr-prefix order. For instance, starting from the root, we obtain for the tree of Fig. 6 the (cyclic) word $bwbbwbbwbbwwwbbbwbbwwbw$, where w (resp. b) denotes a white (resp. black) leaf. We now match the letters w and b of this word as if they were respectively opening and closing brackets:

b w b b w b b w b b w w w b b b w b b w w b w

More precisely, at step 1, each letter w that is followed by a b is matched with this occurrence of b. We then forget all matched letters and repeat the procedure until no more matches are possible. We match accordingly the leaves of T (Fig. 7). As there are more black leaves than white leaves, some black leaves — exactly m of them — remain unmatched: we call them *single*.

Definition 4.4 *An m-Eulerian tree is said to be* balanced *if its root remains single after the matching procedure.*

Proposition 4.5 *Let $m \geqslant 2$. The number of balanced m-Eulerian trees having d_i white vertices of degree mi for $i \geqslant 1$ is*

$$m(m-1)^{f-1} \frac{[(m-1)n]!}{[(m-1)n-f+2]!} \prod_{i \geqslant 1} \frac{1}{d_i!} \binom{mi-1}{i-1}^{d_i},$$

where $n = \sum i d_i$ and $f = \sum d_i$.

Proof. Let A denote the number given by (5). Then mA can be understood either as the number of m-Eulerian trees having a distinguished single leaf, or, by planting the tree at this leaf and distinguishing the former root, as the number of balanced m-Eulerian trees having a black leaf distinguished. As an m-Eulerian tree having d_i white vertices of degree mi has $(m-1)n - f + 2$ black leaves, the proposition follows.

■

Observe that the expressions given in Theorem 2.3 and Proposition 4.5 are identical: hence Theorem 2.3 will follow from Proposition 4.5 via a one-to-one correspondence between balanced Eulerian trees and constellations.

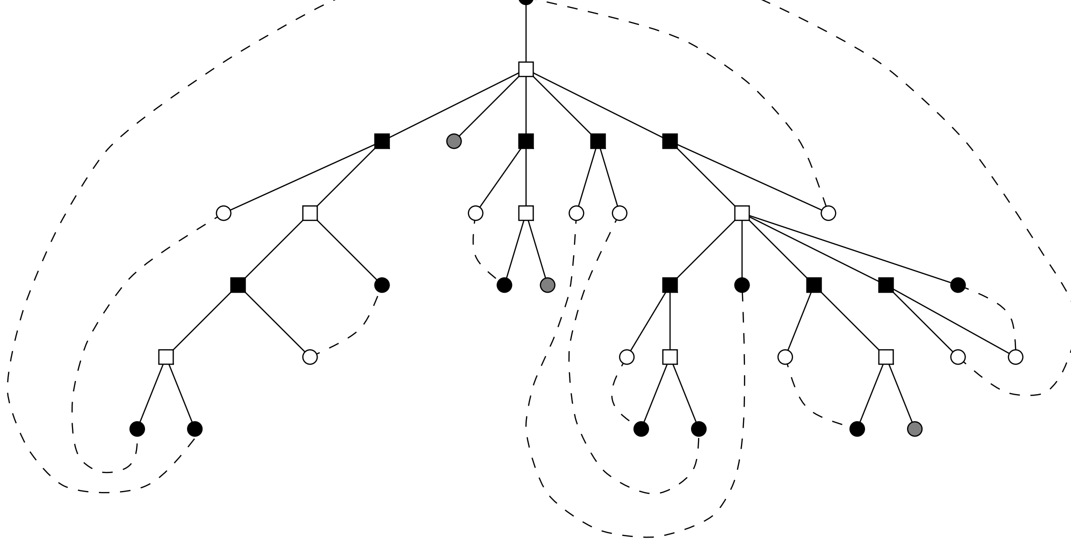

Figure 7: Matching the leaves of the 3-Eulerian tree of Fig. 6. Three leaves remain single.

5 The bijection between balanced Eulerian trees and constellations

5.1 From trees to constellations: the transformation Φ

The transformation of a balanced Eulerian tree T into a constellation $C = \Phi(T)$ is easy to describe. Actually, most of the work has been done already. The construction is exemplified on Fig. 8.

We form a first planar map E_1 by adding edges between the matched leaves of T. We thus obtain the dashed edges of Fig. 8a. Exactly m black leaves remain single. By construction, all of them lie in the same face of E_1; in what follows, we shall often consider E_1 as map on the plane (rather than on the sphere) by taking the convention that the single leaves lie in the infinite face.

We add in the infinite face of E_1 an extra star, having a black center and m rays. Each ray ends with a white leaf. We match these m white leaves with the m single vertices of the tree (dotted lines in Fig. 8a) in cyclic order to obtain a planar map. We mark the dotted edge that ends at the root of the tree. We finally erase all leaves of the underlying tree T and replace dashed and dotted lines by plain lines. By construction, the map we have obtained is a rooted m-Eulerian map E. Taking the dual of E gives a constellation C which we define to be $\Phi(T)$.

Observe that the Eulerian map associated with the tree of Fig. 8 is the map of Fig. 5 and that its dual is the constellation of Fig. 1.

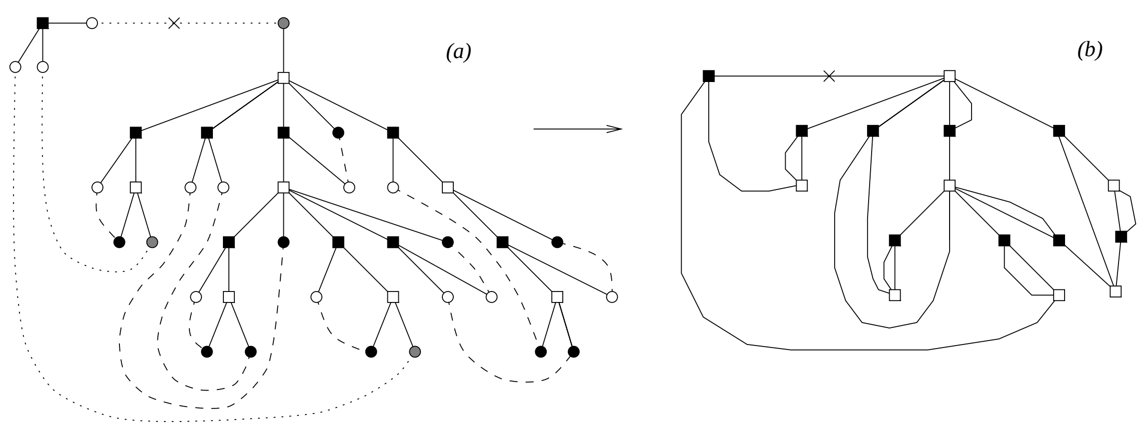

Figure 8: From a balanced 3-Eulerian tree to a 3-Eulerian map.

