Optimal Coding and Sampling of Triangulations

Dominique Poulalhon[1] and Gilles Schaeffer[2⋆]

[1] LIAFA, CNRS/Université Paris 7,
case 7014, 2, place Jussieu, 75251 Paris cedex 05, France,
http://www.liafa.jussieu.fr/~poulalho
[2] LIX, CNRS, École polytechnique, 91128 Palaiseau cedex, France,
http://www.lix.polytechnique.fr/Labo/Gilles.Schaeffer

Abstract. We present a bijection between the set of plane triangulations (*aka* maximal planar graphs) and a simple subset of the set of plane trees with two leaves adjacent to each node. The construction takes advantage of Schnyder tree decompositions of plane triangulations.
This bijection yields an interpretation of the formula for the number of plane triangulations with n vertices. Moreover the construction is simple enough to induce a linear random sampling algorithm, and an explicit information theory optimal encoding.
Finally we extend our bijection approach to triangulations of a polygon with k sides with m inner vertices, and develop in passing new results about Schnyder tree decompositions for these objects.

1 Introduction

This article addresses three problems on triangulations: coding, counting, and sampling. The triangulations that are considered here are finite combinatorial planar triangulations, or equivalently, maximal planar graphs. The results are obtained as consequences of a new bijection, between triangulations and trees in the simple class of plane trees with two leaves adjacent to each node. The extension of our approach to triangulations of a polygon with interior points involves new results about Schnyder tree decompositions.

1.1 Coding

The coding problem was first raised in algorithmic geometry: find an encoding of triangulated geometries which is as compact as possible. As demonstrated by previous work, a very effective "structure driven" approach consists in distinguishing the encoding of the combinatorial structure, – that is, the triangulation – from the geometry – that is, vertex coordinates (see [33] for a survey and [21] for an opposite "coordinate driven" approach). Two main properties of the combinatorial code are then desirable: *compactness*, that is minimization of the bit

⋆ Partially supported by the Research Training Network "Algebraic Combinatorics in Europe", and by the ACI-MD project "Geocomp".

length of code words and *linear complexity* of the complete coding and decoding procedure.

For the fundamental class $\mathcal{T}_n$ of triangulations of a sphere with $2n$ triangles, several codes of linear complexity were proposed, that are based on traversing the triangulation in a determined order and progressively outputting a codeword. Depending on the traversal policy, various bit length $\alpha n(1+o(1))$ were achieved from $\alpha = 4$ in [6, 13, 24], to $\alpha = 3.67$ in [27, 36], and recently $\alpha = 3.37$ bits in [7]. From the point of view of information theory the best α that can be guarantied is $\alpha_0 = \frac{1}{n}\log|T_n| \sim \frac{256}{27} \approx 3.245$ (see below). Other coders, like the valence directed coder of [38], do better on the "tame" triangulations that are typically produced by geometric modeling but offer no good guarantee on their worst case behavior. This of course does not contradict the entropy bound, which only states that among the codewords for all triangulations with $2n$ faces, some may be very short (typically those of "tame" triangulations) but at least one (in fact a majority) have length at least $\alpha_0 n$.

Finally, let us indicate that in the context of succinct data structures, the compactness problem was given radically different solutions. For instance, using separators and the recursive structure of triangulations, a representation was proposed in [25, 28] which can come ε-close to the optimal space requirement for the class of triangulations with a boundary (which has a larger entropy). Using a simpler 3-level approach inspired from [29], a space optimal data structure (allowing constant time queries on the triangulations) was proposed in [12], again for the class of triangulations with a boundary. However this approach focuses on data structures and on the support of queries in constant time rather than on producing simple compact codes for triangulations. In particular they heavily rely on auxiliary tables of size $O(n \cdot \frac{\log\log n}{\log n})$ that use a subdominant but large amount of space.

1.2 Counting

The exact enumeration problem for triangulations was solved by Tutte in the sixties [39]. The number of rooted triangulations with $2n$ triangles, $3n$ edges and $n+2$ vertices is

$$T_n = \frac{2\,(4n-3)!}{n!(3n-1)!}. \tag{1}$$

(This formula gives the above constant $\alpha_0 = \frac{256}{27}$.) More generally Tutte was interested in *rooted planar maps*: embedded planar multigraphs (with a distinguished half-edge to break possible symmetries) considered up to homeomorphisms of the sphere. He obtained several elegant formulas akin to (1) for the number of planar maps with n edges and for several subclasses (bipartite maps, 2-connected maps, 4-regular maps). It later turned out that constraints of this kind lead systematically to explicit enumeration results for subclasses of maps (in the form of algebraic generating functions, see [5] and references therein). A natural question in this context is to find simple combinatorial proofs explaining these results, as opposed to the technical computational proofs *à la Tutte*. This

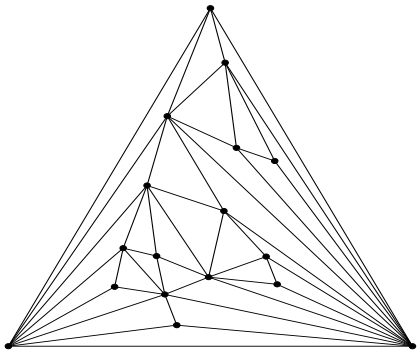

Fig. 1. A random triangulation with 30 triangles.

was done in a very general setting for maps without restrictions on multiple edges and loops [9, 34]. Two main ingredients are at the heart of this approach: dual breadth-first search to go from maps to trees, and a closure operation for the inverse mapping. When loops are forbidden, the first ingredient is no longer appropriate, but it was shown that it can be replaced by bipolar orientations [31, 35]. In particular in [31], we dealt with triangulations of a polygon with inner vertices *with multiple edges allowed*. When multiple edges are forbidden (as for the triangulations of the present paper), the situation appears completely different and neither of the previous methods directly apply.

It should be stressed that planar graphs have in general non-unique embeddings: a given planar graph may underlie many planar maps. This explains that, as opposed to the situation for maps, the entropy of the class of labelled planar graphs with n vertices has been obtained only very recently [23] and only rough bounds are known for unlabelled planar graphs, see [7, 30]. As opposed to this, according to Whitney's theorem, 3-connected planar graphs have an essentially unique embedding. In particular the class of triangulations is equivalent to the class of maximal planar graphs (a graph is maximal planar if no edge can be added without losing planarity).

1.3 Sampling

A perfect (resp. approximate) random sampling algorithm outputs a random triangulation chosen in $\mathcal{T}_n$ under the uniform distribution (resp. under an approximation thereof): the probability to output a specific rooted triangulation T with $2n$ vertices is (resp. is close to) $1/T_n$. Save for an exponentially small fraction of them, triangulations have a trivial automorphism group [32], so that as far as polynomial parameters are concerned, the uniform distribution on rooted or unrooted triangulations are indistinguishable for large n.

This question was first considered by physicists willing to test experimentally properties of two dimensional quantum gravity: it turns out that the proper

discretization of a typical quantum universe is precisely obtained by sampling from the uniform distribution on rooted triangulations [4]. Several approximate sampling algorithms were thus developed by physicists for planar maps, including for triangulations [3]. Most of them are based on Markov chains, the mixing times of which are not known (see however [22] for a related study). A recursive perfect sampler was also developed for cubic maps, but has at least quadratic complexity [1]. More efficient and perfect samplers were recently developed for a dozen of classes of planar maps [5, 36]. These algorithms are linear for triangular maps (with multiple edges allowed) but have average complexity $O(n^{5/3})$ for the class of triangulations.