m-constellations. What can the reverse bijection be? Imagine we start with a rooted m-constellation C (or its dual map E, which is m-Eulerian) and try to construct the corresponding m-Eulerian tree T. What we need to do is select – in a clever way – a set S of edges of E, add two vertices on each of them in such a way the resulting map remains bicolored, and then delete the part of the edge that links these two vertices; this must yield two connected components: an m-star and a balanced m-Eulerian tree. Thus, the central difficulty of the reverse bijection consists in describing the set S of edges of E we need to *open*.

Let us consider again the Eulerian map of Fig. 8b. Looking at Fig. 8a tells us what the set S has to be. Let us draw the set of dual edges, denoted S' (Fig. 9, thick lines). We observe that S' is formed of the root m-gon of the constellation C, on which m trees, denoted $T_1, \ldots, T_m$ are planted. These m trees cover all vertices of C. We shall see that this is a general phenomenon: describing the reverse bijection of Φ boils down to defining a certain covering forest of a constellation, which will be called its *rank forest*.

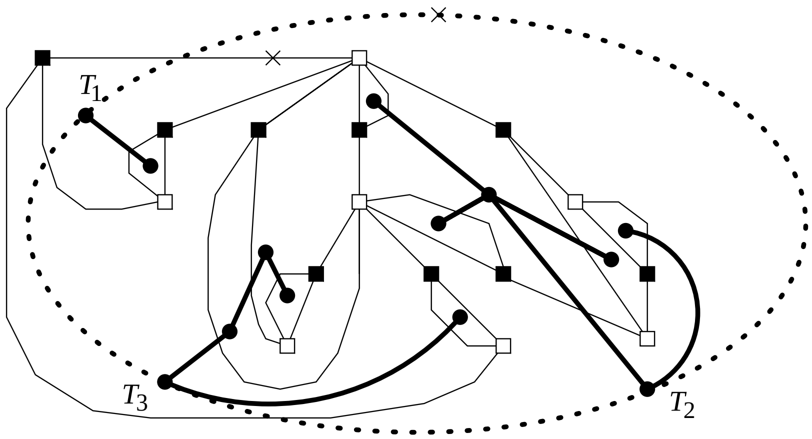

Figure 9: The dual edges of the dashed and dotted edges of Fig. 8a.

5.2 From constellations to trees: the transformation Ψ

Let C be a rooted constellation; let us draw it on the plane in such a way the infinite face is the root m-gon. Let $\bar{C}$ be obtained by orienting the edges of C in counterclockwise direction around m-gons. We define the *rank* $r(v)$ of a vertex v as the length of the shortest (oriented) path of $\bar{C}$ going from a vertex of the root m-gon to v (Fig. 10). The rank of v should not be mixed up with its label $\ell(v) \in \{1, 2, \ldots, m\}$, given by the canonical labelling defined in Section 2. The following lemma tells us how to construct the *rank forest* of a constellation. The principle is simple: we start from the root m-gon and proceed by breadth first search, from right to left. The reader is advised to practice on the example of Fig. 10.

Lemma 5.1 *Let C be a rooted constellation. There exists a unique covering forest F of C, consisting of m trees T_a, $1 \leqslant a \leqslant m$, respectively planted at the vertex labelled a of the root polygon, that satisfies the following four properties.*

1. *The orientation of edges of F induced by the trees T_a (from the roots to the leaves) coincides with their orientation in the oriented map $\bar{C}$.*

2. *The rank increases by one along each edge of F. In other words, the depth of a vertex of T_a is given by its rank.*

Let u be a vertex of C. Properties 1 and 2 imply that u belongs to T_a, where $a = \ell(u) - r(u) \bmod m$.

3. *Assume $r(u) > 0$. All vertices of label $\ell(u) - 1$ and rank $r(u) - 1$ belong to T_a. If we visit them in rl-prefix order, the first one that is adjacent to u is the father of u in T_a.*

4. *Let v be the father of u in T_a. Let e be the edge of T_a that links v to its father. If we visit the edges of C adjacent to v in clockwise order, starting from e, the first one that ends at u belongs to T_a.*

This covering forest will be called the rank forest *of C.*

$1 \leqslant a \leqslant m$, the tree T_a is reduced to the vertex labelled a of the root polygon. We plant T_a by attaching to this vertex a short extra edge that lies in the infinite face of C.

Assume that, after step k, the forest we have obtained is not yet covering C. Let u be a vertex of rank $k+1$. All vertices of rank k and label $\ell(u)-1$ belong to the same tree T_a. We choose the father v of u according to Property 3 of our lemma, and the edge of T_a joining v to u according to Property 4. ■

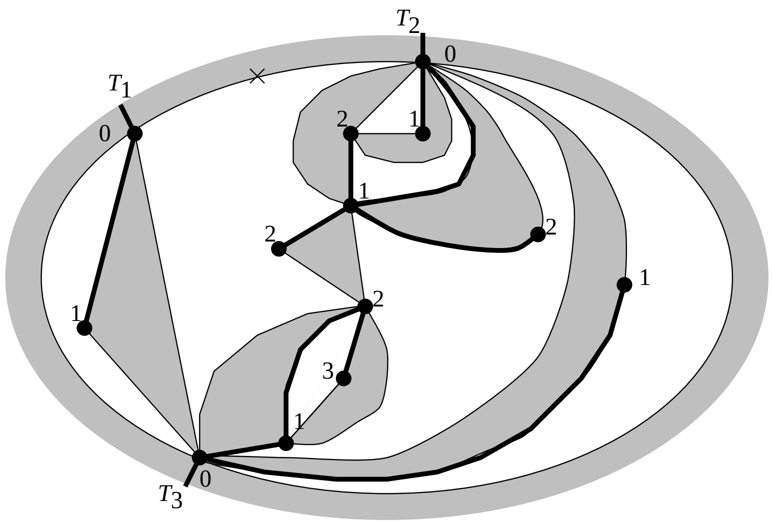

Figure 10: A rooted 3-constellation: the ranks of the vertices and the rank forest.

Once the rank forest of C is constructed, the Eulerian tree $\Psi(C)$ is easy to obtain. Let S' be the set of edges of C that belong either to the rank forest, or to the root m-gon. Let E be the dual map of C, and S be the dual set of S'. On each edge e of S, we add two vertices in such a way the resulting map remains bicolored; we then delete the part of e that links these two vertices. We claim that this provides an m-star and a balanced m-Eulerian tree $\Psi(C)$, which we plant by the root edge of E.

Example. Starting from the constellation of Fig. 10, we obtain the tree of Fig. 8a. □

The next section is devoted to the proof of the following theorem, which, according to Proposition 4.5, implies Theorem 2.3.

Theorem 5.2 *The transformation Φ is a bijection from balanced m-Eulerian trees to m-constellations. The reverse bijection is Ψ. Moreover, if $\Phi(T)=C$ and T has d_i white vertices of degree mi, then C has d_i white faces of degree mi.*

Note that the last statement is clear. It is, unfortunately, the unique obvious statement of this theorem.

Our bijection illustrates a general idea that is developed in [23]: *rooted planar maps are canonical representants of conjugacy classes of planted plane trees.* Here, we say that two trees are *conjugated* if one is obtained from the other by changing the root. This implies that conjugacy classes are simply plane trees, but the terminology originates in the analogy with words. Indeed, conjugating a tree results in conjugating the word obtained by a prefix ordering of its leaves, so that Proposition 4.5 can be seen as an application of the cycle lemma. The motto of [23] could then be stated: *if applying the cycle lemma to words yields trees, applying it to planted trees yields maps.* Besides constellations, planar maps, Eulerian planar maps, nonseparable planar maps and cubic nonseparable maps can indeed be obtained from suitable balanced trees by some matching procedure very similar to Φ.

Remarks.
1. Our main theorem (Theorem 2.3) refines Corollary 2.4 by taking into account the degrees of the white faces. Thanks to our bijection, we can give algebraic equations for another enumeration problem that also refines Corollary 2.4, this time by counting white faces and vertices labelled k, for $1 \leqslant k \leqslant m$. For $\mathbf{i} = (i_1, \ldots, i_m, i_{m+1}) \in \mathbb{N}^{m+1}$, let $a_{\mathbf{i}}$ denote the number of m-constellations having exactly i_{m+1} white

$g(\mathbf{x}) = \sum_{\mathbf{i}} a_{\mathbf{i}} \mathbf{x}^{\mathbf{i}}$ to be the corresponding generating function. Then

$$g(\mathbf{x}) = \int u_1(\mathbf{x}) dx_1 \tag{6}$$

where, for $1 \leqslant k \leqslant m+1$,

$$u_k(\mathbf{x}) = \prod_{i \in [1, m+1],\ i \neq k} \left(x_i + \sum_{j \neq i} u_j(\mathbf{x}) \right). \tag{7}$$

This corresponds to counting factorizations of permutations by the number of cycles of the factors.

To obtain this result, we first observe that, according to the proof of Corollary 2.4, the series $g(\mathbf{x})$ counts $(m+1)$-constellations having only faces of degree $m+1$. Via our bijection, they correspond to balanced $(m+1)$-Eulerian trees having all inner vertices of degree $m+1$. Eq. (6) essentially expresses the generating function for *balanced* Eulerian trees in terms of the generating function for *general* Eulerian trees. The latter can be decomposed in a standard way to provide m trees, which are Eulerian: this decomposition yields (7).

2. What happens to cacti? By Theorem 5.2, an m-cactus formed of n polygons is mapped by Ψ on a balanced m-Eulerian tree having a unique white inner vertex, of degree mn (Fig. 11). Let us visit its children from left to right, writing a letter a for each of the $n-1$ black inner vertices and a letter b for each of the $n(m-1)$ black leaves we meet. We obtain a word $x_1 \cdots x_{nm-1}$. As the tree is balanced, the word $ax_1 \cdots x_{nm-1}b$ is the prefix code of a complete m-ary tree with n inner vertices and $(m-1)n+1$ leaves; this defines a one-to-one correspondence between m-cacti formed of n polygons and complete m-ary trees with n inner vertices, which differs from the one found by Michel Bousquet [2].

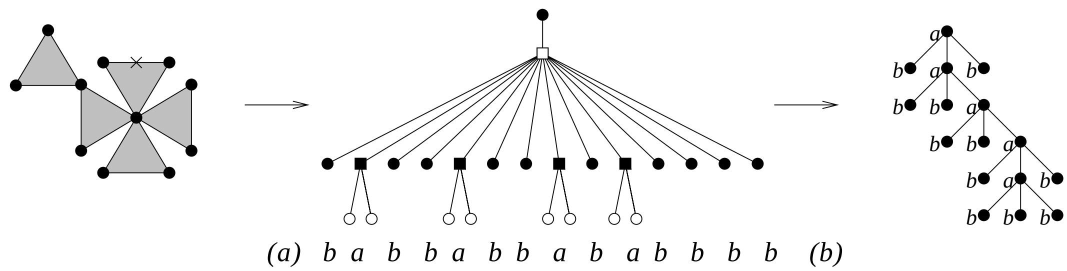

Figure 11: The restriction of our bijection to cacti.

6 Why does it work?

In order to prove Theorem 5.2, we must prove the three following points:

- $\Psi \circ \Phi = id$,
- the transformation Ψ always produces a balanced Eulerian tree,
- $\Phi \circ \Psi = id$.

We prove the first point in Section 6.1, and the next two in Section 6.2. In what follows, we try to use uniform notations for each type of object we meet. In particular, we make an immoderate use of the following conventions.

- **Constellations**: let C be a rooted constellation. We consider C as a map on the plane by choosing the root m-gon as the infinite face. We denote by $\bar{C}$ the oriented version of C (see Section 5.2). A finite face of C will be denoted B or W, depending on whether it is black or white. If this face is of degree mi, its vertices will be called $v_1, v_2, \ldots, v_{mi}$, in *clockwise* order. The rank of v_k will be r_k, and the edge of the face connecting v_k to v_{k+1} will be denoted e_k.
- **The rank forest**: given a constellation C, its rank forest will be denoted F. This forest contains m trees, denoted $T_1, \ldots, T_m$, in such a way that T_a contains the vertex of the infinite face labelled a, for $1 \leqslant a \leqslant m$. The edges of F and of the root m-gon form a set that we denote by S'.