Most random sampling algorithms are usually either based on Markov chains, or on enumerative properties. On the one hand, an algorithm of the first type performs a random walk on the set of configurations until it has (approximately) forgotten its start point. This is a very versatile method that requires little knowledge of the structures. It can even allow for perfect sampling in some restricted cases [40]. However in most cases it yields only approximate samplers of at least quadratic complexities. On the other hand, algorithms of the second type take advantage of exact counting results to construct directly a configuration from the uniform distribution [19]. As a result these perfect samplers often operate in linear time with little more than the number of random bits required by information theory bounds to generate a configuration [2, 15]. It is very desirable to obtain such an algorithm when the combinatorial class to be sampled displays simple enumerative properties, like Formula (1) for triangulations.

1.4 New results

The central result of this paper is a one-to-one correspondence between the triangulations in $\mathcal{T}_n$ and the *balanced trees* of a new simple family $\mathcal{B}_n$ of plane trees. We give a linear *closure algorithm* that constructs a triangulation out of a balanced tree, and conversely, a linear *opening algorithm* that recovers a balanced tree as a special depth-first search spanning tree of a triangulation endowed with its minimal Schnyder tree decomposition. Schnyder tree decompositions were introduced by Schnyder [37] to compute graph embeddings and have proved a fundamental tool in the study of planar graphs [8, 10, 17, 26]. The role played in this paper by the minimal Schnyder tree decompositions of triangulations is akin to the role of breadth-first search spanning trees in planar maps [9, 34, 36], and of minimal bipolar orientations in 2-connected maps [31, 35], however the closure algorithm is very different from the closure used in the latter works. Our bijection allows us to address the three previously discussed problems.

From the coding point of view, our encoding in terms of trees preserves the entropy and satisfies linearity: each triangulation is encoded by one of the $\binom{4n}{n}$ bit strings of length $4n$ with sum of bits equal to n. Optimal compactness can then be reached, still within linear time, using for instance [7, Lemma 7] or entropy encoders. It should be observed also that the code is produced using only two simple traversals of the triangulation (one to compute a Schnyder tree

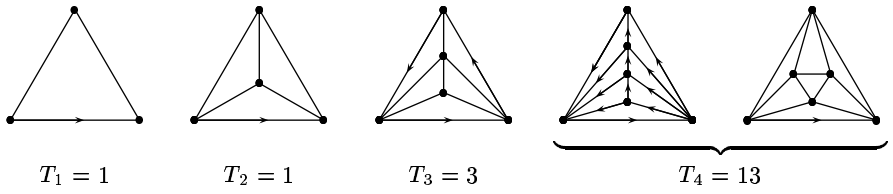

Fig. 2. The smallest triangulations with their inequivalent rootings.

decomposition and the second to produce our code), and that decoding can be done incrementally as bits become available.

From the exact enumerative point of view, the outcome of this work is a bijective derivation of Formula (1), giving it a simple interpretation in terms of trees. As far as we know, this is the first such bijective construction for any natural family of 3-connected planar graphs.

As far as random sampling is concerned, we obtain a linear time algorithm to sample random triangulations according to the (perfect) uniform distribution. In practice the speed we reach is about 100,000 vertices per second on a standard PC and triangulations with millions of vertices can be generated.

Since our central bijection heavily relies on Schnyder tree decompositions, it is natural to ask whether such structures exist for more general classes of graphs. This question was given a first answer in [17, 18], where Schnyder tree decompositions are developed for 3-connected planar graphs. As we showed in [20], Felsner's work allows the adaptation of the present approach to 3-connected planar graphs, although the technical details are much more involved. Another direction consists in considering triangulations of a surface with higher genus. We take what we consider as a first step in this direction by showing how the extension of our bijection to triangulations of a polygon leads to a Schnyder tree structure for these triangulations.

2 A one-to-one correspondence

Let us first recall some definitions, illustrated by Figure 2.

Definition 2.1. *A* planar map *is an embedding of a connected planar graph in the oriented sphere. It is* rooted *if one of its edges it distinguished and oriented; this determines a root edge, a root vertex (its origin) and a root face (to its right).*

A triangular map *is a rooted planar map with all faces of degree 3. It is a* triangulation *if moreover it has no loop nor multiple edge. A triangular map of size n has $2n$ triangular faces, $3n$ edges and $n+2$ vertices; the three vertices incident to the root face are called* outer, *as opposed to the $n-1$* inner *other ones. The set of triangulations of size n is denoted by $\mathcal{T}_n$.*

Drawing a planar map in the plane requires the choice of one face to become the unbounded face. A canonical choice is to use the root face as unbounded

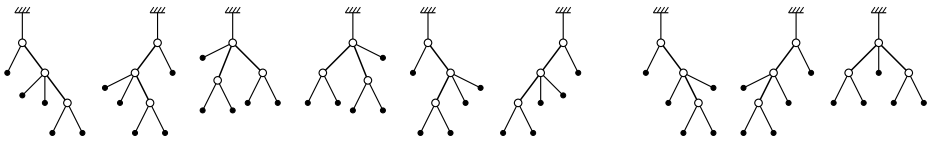

Fig. 3. The 9 elements of the set $\mathcal{B}_3$: the 6 distinct rootings of an asymmetric unrooted tree, and the 3 rootings of a symmetric one.

face, and we shall make this choice most of the time, except in Section 5. Once a planar map is embedded in the plane, the conventional orientation of the plane induces a natural positive orientation for cycles of the map, and allows us to define their *interior* (the bounded component of the plane that they define) and their *exterior* (the unbounded component).

2.1 From trees to triangulations

In view of Formula (1), it seems natural to ask for a bijection between triangulations and some kind of quaternary trees: indeed the number of such trees with n nodes is well known to be $\frac{(4n)!}{n!(3n+1)!}$. It proves however more interesting to consider a less classical family of plane trees.

First recall that *plane trees* can be defined as rooted planar maps with one single face and with a root vertex of degree 1. It will prove useful to make a distinction between *nodes* (vertices of degree at least 2) and *leaves* (vertices of degree 1), and between *inner edges* (connecting two nodes) and *pendant edges* (connecting a node to a leaf). The following definition is illustrated by Figure 3.

Definition 2.2. *Let $\mathcal{B}$ denote the set of plane trees with two leaves adjacent to each node; the number of nodes of a tree in $\mathcal{B}$ is called its* size*, and $\mathcal{B}_n$ denotes the subset of $\mathcal{B}$ consisting of trees of size n.*

The partial closure We introduce here a *partial closure* operation on trees in $\mathcal{B}$ (and more generally on planar maps with pendant edges) that merges leaves to nodes in order to create triangular faces.

Let B be a tree in $\mathcal{B}$. The border of the unbounded face consists of inner and pendant edges. An *admissible triple* is a sequence (e_1, e_2, e_3) of two successive inner edges followed by a pendant one in counterclockwise direction around the unbounded face. An admissible triple is thus formed of three edges $e_1 = (v, v')$, $e_2 = (v', v'')$ and $e_3 = (v'', \ell)$ such that v, v' and v'' are nodes and ℓ is a leaf.