- **Eulerian trees** are of course denoted T.
- Let a and b be two integers; the notation $a \equiv b$ means that a and b are equal modulo m.

Before we go into the details of the proof, we need to give an alternative characterization of the map E_1 obtained by matching the edges of a balanced Eulerian tree T (see Section 5.1). We observe that E_1 is the unique map on the plane that can be obtained by adding edges to T with the following rules:

– each edge e of $E_1 \setminus T$ joins a white leaf w of T to a black leaf b; moreover, the tree T lies *to the left* of e, in the following sense: the map formed by T and e has two faces, and

if we turn around w in counterclockwise order, we meet
successively the infinite face, the edge e, and the finite face. (8)

– each white leaf of T is incident to exactly one edge of $E_1 \setminus T$; each black leaf is incident to at most one edge of $E_1 \setminus T$;

– the m black leaves of T that are not adjacent to an edge of $E_1 \setminus T$ lie in the infinite face.

6.1 $\Psi \circ \Phi = id$

The above statement is equivalent to the following proposition.

Proposition 6.1 *Let T be a balanced Eulerian tree, $C = \Phi(T)$ the corresponding constellation, and E the dual map of C. Let S be the set of edges of E that were not edges of T, and let S' be the dual set of edges. Then the elements of S' are the edges of the root m-gon and of the rank forest of C.*

Proof. As T is a tree, it defines a unique face, and thus the figure formed by the edges of S' is connected. Moreover, the edges of the root m-gon clearly belong to S': they are the duals of the edges connecting T to the extra star (dotted edges in Fig. 8a). The root m-gon is the only cycle of S', because any cycle of S' separates two non empty connected components of the figure formed of T and the extra star. Let G be the smallest subgraph of C containing the edges of S' that are not in the root m-gon and all the vertices of the root m-gon. The above remarks prove that G is a covering forest of C consisting of m trees $T_1, \ldots, T_m$, such that T_a is planted at the vertex labelled a of the root polygon. We want to prove that G is the rank forest of C, that is, satisfies the four properties of Lemma 5.1.

The first property is guaranteed by Property (8) above. Let e' be an edge of T_a, and e its dual. Combining the fact that the tree T lies to the left of e and does not intersect T_a, we conclude that the orientation of e' given by T_a coincides with its orientation in $\bar{C}$ (Fig. 12).

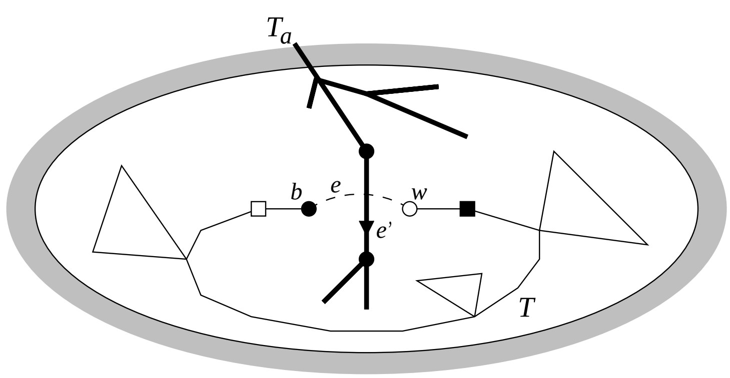

Figure 12: The top-down orientation in T_a coincides with the orientation in $\bar{C}$.

We shall prove the three other properties of Lemma 5.1 by a common counting argument. Let us begin with some generalities. Let e' be an (oriented) edge of $\bar{C}$ going from w to u, and assume that e' does not belong to S'. Let e be the dual edge of e'. Then e is an inner edge of the tree T. Let us add on e two new vertices (Fig. 13). The edge e' splits T into two trees T' and T'', respectively planted at a black leaf and at a white one. The definition of an Eulerian tree (Definition 4.2) being local, we observe that:

– otherwise, T'' is reduced to a single inner vertex, and T' has $2m$ more black leaves than white leaves. In both cases, the difference between the number of black and white leaves of T' is αm, with $\alpha = 1$ or 2.

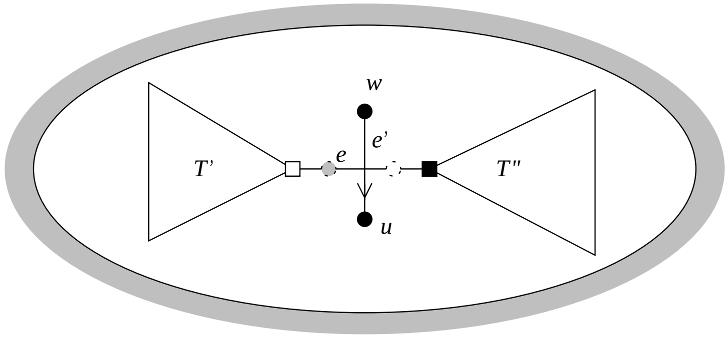

Figure 13: An edge of C not belonging to S' splits T into two trees.

Let us now prove by induction on the rank k that all vertices of C satisfy Properties $2, 3$, and 4 of Lemma 5.1. Clearly, there is nothing to prove for $k = 0$. Assume the properties are true for all vertices of rank smaller than $k > 0$. Let u be a vertex of rank k. Let T_a be the tree of G that contains u. According to the first property of Lemma 5.1 (which we have proved), the depth of u in T_a is at least k, say $k + d$ with $d \geqslant 0$.

Let us first prove *ab absurdo* that $a \equiv \ell(u) - r(u)$. Assume $\ell(u) - r(u) \equiv b \not\equiv a$. By definition of the rank, there exists an oriented edge e' of $\bar{C}$ that goes from a vertex w of rank $k - 1$ and label $\ell(u) - 1$ to u. By assumption, w lies at depth $k - 1$ in the tree T_b (as $b \equiv \ell(w) - r(w) \equiv \ell(u) - r(u)$), and hence the edge e' cannot belong to S' (Fig. 14). We define the trees T' and T'' as above. Property (8) implies that the edges of T_b that lie on the path that links the root of T_b to w are the duals of the (dashed) edges that match black leaves of T' to white leaves of T''. Hence, the number of black leaves of T' matched with a white leaf of T'' is $k - 1$. Similarly, the number of white leaves of T' matched with a black leaf of T'' is $k + d$. Let r denote the number of single black leaves of T that lie in T'. Of course, $r \leqslant m$. The difference between the number of black and white leaves of T' is $(k - 1) + 1 + r - (k + d) = \alpha m$, which gives $r = d + \alpha m \geqslant m$. This implies that $r = m$: all single leaves of T lie in T'. But this is impossible: as $a \neq b$, at least one black leaf of T'' is single. Hence $\ell(u) - r(u) \equiv a$.