The *local closure* of such an admissible triple (e_1, e_2, e_3) consists in merging the leaf ℓ with the node v so as to create a bounded face of degree 3. The pendant edge $e_3 = (v'', \ell)$ then becomes an inner edge (v'', v). For instance the first three edges after the root around the unbounded face of the tree of Figure 4(a) form an admissible triple, and the local closure of this triple produces the planar map

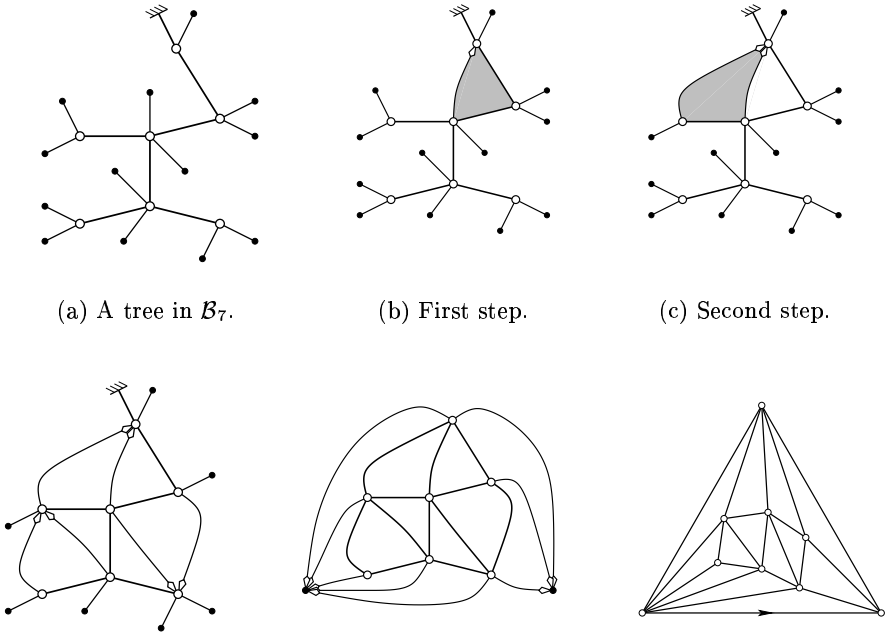

(a) A tree in $\mathcal{B}_7$. (b) First step. (c) Second step.

(d) Partial closure. (e) Two more vertices. (f) Complete closure.

Fig. 4. Complete closure construction on an element of $\mathcal{B}_7$.

of Figure 4(b). In turn, the first three edges of this map form a new admissible triple, and its local closure yields the map of Figure 4(c).

The *partial closure* $\tilde{B}$ of a tree B is the result of the greedy recursive application of local closure to all available admissible triples. The partial closure of the tree of Figure 4(a) is shown on Figure 4(d). At a given step of the construction, there are usually several admissible triples, but their local closures are independent so that the order in which they are performed is irrelevant and the final map $\tilde{B}$ is uniquely defined.

Balanced trees Any tree B in $\mathcal{B}$ has two more pendant edges than sides of inner edges incident to the unbounded face, and this property is preserved by local closures. The map $\tilde{B}$ has no admissible triple anymore, but some leaves remain unmatched. Therefore in its unbounded face no two inner edges can be consecutive: each inner edge lies between two pendant edges, as illustrated by Figures 4(d) and 5(a). More precisely the pendant edges and sides of inner edges alternate except at two special nodes: these two nodes v_0 and v_0' are each incident to two pendant edges with leaves ℓ_1, ℓ_2 and ℓ_1', ℓ_2' such that ℓ_2 (resp. ℓ_2')

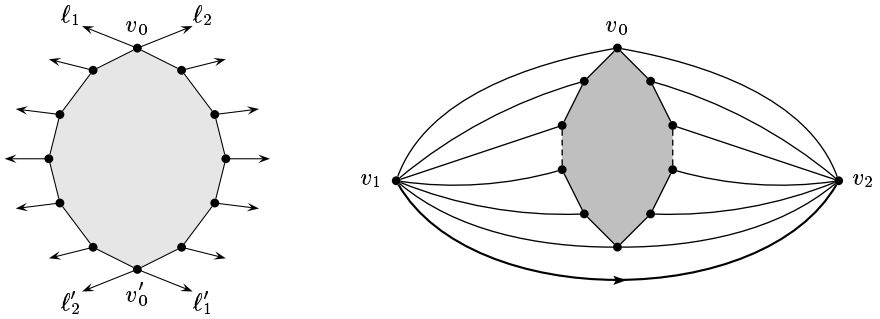

(a) After partial closure. (b) After complete closure.

Fig. 5. Generic situation.

follows ℓ_1 (resp. ℓ'_1) in clockwise order around the unbounded face as shown in Figure 5(a).

Observe that the partial closure of a tree is defined regardless of which of its leaves is the root. A tree B in $\mathcal{B}$ is *balanced* if its root leaf is one of the two leaves ℓ_1 or ℓ'_1 of its partial closure $\tilde{B}$. Let $\mathcal{B}^*$ be the subset of balanced trees in $\mathcal{B}$. The fourth, sixth and eighth trees in Figure 3 are balanced. The following immediate property shall be useful later on.

Property 2.3. Let B be a balanced tree; then each local closure merges a leaf ℓ and a node v such that ℓ comes after v in the left-to-right preorder on B. The closure edge e, naturally oriented to v, forms with some edges of B a cycle whose interior is on the right of e.

The complete closure Let B be a balanced tree in $\mathcal{B}^*$, and call v_0 and v'_0 the two special nodes of $\tilde{B}$ that are adjacent to the leaves ℓ_1, ℓ_2, ℓ'_1 and ℓ'_2. The *complete closure* of B is obtained from its partial closure in the following way (see Figures 4 and 5(b)):

1. merge ℓ_1, ℓ'_2 and all leaves in between (counterclockwise around the unbounded face) at a new vertex v_1;
2. merge ℓ'_1, ℓ_2 and all leaves in between at a new vertex v_2;
3. add an (oriented) root edge (v_1, v_2), from v_1 to v_2.

The result of this complete closure is clearly a triangular map (see Definition 2.1), which we denote $\bar{B}$. We shall prove the following theorem in Section 3:

Theorem 2.4. *The complete closure is a one-to-one correspondence between the set $\mathcal{B}^*_n$ of balanced plane trees with n nodes and two leaves adjacent to each node, and the set $\mathcal{T}_n$ of rooted triangulations of size n.*

Although the constructions are formally unrelated, the terminology we use here is reminiscent of [9, 31, 34], where bijections were proposed between some trees and planar maps with multiple edges.

Apart from the orientation of the root, the complete closure can be more generally defined for any tree in $\mathcal{B}_n$: it consists of Steps 1 and 2 above, together with the adjunction of a marked but non oriented edge (v_1, v_2). The result $\bar{B}^0$ of this unrooted complete closure of a tree then only depends on the unrooted tree B^0 underlying B, and not on which of the $2n$ leaves is the root of the tree B.

In general $2n$ rooted trees are associated to a given unrooted tree B^0 (or n in the exceptional case where B^0 has a global symmetry). As already indicated, these $2n$ trees have the same image $\bar{B}^0$ by unrooted complete closure. This image is a triangular map with a marked (non oriented) edge (v_1, v_2), to which can be associated two (or exceptionally one) rooted triangular maps. Unrooted complete closure thus defines a "$2n$-to-2" correspondence between rooted trees and rooted triangulations. However we prefer to deal with *balanced* trees in order to have a plain one-to-one correspondence.

2.2 From triangulations to trees

Covering trees Given a triangular map T and a subset E' of its edge set $E(T)$, a *covering tree* of E' is a tree B that can be obtained from T by deleting the edges of $E(T) \setminus E'$ and opening some edges of E in order to form pendant edges.

In particular the closure of a tree B in $\mathcal{B}$ produces a triangular map T for which B is a covering tree. Conversely, if a triangular map T is to be obtained by closure from a tree B, this tree must be looked for among covering trees of T.

Minimal Schnyder tree decomposition We shall use the following notion, due to Schnyder [37].

Definition 2.5. *Let T be a triangulation, with root edge (v_1, v_2), and with v_0 its third outer vertex. A* Schnyder tree decomposition *of T is a colouring of its inner edges in three colours c_0, c_1 and c_2 satisfying the following conditions:*

- *for each $i \in \{0, 1, 2\}$, edges of colour c_i form a spanning tree of $T \setminus \{v_{i+1}, v_{i+2}\}$ rooted on v_i; this induces an orientation of edges of colour c_i toward v_i, such that each vertex has exactly one outgoing edge of colour c_i;*
- *around each inner vertex, outgoing edges of each colour always appear in the cyclic order shown on Figure 6(a), and entering edges of colour c_i (if any) appear between outgoing edges of the two other colours.*

From now on, this second condition is referred to as Schnyder condition.

The Schnyder rule can be restated in the following way. For any vertex v, let us define a *corner* at v as a couple of edges (e, e') incident to v such that e' immediately follows e when turning clockwise around v. A corner (e, e') at v satisfies the *local Schnyder rule* if

(a) Schnyder rule for a vertex.

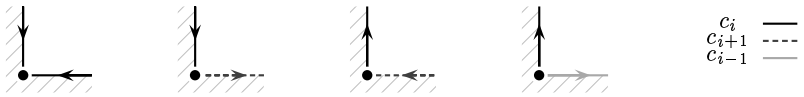

(b) Local Schnyder rules for a corner.

Fig. 6. Property of a Schnyder tree decomposition.

- either e is ingoing at v with colour c_i and e' is ingoing at v with colour c_i,
- or e is ingoing at v with colour c_i and e' is outgoing at v with colour c_{i+1},
- or e is outgoing at v with colour c_i and e' is ingoing at v with colour c_{i+1},
- or e is outgoing at v with colour c_i and e' is outgoing at v with c_{i-1}.

These four types of Schnyder corners are shown on Figure 6(b). A vertex v with outdegree 3 satisfies the Schnyder rule if and only if any corner at v satisfies the local Schnyder rule, so that both are equivalent as far as triangulations are concerned. However this local point of view will be useful in Section 5.

Schnyder tree decompositions of triangulations satisfy a number of nice properties [14, 17, 37], among which we shall use the following ones:

Proposition 2.6. – *Every triangulation has a Schnyder tree decomposition.*
- *The set of Schnyder tree decompositions of a triangulation can be endowed with an order for which minimal and maximal elements are unique.*
- *The minimal Schnyder tree decomposition of a triangulation T is the unique Schnyder tree decomposition of T that has no counterclockwise circuit.*
- *The minimal Schnyder tree decomposition of a triangulation can be computed in linear time.*

Depth-first search opening Let T be a triangulation endowed with its minimal Schnyder tree decomposition, and let $E(T)$ denote the set of edges of T. Let (v_1, v_2) be its root edge, oriented from v_1 to v_2, and v_0 be the third outer vertex. We construct a covering tree on $E(T) \setminus \{(v_1, v_2)\}$ using a right-to-left depth-first search traversal of T modified to traverse edges only in the direction opposite to their orientation in the Schnyder tree decomposition:

1. delete (v_1, v_2), and detach (v_0, v_1) and (v_0, v_2) from v_1 and v_2 to form two leaves ℓ_1, ℓ_2 attached to v_0,
2. set $v \leftarrow v_0$ and $e \leftarrow (v_0, \ell_2)$,
3. as long as $e \neq (v_0, \ell_1)$, repeat:
   (a) $e' \leftarrow (v, u)$, the edge after e around v *in clockwise direction*,
   (b) *special orientation test:* if e' is unmarked and oriented from v to u in the Schnyder tree decomposition, then mark e', detach it from u to create a leaf ℓ attached to v, and set $e \leftarrow (v, \ell)$,
   (c) *standard traversal condition:* if e' is unmarked and u was already visited, then detach e' from u to create a leaf ℓ attached to v and set $e \leftarrow (v, \ell)$,
   (d) otherwise, mark e' and set $e \leftarrow (u, v)$ and $v \leftarrow u$.

As usual for a depth-first search, Step 3c prevents the opening algorithm from creating a cycle of marked edges. The set of marked edges therefore forms a tree and the algorithm clearly terminates. Let $S(T)$ be the tree containing all marked edges. Without Step 3b, the opening algorithm would be a standard right-to-left depth-first search, the inner edges of $S(T)$ would form a spanning tree of $T \setminus \{v_1, v_2\}$, and $S(T)$ would be a covering tree of $E(T) \setminus \{(v_1, v_2)\}$. A priori, because of Step 3b, one could expect the opening algorithm to visit only part of the edges. We shall prove in fact the following proposition:

Proposition 2.7. *For any triangulation T, the tree $S(T)$ is a covering tree of $E(T) \setminus \{(v_1, v_2)\}$ and belongs to $\mathcal{B}^*$. Moreover it is the unique balanced tree having the triangulation T as complete closure.*

Because of the minimal orientation of T (without counterclockwise circuit), we shall see that the condition of Step 3c is actually never satisfied. This line of the algorithm could thus as well be ignored: it was included only to make clear the fact that the algorithm terminates.

3 Proofs

3.1 The closure produces a triangulation

At all steps of the closure construction, unmatched leaves are incident to the unbounded face and each local closure creates a triangular face without violating planarity. It is thus clear that the complete closure yields a triangular map $\bar{B}$ with outer vertices v_0, v_1 and v_2, and with exactly two more vertices than B has nodes. Let us show that $\bar{B}$ is indeed a triangulation, *i.e.* that it has no multiple edge.

Let B be a balanced tree in $\mathcal{B}^*$. By definition the root leaf ℓ_1 of B is immediately followed around v_0 in clockwise direction by a second leaf ℓ_2. Assign colour c_1 to ℓ_1, colour c_2 to ℓ_2, and colour c_0 to the other edges incident to v_0. Upon orienting all inner edges of B toward v_0 and all pendant edges toward their leaf, all vertices but v_0 have three outgoing edges. Since the tree B is acyclic, its orientation induces a unique colouring of edges satisfying the Schnyder condition at all vertices but v_0, which we call the *canonical Schnyder colouring* of B. An example is shown on Figure 7(a).

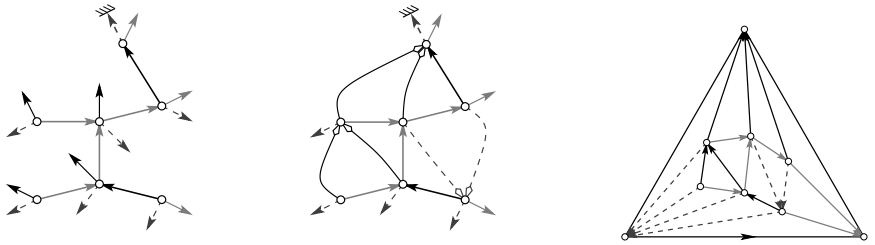

(a) On the tree, (b) the partial closure, (c) and the complete closure.

Fig. 7. Colouring of the example of Figure 4.

Lemma 3.1. *The orientation and colouring of edges still satisfy the Schnyder condition on each node but v_0 after the partial closure of B.*

Proof. This lemma is checked iteratively, by observing that each face created during the partial closure falls into one of the four types indicated on Figure 8 (up to cyclic permutation of colours). Indeed, consider an admissible triple (e_1, e_2, e_3). Assuming without loss of generality that the pendant edge e_3 to be closed is of colour c_0, only two colours are possible for e_2 in view of the Schnyder condition at v''. In each case again, only two colours are possible for e_1. Finally in all four cases, the merging of ℓ into v does not contradict the Schnyder condition at v. □

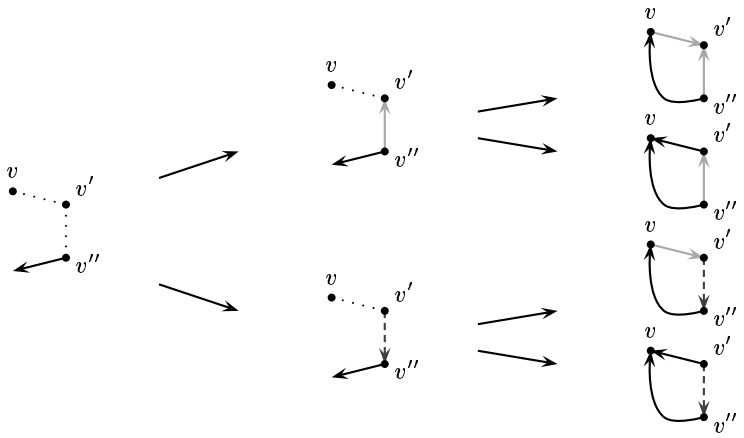

Fig. 8. The different cases of closure of a leaf.