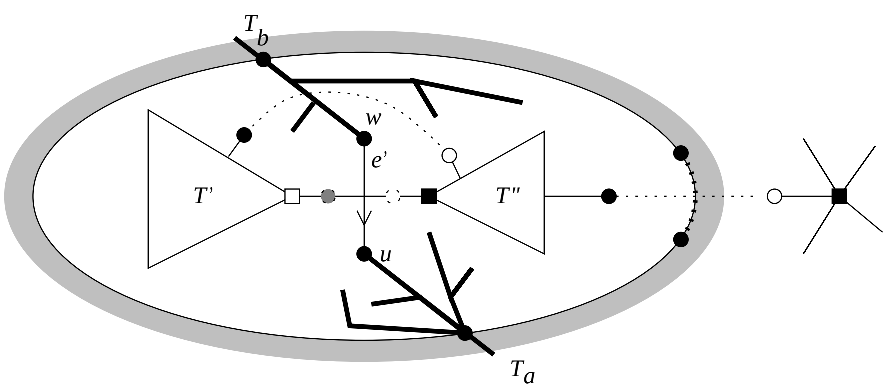

Figure 14: Assuming $a \neq b$.

Let us now prove that u lies at depth k in T_a, *i.e.*, that $d = 0$. Let e' be any oriented edge of $\bar{C}$ going from a vertex w of rank $k - 1$ and label $\ell(u) - 1$ to u. By assumption, w occurs at depth $k - 1$ in T_a. If e' belongs to T_a, then u occurs at depth k in T_a. Otherwise, let x be the common ancestor of u and w in T_a having the largest depth, say δ (Fig. 15). Again, counting black leaves and white leaves of T' gives $(k - 1 - \delta) + r + 1 - (k + d - \delta) = \alpha m$, hence $r = d + \alpha m$. As $r \leqslant m$, this implies that $d = 0$.

The identity $r = d + \alpha m$ also gives $\alpha = 1$ and $r = m$. This implies that T' lies in the infinite face of the map formed by T_a and the edge e'. In particular, if the father v of u in T_a is not w, then v comes

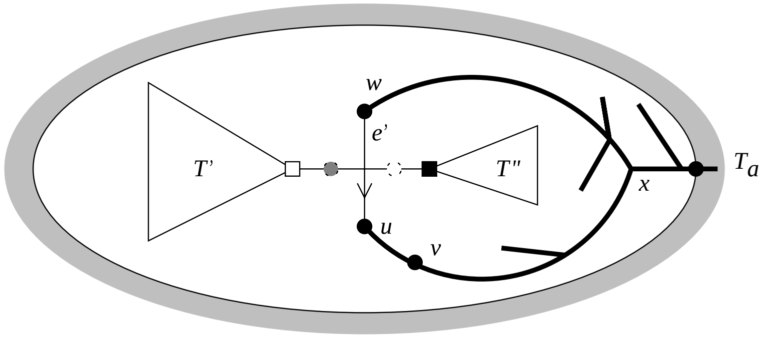

Figure 15: The forest G satisfies the third property of Lemma 5.1.

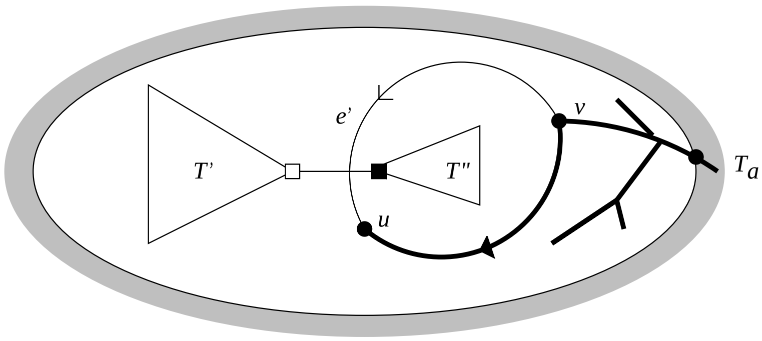

Figure 16: The forest G satisfies the fourth property of Lemma 5.1.

6.2 The construction Ψ provides a balanced Eulerian tree, and $\Phi \circ \Psi = id$

Proposition 6.2 *Let C be an m-constellation. Applying the algorithm Ψ described in Section 5.2 provides a balanced m-Eulerian tree T such that $\Phi(T) = C$.*

Lemma 6.3 *The graph $T = \Psi(C)$ is a planted tree.*

Proof. The subgraph of C formed by the edges of S' (*i.e.*, the rank forest F and the edges of the root m-gon) is connected. Hence, given two vertices of C, there exists a path formed of edges of S' that connects them. In terms of the Eulerian map E, this means that we can connect any two faces of E by a path that does not intersect the graph U obtained from E by opening the edges of S. This implies that U has a single face, and hence, is a forest.

The number of connected components of this forest is $v(U) - e(U)$, where $v(U)$ and $e(U)$ denote respectively the number of vertices and edges of U. In order to determine these numbers, we need to know how many edges of C have been opened, that is, the cardinality of S'. As F is a covering forest of C composed of m trees, it contains $v(C) - m$ edges. Hence, the set S' has cardinality $v(C)$. This gives $v(U) = v(E) + 2v(C)$ and $e(U) = e(E) + v(C)$. Hence,

$$\begin{aligned} v(U) - e(U) &= v(E) + v(C) - e(E) \\ &= v(E) + f(E) - e(E) \\ &= 2 \end{aligned}$$

where $f(E)$ stands for the number of faces of E, and the last equality is a consequence of Euler's characteristic formula.

Thus the forest U contains exactly two trees: one of them is the m-star, and the other is $T = \Psi(C)$.

white vertex is a multiple of m. We now have to prove that the conditions stated in Definition 4.2 are fulfilled. This will be done in Lemmas 6.5 and 6.6.

Property 6.4, which is geometrically clear (see Fig. 17), will play an important role in the proof. It means that, when we visit the vertices of a face of C in clockwise order, we meet successively vertices of the tree T_1, then vertices of T_2, and so on until we meet vertices of T_m.

Property 6.4 *Let f be a (white or black) finite face of C of degree mi. We can denote its vertices by $v_1, \ldots, v_{mi}$, in clockwise order, in such a way that there exists a sequence $1 = k_1 \leqslant k_2 \leqslant \cdots \leqslant k_m \leqslant k_{m+1} = mi + 1$ such that for $1 \leqslant j \leqslant m$,*

$$\{k : v_k \in T_j\} = \{k_j, k_j + 1, \ldots, k_{j+1} - 1\}.$$

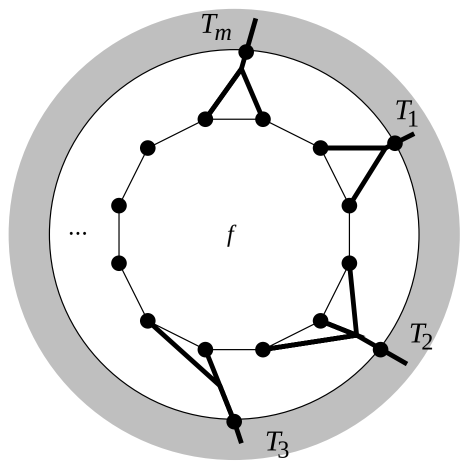

Figure 17: The trees $T_1, \ldots, T_m$ appear in clockwise order around any face of C.

We first establish the properties of the black vertices of the tree $T = \Psi(C)$.

Lemma 6.5 *Let v be an inner black vertex of T. Then v has inner degree 1 or 2.*

Proof. Let us translate the above statement in terms of the constellation C: let B be an m-gon of C, different from the root m-gon. Then at most 2 of its edges do not belong to the covering forest F.

Observe that, if u and v are any two adjacent vertices of B, taken in clockwise order, then by definition of the rank $r(u) \leqslant r(v) + 1$. We need to study separately two cases, depending on whether the equality holds for an edge of B not belonging to F.

First case: there exists in B an edge e, not belonging to F, whose ends u and v, taken in clockwise order, are such that $r(u) = r(v) + 1$. We call e a special edge.

As $\ell(u) \equiv \ell(v) + 1$, the vertices u and v belong to the same tree of the forest F, say T_a. Let w be the father of u in T_a. Then w has rank $r(v)$ and label $\ell(v)$. According to the properties satisfied by T_a (Lemma 5.1),

- either $w \neq v$ and w comes first in the right-to-left prefix order of T_a (Fig. 18*a*),
- or $w = v$ and the edge e' that connects v and u in T_a comes before e when we visit the edges of C that are adjacent to v in clockwise order, starting from the edge that connects v to its father in T_a (Fig. 18*b*).

Fig. 18 and Property (8) imply that all vertices of B belong to T_a, and show that it is topologically impossible to have a second special edge in B. All vertices x of B satisfy $\ell(x) - r(x) \equiv a$. In particular, any two different vertices have different ranks. Let us denote the vertices of B by $v_1, \ldots, v_m$, in clockwise order, in such a way that $r_1 < r_k$ for $1 < k \leqslant m$. As $r_k \leqslant r_{k+1} + 1$ for $1 \leqslant k \leqslant m$, we conclude, by induction on i, that $r_{m-i} = r_1 + i + 1$ for $0 \leqslant i < m - 1$.

Let e_i be the edge of B connecting v_i and v_{i+1}, for $1 \leqslant i \leqslant m$. Assume the special edge e is e_j. As e_j is the *unique* special edge, the fact that $r_k = r_{k+1} + 1$ for $2 \leqslant k \leqslant m$ implies that the edges $e_2, \ldots, e_{j-1}$

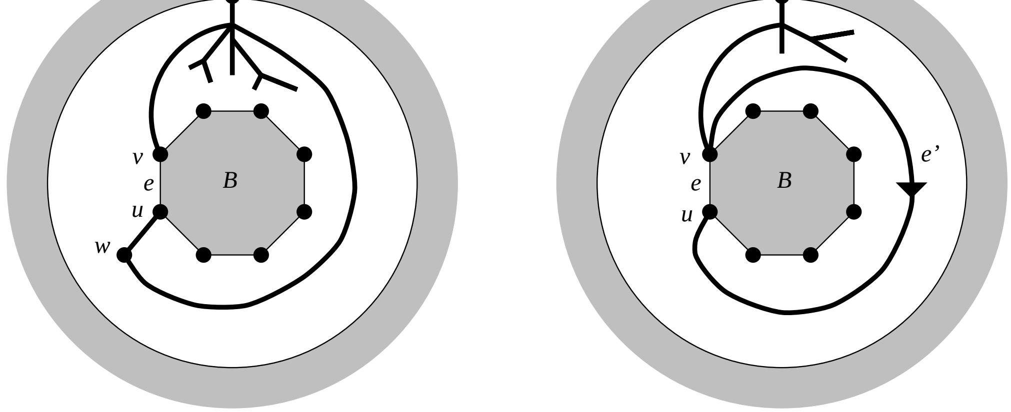

Figure 18: How a special edge looks like.

and $e_{j+1}, \ldots, e_m$ belong to F. Therefore, the m-gon B has at most two edges (e_1 and e_j) that do not belong to F.

Second case: any edge e of B not belonging to F is such that $r(u) < r(v) + 1$, where u and v are the ends of e, taken in clockwise order.

In this case, we are going to prove our lemma *ab absurdo*. Assume at least three edges of B do not belong to F. Let us denote $v_1, \ldots, v_m$ the vertices of B (in clockwise order), and let e_i be the edge of B connecting v_i and v_{i+1}, for $1 \leqslant i \leqslant m$. Assume the three edges e_1, e_i and e_j do not belong to F, with $1 < i < j \leqslant m$.