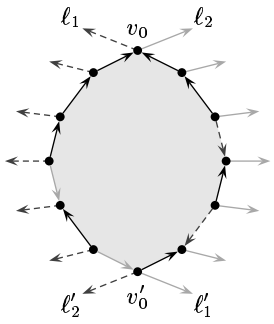

Fig. 9. Colouring of the partial closure.

Lemma 3.2. *After the complete closure, the Schnyder condition is satisfied at each inner vertex, and, apart from the three outer edges, each outer vertex v_i is incident only to entering edges of colour c_i.*

Proof. As illustrated by Figure 9, the Schnyder condition on nodes along the border of the partial closure implies that all pendant edges between ℓ_1 and ℓ'_2 (resp. ℓ_2 and ℓ'_1) are of colour c_1 (resp. c_2). This is readily checked iteratively by a case analysis akin to the previous one. □

Let us define a *Schnyder coloured orientation* of a planar map as an orientation and a 3-colouring of its inner edges as in Lemma 3.2. The following proposition can be seen as an independent result on Schnyder tree decompositions.

Proposition 3.3. *Any triangular map endowed with a Schnyder coloured orientation is in fact a triangulation endowed with a Schnyder tree decomposition.*

Proof. Let us first consider the colour c_0 of the outer vertex v_0. By Schnyder condition each inner vertex has exactly one outgoing edge of colour c_0. In particular any cycle of edges of colour c_0 is in fact a circuit. Moreover from each inner vertex originates a unique oriented path of colour c_0, ending either in v_0 or on a circuit of colour c_0.

Now consider two paths with distinct colours, say c_0 and c_1. In view of the Schnyder colouring, a crossing between these two paths is necessarily of the type:

Hence two such paths can not cross more than once, otherwise this rule would be violated at least at one crossing (where crossing is taken in the (weak) sense of having one vertex in common, even if this is just the origin of the path.) In particular, this excludes multiple edges with different colours.

On the other hand, since each vertex has only one outgoing edge of each colour, two monochrome circuits of the same colour can not cross at all. As a

consequence, any two monochrome circuits must be vertex disjoint. Therefore monochrome circuits are ordered by inclusion with respect to the unbounded face.

Now consider a vertex v on an innermost circuit C. The Schnyder condition at v provides an edge e going out of v into the inner region delimited by C. Since this region contains no monochrome circuit, the oriented path extending e has to cross C a second time, in contradiction with the previous discussion. This excludes monochrome circuits and proves that, for each $i = 0, 1, 2$, edges of colour c_i form an oriented tree rooted at v_i. In particular multiple edges are excluded, and the colouring satisfies the definition of a Schnyder tree decomposition. □

Thus the complete closure of B is a triangulation endowed with a Schnyder tree decomposition, but not any Schnyder tree decomposition:

Property 3.4. If a (triangular) face of $\tilde{B}$ is oriented so that its sides form a circuit, then this circuit is necessarily oriented in the clockwise direction. More generally, each circuit in $\tilde{B}$ is created by the closure of a (last) leaf, the orientation of which imposes on the circuit to be clockwise.

Combining Proposition 3.3 and Property 3.4, we obtain the following corollary that concludes the first part of the proof.

Corollary 3.5. *The closure maps any balanced tree B of $\mathcal{B}_n^*$ endowed with its canonical Schnyder colouring on a triangulation with $n+2$ vertices endowed with its minimal Schnyder tree decomposition.*

3.2 The depth-first search opening is inverse to closure

The following Lemmas 3.6 to 3.9 imply Proposition 2.7, and, together with Corollary 3.5, conclude the proof of Theorem 2.4.

Lemma 3.6. *The depth-first search opening visits all vertices of $T \setminus \{v_1, v_2\}$.*

Proof. Assume that the inner vertex v is not visited by the opening algorithm, that is to say, v does not belong to $S(T)$. By definition of Schnyder tree decompositions, there is a unique oriented path P of colour c_0 starting in v and ending in v_0. Let t be the last vertex on P that does not belong to $S(T)$, and $u \in S(T)$ the next vertex on P. Then the edge between t and u is oriented toward u. Let (s, u) be the first edge before (t, u) in clockwise direction around u such that s belongs to $S(T)$. Then (s, u) is oriented from s to u, and is the first edge of a backward oriented path P' from u to an earlier visited vertex w. But this path forms with P a counterclockwise circuit, which contradicts the minimality of the orientation. □

Lemma 3.7. *The conditions of Step 3c are never satisfied.*

Proof. Consider the first time the conditions of Step 3c are satisfied. Up to that point an oriented tree S was constructed that contains v and u but not the edge (v, u). Since the unmarked edge (v, u) was not considered by Step 3b, it is oriented from u toward v.

Let E be the set of edges that were already cut by Step 3b. Then S is the initial part of the right-to-left depth-first search tree of $T \setminus E$. In particular, only ancestors of v in S are partially visited, and since the edge (u, v) is probed from v, u is an ancestor of v in that tree (and in $S(T)$). But then the tree $S(T)$ contains an oriented path from v to u, that forms a counterclockwise circuit with (u, v). This contradicts the minimality of the orientation. □

Lemma 3.8. *Edges that are cut by the opening algorithm lie on the left-hand side of the tree, in the sense of Property 2.3. Hence the complete closure of $S(T)$ is T.*

Proof. As already observed, as the algorithm proceeds, the tree that is constructed can be thought of as the right-to-left depth-first search tree of a submap of T. In particular when the algorithm probes edge $e' = (v, u)$ from vertex v, the vertex u must be before v in the left-to-right preorder, as in Property 2.3.

To check that the complete closure of $S(T)$ is T, it is sufficient to check that a cut edge would be properly closed by the local application of the closure algorithm. Since cut edges are bordered on one side by the unbounded face and the final tree is a spanning tree, then the other face is bounded, that is, triangular. Hence when $e' = (v, u)$ is cut, the vertex u lie two corners away from v along the unbounded face in clockwise direction, as specified for admissible triples. □

Lemma 3.9. *At most one covering tree of $E(T) \setminus \{(v_1, v_2)\}$ satisfies Property 2.3.*

Proof. Assume there are two such trees S and S'. Consider a left-to-right depth-first search traversal of both trees in parallel. Let $e = (v, u)$ be the first edge met that belongs to one of them – say S – and not to the other one. The tree S' being also a spanning tree, there exists in S' a path from u to v_0, the first edge of which (u, t) is oriented from u toward t. This orientation forbids the edge (u, t) to belong to the tree S; thus it corresponds in that tree to the closure of a leaf of u. But since the edge (v, u) has been visited before (u, t) in the depth-first search traversal, this contradicts Property 2.3. □

4 Applications

4.1 An explicit optimal code for triangulations

As a first byproduct of Theorem 2.4, we obtain an encoding of triangulations in $\mathcal{T}_n$ by balanced trees in $\mathcal{B}_n$. Since a triangulation can be endowed with its minimal Schnyder tree decomposition in linear time (Proposition 2.6), the tree code can be obtained in linear time. Elements of $\mathcal{B}_n$ can themselves be coded by

bit strings of length $4n-2$ and weight $n-1$ using a trivial variant of the usual prefix code for trees:

Lemma 4.1. *A tree B in $\mathcal{B}$ can be linearly represented by the word $s(B)$ that is obtained by writing 1 for "down" steps along inner edges, and 0 for leaves and for "up" steps along inner edges, during a right-to-left depth-first search traversal.*

Hence we obtain a code for triangulations in $\mathcal{T}_n$ which is a subset of the set S of bit strings with length $4n-2$ and weight $n-1$. According to [7, Lem. 7] it can be given in linear time a representation as a bit string of length $\log|S| + o(n) \sim \log\binom{4n}{n} \sim \frac{256}{27}n$.

4.2 A bijective proof of Formula (1)

The enumeration of simple families of trees is a classical problem.

Proposition 4.2. *The set $\mathcal{B}_n$ has cardinality $\frac{2}{4n-2}\cdot\binom{4n-2}{n-1}$.*

Proof. As for classical prefix code of trees, the code words corresponding to trees of $\mathcal{B}_n$ can be easily characterized: they are the bit strings of length $4n-2$ with weight $n-1$ such that any proper prefix u satisfies $3|u|_1 - |u|_0 > -2$, where $|u|_i$ denotes the number of occurrences of the letter i in u. Now the number of such bit strings is readily obtained by the cycle lemma (see [16]): in each cyclic class of words with length $4n-2$ and weight $n-1$, exactly 2 elements among $4n-2$ are code words (or 1 among $2n-1$ for symmetric classes). □

Now as seen in Section 2.1, any tree in $\mathcal{B}_n$ has two particular leaves among its $2n$ ones, and it is balanced if and only if one of these two is its root. Hence the proportion of balanced trees in $\mathcal{B}_n$ is $\frac{2}{2n}$. From Theorem 2.4 we obtain:

Theorem 4.3. *The number of triangulations with $2n$ triangles, $3n$ edges and $n+2$ vertices is $\frac{2}{2n}\cdot\frac{2}{4n-2}\cdot\binom{4n-2}{n-1}$, which is exactly Formula (1).*

4.3 Linear time perfect random sampling of triangulations

The closure construction provides a sampling algorithm with linear complexity:

1. generate a random bit string of length $4n-2$ and weight $n-1$;
2. choose randomly one of its two cyclic shift that code an element of $\mathcal{B}_n$;
3. decode this word to construct the corresponding tree;
4. construct its partial closure by turning around the tree; using a stack, this can be done in at most two complete turns, hence in linear time;
5. complete the closure and choose a random orientation for the edge (v_1, v_2).

Proposition 4.4. *This algorithm produces in linear time a random triangulation uniformly chosen in $\mathcal{T}_n$.*

Observe that Steps 1–3 correspond to a special case of the classical algorithm described *e.g.* in [2] for sampling trees.

5 Triangulations of a polygon

5.1 A counting formula and cycles of trees

The triangulations we considered in the previous sections have the topology of the sphere. It is also possible to define triangulations with the topology of other surfaces (of higher genus and/or with boundary). We deal here with triangulations of a disk, that is, of a surface of genus 0 with a connected boundary. Upon gluing a face to this boundary, they can be considered as planar maps with all faces but one of degree 3.

Definition 5.1. *A* rooted near-triangular map *is a rooted planar map with all non-root faces of degree 3, and with a root face of degree k for some $k \geq 3$. An* annular triangular map *is a rooted near-triangular map with a marked non-root face.*

Rooted near-triangulations *and* annular triangulations *are defined accordingly by the restriction that loops and multiple edges are forbidden. The sets of rooted near-triangulations and of annular triangulations with a root face of degree k, m inner vertices and $2m+k-2$ inner triangular faces are respectively denoted $\mathcal{T}_{m,k}$ and $\mathcal{T}^{\bullet}_{m,k}$.*

Near-triangulations can be also viewed as internally 3-connected triangulations of a polygon, *i.e.* such that chords are allowed but at least 3 inner vertices must be removed to disconnect the map. Except in Section 5.4, we shall always represent annular triangular maps and annular triangulations in the plane with their marked face taken as unbounded face (instead of the root face). However we shall keep the terminology *inner vertices* for the vertices of an annular map that are not on the polygon, even though this terminology makes more sense when the root face is used as unbounded face. Similarly, given a vertex v of the polygon and two incident edges e_1 and e_2, the *inner* wedge between e_1 and e_2 at v is among the two wedges delimited by e_1 and e_2, the one that does not intersect the interior of the polygon.

The following formula is due to Brown:

Theorem 5.2 (Brown [11]). *The number of triangulations of a rooted k-gon with m inner vertices and $2m+k-2$ inner triangular faces is*

$$|\mathcal{T}_{m,k}| = \frac{2(2k-3)!}{(k-1)!(k-3)!}\frac{(2k+4m-5)!}{m!(2k+3m-3)!}$$

The number of annular triangulations with root face of degree k and m inner vertices is $(2m+k-2)$ times this number, that is:

$$|\mathcal{T}^{\bullet}_{m,k}| = (2m+k-2)|\mathcal{T}_{m,k}| = \binom{2k-4}{k-3} \cdot \frac{2k-3}{3m+2k-3}\binom{4m+2k-4}{m}.$$

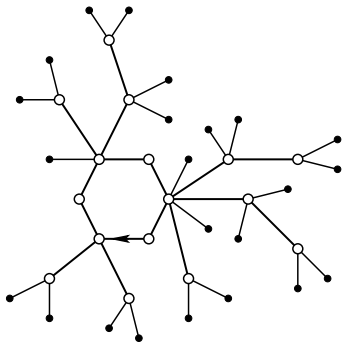

Fig. 10. Example of an hexagonal forest with 10 inner vertices.

This formula agrees with Formula 1: for $k = 3$ the total number of vertices is $m + 3 = n + 2$ and

$$|\mathcal{T}^{\bullet}_{m,3}| = \frac{3}{3m+3}\binom{4m+2}{m} = \frac{1}{n}\frac{(4n-2)!}{(n-1)!(3n-1)!} = (2n-1)|\mathcal{T}_n|.$$

In view of Theorem 5.2, we define the following family of planar maps akin to trees, as illustrated by Figure 10.

Definition 5.3. *A* k-gonal forest *is a rooted planar map with two faces satisfying the following conditions:*

- *The border of the root face is a cycle of length* k*, which is referred to as the* polygon *or the* k-gon *of the forest.*
- *There are* $k-3$ *pendant edges attached to the* k*-gon.*
- *Any other edge incident to the* k*-gon is the root of a planted plane tree in the set* $\mathcal{B}^{\perp}$ *of trees planted on a leaf and such that every node is adjacent to two non-root leaves. The nodes of these subtrees are referred to as the* inner vertices *of the forest.*

Let $\mathcal{B}_{k,m}$ *denote the set of* k*-gonal forests with* m *inner vertices.*

From now on a k-gonal forest will always be represented in the plane taking the non root face as unbounded face.