For $1 \leqslant k \leqslant m$, we have $r_k \leqslant r_{k+1} + 1$. Moreover, the inequality is strict, by assumption, when $k \in \{1, i, j\}$. Thus:

$$r_1 < r_2 + 1 \leqslant r_3 + 2 \leqslant \cdots \leqslant r_i + i - 1 < r_{i+1} + i \leqslant \cdots \leqslant r_j + j - 1 < r_{j+1} + j \leqslant \cdots \leqslant r_1 + m.$$

Let $p = \ell(v_1)$ be the label of v_1. Then for $1 \leqslant k \leqslant m$, we have $\ell(v_k) \equiv p - k + 1$. Hence,

$$\begin{aligned} p - r_1 > p - 1 - r_2 \geqslant \cdots \geqslant p - i + 1 - r_i > p - i - r_{i+1} \geqslant \cdots \\ \cdots \geqslant p - j + 1 - r_j > p - j - r_{j+1} \geqslant \cdots \geqslant p - m - r_1. \end{aligned} \tag{9}$$

Recall that a vertex v belongs to the tree T_a if and only if $a \equiv \ell(v) - r(v)$. Let T_a (resp. T_b, T_c) be the tree of F containing v_1 (resp. v_i, v_j), and assume v_1 has been chosen in such a way that $a \geqslant b$ and $a \geqslant c$. There exist three integers k, k', k'' such that $p - r_1 = a + km$, $p - i + 1 - r_i = b + k'm$ and $p - j + 1 - r_j = c + k''m$. According to (9), we have $a + km > b + k'm > c + k''m > a + (k-1)m$. This imposes $k = k' = k''$ and $a > b > c$, which contradicts Property 6.4.

■

We now establish the properties of the white vertices of the tree $\Psi(C)$.

Lemma 6.6 *Let v be an inner white vertex of $T = \Psi(C)$, of degree mi. Exactly $i-1$ of its inner neighbours have inner degree* 1.

Proof. Let us translate the above statement in terms of the constellation C. Let W be a white face of C of degree mi. Every edge e of W belongs to a unique m-gon denoted $B(e)$. Let κ be the number of edges e of W satisfying the following condition:

$$e \text{ is the only edge of } B(e) \text{ that does not belong to the forest } F. \tag{10}$$

We wish to prove that $\kappa = i - 1$.

1. We first prove that κ is at most i. Let $v_1, \ldots, v_{mi}$ denote the vertices of W, taken in clockwise order. For $1 \leqslant k \leqslant mi$, let e_k denote the edge of W connecting v_k and v_{k+1}, and let r_k be the rank of v_k. Then

can write $r_k \geqslant r_{k+1} + a_k$, where $a_k = m - 1$ if e_k satisfies (10) and $a_k = -1$ otherwise. We have

$$r_1 \geqslant r_2 + a_1 \geqslant r_3 + a_1 + a_2 \geqslant \cdots \geqslant r_{mi} + a_1 + \cdots + a_{mi-1} \geqslant r_1 + a_1 + \cdots + a_{mi}, \tag{11}$$

which implies that $a_1 + \cdots + a_{mi} \leqslant 0$. As κ edges of W satisfy (10), we have $a_1 + \cdots + a_{mi} = \kappa(m-1) - (mi - \kappa) = m(\kappa - i)$, and hence $\kappa \leqslant i$.

2. Let us now rule out the hypothesis $\kappa = i$. If $\kappa = i$, then all the inequalities in (11) become equalities. In other words, for every edge e_k of W that does not satisfy (10), $r_{k+1} = r_k + 1$. As $\ell_{k+1} \equiv \ell_k + 1$, Lemma 5.1 shows that all vertices of W belong to the same tree of F, say T_a.

Assume that v_1 is a vertex of W of minimal rank: $r_1 \leqslant r_k$ for all k. We split the vertices of W in two classes: right and left vertices, as follows (Fig. 19).

- By convention, v_1 is a right vertex; let e_0 be the edge of T_a ending at v_1.
- If v_k comes before v_1 in the right-to-left prefix order of T_a, then v_k is a right vertex.
- If v_k lies in the subtree of T_a of root v_1 and the branch of T_a going from v_1 to v_k begins with an edge e lying between e_0 and e_1 (in clockwise order), then v_k is a right vertex.
- In all other cases, v_k is a left vertex.

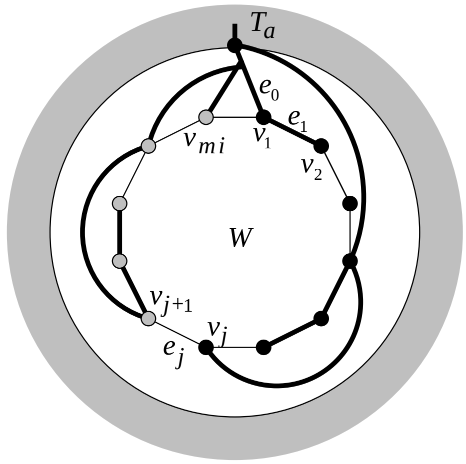

Figure 19: Left and right vertices of a white face.

Observe that, if v_k and v_ℓ are respectively right and left vertices, then v_k comes before v_ℓ in the rl-prefix order of T_a. Moreover, if e_j satisfies (10) and v_j is a right vertex, then $j \neq mi$ and v_{j+1} is a right vertex. Let $j = \max\{k : v_k \text{ is a right vertex}\}$. Then e_j cannot satisfy (10). Consequently, $r_{j+1} = r_j + 1$. In passing, note that this implies that $j \neq mi$, since r_1 is minimal; in other words, v_{j+1} is a left vertex.

The father of v_{j+1} in T_a is a vertex v, of rank r_j. If $v \neq v_j$, then, by construction of T_a, v comes before v_j in the rl-prefix order of T_a, and so does v_{j+1}. This contradicts the fact that v_{j+1} is a left vertex. If $v = v_j$, then $j = 1$, because v_1 is the only right vertex of W that might have left sons. By construction of T_a, the edge of T_a that links v_1 and v_2 lies between e_0 and e_1, which makes v_2 a right vertex, and gives a contradiction. Hence $\kappa \leqslant i - 1$.

3. Let us turn back to the tree $T = \Psi(C)$. Let d_i be the number of white vertices of T of degree mi. For $j \in \{1, 2\}$, let b_j denote the number of inner black vertices of T of inner degree j. We have proved that an inner white vertex of T of degree mi has at most $i - 1$ inner neighbors of inner degree 1. This implies that $b_1 \leqslant \sum_i (i-1) d_i$. We are going to prove that the equality actually holds: $b_1 = \sum_i (i-1) d_i$, and this will conclude the proof of our lemma.