Proposition 5.4. *The number of* k*-gonal forests with* m *inner vertices is*

$$|\mathcal{B}_{m,k}| = \binom{2k-4}{k-3} \cdot \frac{2k-3}{3m+2k-3}\binom{4m+2k-4}{m}.$$

Proof. A proof very similar to the proof of Proposition 4.2 allows to show that the cardinality of the set $(\mathcal{B}^{\perp})^{*}_{m}$ of sequences of trees of $\mathcal{B}^{\perp}$ with a total of m nodes is $\frac{1}{3m+1}\binom{4m}{m}$, and more generally that the set of p-uples of such sequences of trees with a total of m nodes has cardinality $\frac{p}{3m+p}\binom{4m+p-1}{m}$. The proposition then follows from the way to produce uniquely all k-gonal forests with m inner vertices:

- start with a rooted k-gon, $k \geq 3$;
- attach $k-3$ pendant edges in the unbounded face of the k-gon, in one of the $\binom{2k-4}{k-3}$ possible ways; this produces $2k-3$ *places*, where subtrees can be attached;
- select a $(2k-3)$-uple of sequences of trees, among the $\frac{2k-3}{3m+2k-3}\binom{4m+2k-4}{m}$ possible ones, and attach these sequences in the previous $2k-3$ places.

□

5.2 Revisiting the closure

Let us apply local closures, as defined in Section 2.1, to k-gonal forests. Recall that the local closure of a pendant edge e takes place in clockwise direction around the tree, thus creating a bounded triangular face on the right hand side of e. The partial closure of a k-gonal forest is obtained by greedy iteration of local closure until no more local closures are possible.

In a k-gonal forest with m inner vertices, the number of pendant edges is by definition $k-3+2m$, while the number of sides of inner edges incident to the unbounded face is $k+2n$. There are thus three more sides of edges than leaves, and this property is preserved by local closures. In particular as long as there are leaves, there exists one leaf that is followed by at least two sides of inner edges, that is, there exists one leaf eligible for local closure. The iteration of local closures thus stops when all leaves have been matched. At this point, three sides of edges remain, so that the unbounded face is triangular.

The complete closure of a k-gonal forest is the annular map obtained by partial closure upon marking the unbounded face. The analog of Theorem 2.4 is then the following.

Theorem 5.5. *Complete closure is a one-to-one correspondence between the set $\mathcal{B}_{k,m}$ of k-gonal forests with m inner vertices, and the set $\mathcal{T}^*_{k,m}$ of annular triangulations with root face of degree k and m inner vertices.*

5.3 α-orientations

Given a k-gonal forest, orient the k-gon in clockwise direction, the inner edges of the forest toward the k-gon and the pendant edges toward their leaf. Then by construction the closure does not change the outdegree of vertices, so that after closure, each inner vertex has outdegree 3 and each vertex of the polygon with i pendant edges has outdegree $i+1$.

Proposition 5.6. *The image of a k-gonal forest by closure with orientations is an annular triangular map with root face of degree k endowed with an orientation such that there is no counterclockwise circuit (when the marked face is taken as unbounded face).*

Proof. The fact that the orientation has no counterclockwise circuit follows from the same remark as in Property 3.4: a circuit must be completed by the local closure of a last edge, which by definition cannot occur in counterclockwise direction. □

Orientation without counterclockwise circuits play a particular role in Felsner's theory of α-orientations [18]. Let us recall some elements of this theory.

Definition 5.7. *Let α be a valuation of the vertices of a planar map T such that for a vertex v with degree d, $\alpha(v) \in \{0, \ldots, d\}$. Then an* α-orientation *is an orientation of the edges of the map T such that each vertex v has outdegree $\alpha(v)$.*

Given a planar map T and a valuation α such that T admits an α-orientation, the set of all α-orientations can be partially ordered by the relation induced by the operation of reversing circuits. (Recall again that the clockwise or counterclockwise character of a circuit is defined with respect to a choice of unbounded face.)

Theorem 5.8 (Felsner [18], Ossona de Mendez [14]). *Circuit reversal endows the set of α-orientations of a planar map with a lattice structure. In particular there is a unique minimal element of the lattice, that is, a unique α-orientation without counterclockwise circuit (with respect to a given choice of unbounded face).*

In terms of α-orientations, the result of Proposition 5.6 is therefore that the closure of a k-gonal forest is an annular triangular map endowed with its minimal α-orientation with respect to the following valuation:

- for any inner vertex v, $\alpha(v) = 3$,
- for any vertex v of the k-gon holding i pendant edges, $\alpha(v) = i + 1$.

In particular the definition of the orientation of a k-gonal forest implies that the k-gon forms a clockwise circuit in the minimal α-orientation of its closure. In fact this is a general property of α-orientations of annular triangular maps for a large class of valuations α:

Proposition 5.9. *Let T be an annular triangular map and α a valuation such that $\alpha(v) \geq 1$ if v belongs to the root polygon and $\alpha(v) = 3$ otherwise. If T can be endowed with an α-orientation, then the root polygon forms a clockwise circuit in the minimal α-orientation of T (for any choice of unbounded face that is not the root face).*

Proof. Consider the map T' obtained from T by contracting the k edges of the polygon, and define a valuation α' by $\alpha'(v_0) = \sum_v(\alpha(v) - 1)$ where the summation runs over the k vertices of the polygon and v_0 is the contraction vertex, and $\alpha'(v) = \alpha(v)$ otherwise.

Starting with the unique minimal α'-orientation of T', one produces an α-orientation of T without counterclockwise circuit by replacing the polygon as a clockwise circuit (one easily checks that the circuit cannot create new counterclockwise circuits). The uniqueness of this α-orientation of T without ccw circuit ends the proof. □

5.4 Schnyder tree decompositions for annular domains

The next step is to introduce the analog of Schnyder tree decompositions. Let B be a k-gonal forest endowed with its orientation, and let T be its closure, which inherits a minimal α-orientation.

The colouring of a k-gonal forest is defined by setting the colour of the root to c_0 and then propagating the colours using the local Schnyder rule at all corners except inside the polygon.

Proposition 5.10. *A k-gonal forest admits a unique colouring such that the root has colour c_0 and the local Schnyder rule is satisfied at all corners except inside the polygon.*

Proof. Let us number up from 0 the edges of the polygon, in clockwise direction starting with the root, and the vertices, starting with the endpoint of the root. By construction, the outgoing edges from the vertices of the polygon are the k edges of the polygon and the $k-3$ pendant edges. The application of the local Schnyder rule to all the inner corners of the ℓth vertex implies that the colour of the $(\ell+1)$th edge is c_{i+1-j} if the vertex carries j pendant edges and the ℓth edge has colour c_i. The propagation of this relation in the clockwise direction from edge 0 (which is the root) to edge k (which is again the root) yields the condition $c_0 = c_{k-(k-3)}$ which is coherent since colours are understood modulo 3.

Finally all other vertices have outdegree 3, so that the rest of the propagation uses the standard Schnyder rule. □

Lemma 5.11. *After the complete closure of a k-gonal forest, the colouring and orientation induced on the resulting annular triangular map satisfy the local Schnyder rule at all corners except inside the polygon.*

Proof. The proof of Lemma 3.2 immediately adapts, since it relies only on the local Schnyder rule at the 3 corners involved in a local closure, as illustrated by Figure 8. □

For the purpose of the next lemma, it is more convenient to change our representation of the map in the plane by taking the root face as unbounded face: in particular the notion of interior and exterior of a cycle are taken in this proof with respect to the root face.

The following lemma is the analog of Proposition 3.3 for annular triangular maps.

Lemma 5.12. *Let T be an annular triangular map with a colouring and an orientation of edges such that the local Schnyder rule is satisfied at all corners except inside the polygon. Then for each $i = 0, 1, 2$, the set of edges of colour c_i forms a forest of trees, and each of these trees has a root on the polygon toward which all edges are oriented.*

Proof. The proof is essentially the same as the proof of Proposition 3.3 but involves some technical details due to the fact that some vertices may have outdegree larger than 3.

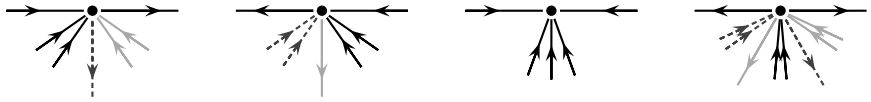

(a) (b) (c) (d)

Fig. 11. Possible configurations on vertices of a monochrome c_0 cycle with no c_0 edge toward the interior.