Let ℓ_b (resp. ℓ_w) denote the number of black (resp. white) leaves of T. We know that $\ell_b = \ell_w + m$. Let us count in two different ways the edges of T, using the fact that each of them has a white end and a black end. We find:

$$\ell_w + m \sum_i i d_i = \ell_b + m(b_1 + b_2),$$

and hence $b_1 + b_2 = -1 + \sum_i i d_i$. Now, every inner white vertex of T, but one, is the child of an inner black vertex of inner degree 2; this gives $\sum d_i = 1 + b_2$, and finally $b_1 = -1 + \sum i d_i + 1 - \sum d_i = \sum (i-1) d_i$.
$\blacksquare$

6.2, what remains to prove is the following lemma.

Lemma 6.7 *The m-Eulerian tree $T = \Psi(C)$ is balanced, and $\Phi(T) = C$.*

Proof. Let us alter the procedure that leads from C to $T = \Psi(C)$ as follows: having added a pair of vertices on each edge e of S, we delete the central part of e (which connects the two new vertices) if e is dual to an edge of the root m-gon of C. Otherwise, we just replace the central part of e by a dashed edge. We need to prove that the dashed edges realize the matching of edges of T. This will imply that the root edge of T, which is not incident to a dashed edge, is a single leaf of T, *i.e.*, that T is balanced, and will also prove that $\Phi(T) = C$.

What we want to prove is that the dashed edges satisfy the properties described at the beginning of Section 6. The only property that is not completely clear is (8), stating that T lies to the left of the dashed edges, when we visit them from their white end to their black one.

Let e be a dashed edge: its dual e' belongs to a tree T_a of the rank forest F. Recall that T_a does not intersect T. According to Property 1 of Lemma 5.1, the tree T has to be on the left of e (Fig. 12).

■

Aknowledgements. We are grateful to Alexander Zvonkin who introduced us to constellations via his exciting talks and many interesting discussions. We also thank Volker Strehl for his interest in this paper.

References

[1] M. Bóna, M. Bousquet, G. Labelle and P. Leroux, Enumeration of m-ary cacti, *Adv. in Appl. Math.* **24** No.1 (2000) 22–56.

[2] M. Bousquet, *Espèces de structures et applications au dénombrement de cartes et de cactus planaires*, Ph.D. thesis, Université du Québec à Montréal, 1998. Publications du LaCIM, vol. 24, Montréal, 1999.

[3] R. Cori, *Un code pour les graphes planaires et ses applications*, *Astérisque*, vol. 27, Société Mathématique de France, 1975.

[4] J. Dénes, The representation of a permutation as the product of a minimal number of transpositions and its connection with the theory of graphs, *Publ. Math. Inst. Hungar. Acad. Sci.* **4** (1959) 63–70.

[5] C. L. Ezell, Branch point structure of covering maps onto nonorientable surfaces, *Trans. Amer. Math. Soc.* **243** (1978) 123–133.

[6] V. V. Goryunov and S. K. Lando, On enumeration of meromorphic functions on the line, preprint 1997.

[7] I. P. Goulden and D. M. Jackson, The combinatorial relationship between trees, cacti and certain connection coefficients for the symmetric group, *Europ. J. Combin.* **13** (1992) 357–365.

[8] I. P. Goulden and D. M. Jackson, Transitive factorizations into transpositions and holomorphic mappings of the sphere, *Proc. Amer. Math. Soc.* **125** (1997) 51–60.

[9] I. P. Goulden and D. M. Jackson, *Combinatorial Enumeration*, John Wiley and Sons, New York, 1983.

[10] I. P. Goulden and D. M. Jackson, Symmetric functions and Macdonald's result for top connection coefficients in the symmetric group, *J. Algebra* **166** (1994) 364–378.

[11] I. P. Goulden and S. Pepper, Labelled trees and factorizations of a cycle into transpositions, *Discrete Math.* **113** (1993) 263–268.

[12] N. Hanusse, *Cartes, constellations et groupes : quelques questions algorithmiques*, Ph.D. thesis, Université Bordeaux 1, 1997.

[13] F. Harary, G. Prins and W. T. Tutte, The number of plane trees, *Indag. Math.* **26** (1964) 319–329.

[14] A. Hurwitz, Über Riemann'sche Flächen mit gegebenen Verzweignungspunkten, *Math. Ann.* **39** (1891) 1–60.

(P. Erdős, A. Rényi, V. Sós eds.), North-Holland, Amsterdam, 1970, vol. 2, [illegible]

[16] G. Jones and D. Singerman, Belyi functions, hypermaps and Galois groups, *Bull. London Math. Soc.* **28** (1996) 561–590.

[17] V. A. Liskovets, A census of non-isomorphic planar maps, in "Algebraic methods in graph theory" (L. Lovász, V. Sós eds.), *Coll. Math. Soc. J. Bolyai*, North-Holland, Amsterdam, 1981, vol. 2, 479–494.

[18] M. Lothaire, *Combinatorics on words*, vol. 17 of *Encyclopedia of Mathematics and its Applications*, Cambridge University Press, 1984.

[19] W. S. Massey, *Algebraic topology: an introduction*, *Graduate Texts in Mathematics*, Vol. 56, Springer Verlag, 1967.

[20] A. D. Mednykh, Nonequivalent coverings of Riemann surfaces with a prescribed ramification type, *Siberian Math. J.* **25** (1984) 606–625.

[21] P. Moszkowski, A solution of a problem of Dénes, a bijection between trees and factorizations of cyclic permutations, *European J. Combin.* **101** (1989) 13–16.

[22] G. Schaeffer, Bijective census and random generation of Eulerian planar maps with prescribed vertex degrees, *Electronic J. Combinatorics* (`http://www.combinatorics.org/`) **4** (1997) R20.

[23] G. Schaeffer, *Conjugaison d'arbres et cartes combinatoires aléatoires*, Ph.D. thesis, Université Bordeaux 1, 1998.

[24] J.-P. Serre, *Topics in Galois theory*, Jones and Bartlett Publishers, Boston, 1992.

[25] V. Strehl, Minimal transitive products of transpositions: the reconstruction of a proof of A. Hurwitz, *Séminaire Lotharingien de Combinatoire* (`http://cartan.u-strasbg.fr/~slc/`) **37** (1996).

[26] W. T. Tutte, A census of slicings, *Canad. J. Math.* **14** (1962) 708–722.