Let us first observe that inner vertices have outdegree 1 in each colour, so that any monochrome oriented path must end on the polygon or on a monochrome *inner circuit* (that is, a circuit visiting only inner vertices). But the proof of Proposition 3.3 immediately rules out the existence of such inner circuits, so that all monochrome oriented paths end on the polygon.

Let us now consider a monochrome cycle C, say of colour c_0. It must visit at least one vertex of the polygon, otherwise it would be a circuit since inner vertices have outdegree 1.

Consider now a vertex x of the cycle C, and assume there is an edge e of colour c_0 going out of x inside the cycle C. Then the monochrome oriented path starting with e defines with a part of C a new cycle C' of colour c_0 such that the interior of C' is strictly included in the interior of C.

Upon iterating the previous construction, we can thus take C such that there is no edges of colour c_0 going out of a vertex of C toward its interior. In view of the local Schnyder rule, this creates limitations on the colours of edges that can leave or arrive to a vertex of C on the inside. The possible configurations are listed in Figure 11. If C is a circuit, we obtain a contradiction as in Proposition 3.3.

Therefore C is not a circuit, and then it can be decomposed as the concatenation of $2k$ oriented paths: $C = P_1P_2 \cdots P_{2k}$ such that the P_{2i+1} are clockwise around the interior of C and the P_{2i} are counterclockwise around the interior of C. Let x_i be the common vertex of P_i and P_{i+1}. Then generic vertices of path P_{2i+1} are of type 11(a), generic vertices of path P_{2i} are of type 11(b), vertices x_{2i+1} are of type 11(c) vertices x_{2i} are of type 11(d).

Considering outgoing edges of vertices x_2i, it is then easy to check that these constraints cannot all be satisfied in the plane. □

Lemma 5.13. *Let T be an annular triangular map with a colouring and an orientation of edges such that the local Schnyder rule is satisfied at all corners except inside the polygon. Then two monochrome oriented paths cannot meet twice. In particular T has no multiple edges, that is, T is an annular triangulation.*

Proof. The proof is very similar to the previous one. Consider two monochrome oriented paths P_0 of colour c_0 and P_1 of colour c_1 that meet twice. They define a bicolour cycle C. Using an iteration as in the proof of the previous lemma, one can easily exclude, without loss of generality, that there exists inside C an edge of colour c_0 originating from a vertex of P_0 or an edge of colour c_1 originating from

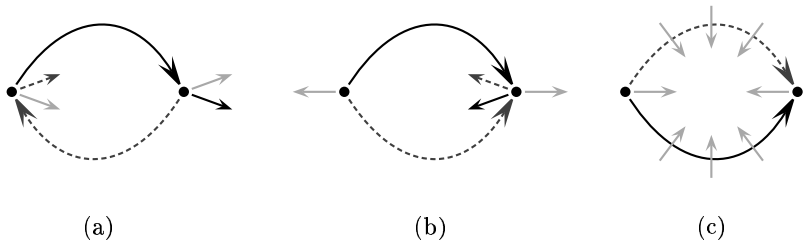

Fig. 12. Possible bicolour cycles.

a vertex of P_1. Now C can be either (a) a circuit, or (b) a cycle with P_0 turning clockwise and P_1 counterclockwise, or (c) the opposite. Figure 12 illustrates the fact that these three cases are impossible.

□

5.5 Injectivity

For this paragraph, we return to the representation of annular triangulations in the plane with their marked face taken as unbounded face.

Let us define a *covering k-gonal forest* of an annular triangulation as a k-gonal forest obtained from the annular triangulation by detaching some edges from their endpoint in order to form leaves.

Lemma 5.14. *At most one covering k-gonal forest of T has only clockwise closure edges.*

Proof. The proof is essentially the same as for Lemma 3.9 but since the root edge is not necessarily incident to the unbounded face, some details must be adapted: in particular, the depth-first search traversal of both trees will be replaced by a contour walk around them. Observe that the border of the unbounded face of T contains at least one clockwise edge. Starting from the left hand side of this edge, perform a clockwise walk around both trees, that is, follow edge sides in clockwise direction (in particular the two sides of a pendant edge are visited successively, while for an inner edge the whole subtree on one extremity is visited between the visit of the first side and that of the second side).

The first time this walk differs between the two trees, the same contradiction appears as for Lemma 3.9, using the formulation of Property 2.3 in terms of interiors of cycles. □

This lemma proves the injectivity of the closure. The rest follows from equality of cardinalities.

Acknowledgments. We thank Nicolas Bonichon for sharing his invaluable knowledge of minimal Schnyder tree decompositions, and Cyril Gavoille for pointing out Lemma 7 of [7]. Many thanks also to Luca Castelli Aleardi and Eric Fusy for very useful conversations.

References

1. M.E. Agishtein and A.A. Migdal. Geometry of a two-dimensional quantum gravity: numerical study. *Nucl. Phys. B*, 350:690–728, 1991.
2. L. Alonso, J.-L. Remy, and R. Schott. A linear-time algorithm for the generation of trees. *Algorithmica*, 17(2):162–183, 1997.
3. J. Ambjørn, P Białas, Z. Burda, J. Jurkiewicz, and B. Petersson. Effective sampling of random surfaces by baby universe surgery. *Phys. Lett. B*, 325:337–346, 1994.
4. J. Ambjørn, B. Durhuus, and T. Jonsson. *Quantum geometry.* Cambridge Monographs on Mathematical Physics. Cambridge University Press, Cambridge, 1997.
5. C. Banderier, P. Flajolet, G. Schaeffer, and M. Soria. Planar maps and Airy phenomena. In *ICALP*, pages 388–402, 2000.
6. N. Bonichon. A bijection between realizers of maximal plane graphs and pairs of non-crossing dyck paths. In *FPSAC*, 2002.
7. N. Bonichon, C. Gavoille, and N. Hanusse. An information-theoretic upper bound of planar graphs using triangulation. In *STACS*, 2003.
8. N. Bonichon, B. Le Saëc, and M. Mosbah. Wagner's theorem on realizers. In *ICALP*, pages 1043–1053, 2002.
9. M. Bousquet-Melou and G. Schaeffer. The degree distribution in bipartite planar maps: applications to the Ising model. 2002, `arXiv:math.CO/0211070`.
10. E. Brehm. 3-orientations and Schnyder 3-tree-decompositions. Master's thesis, FB Mathematik und Informatik, Freie Universität Berlin, 2000.
11. W. Brown. Enumeration of triangulations of the disk. *Proc. London Math. Soc. (3)*, 14:746–768, 1964.
12. L. Castelli Aleardi, O. Devillers, and G. Schaeffer. Compact representation of triangulations. Submitted, 2004.
13. R. C.-N. Chuang, A. Garg, X. He, M.-Y. Kao, and H.-I Lu. Compact encodings of planar graphs via canonical orderings. In *ICALP*, pages 118–129, 1998.
14. H. de Fraysseix and P. Ossona de Mendez. Regular orientations, arboricity and augmentation. In *Graph Drawing*, 1995.
15. P. Duchon, P. Flajolet, G. Louchard, and G. Schaeffer. Random sampling from boltzmann principles. In *ICALP*, pages 501–513, 2002.
16. A. Dvoretzky and Th. Motzkin. The asymptotic density of certain sets of real numbers. *Duke Math. J.*, 14:315–321, 1947.
17. S. Felsner. Convex drawings of planar graphs and the order dimension of 3-polytopes. *Order*, 18:19–37, 2001.
18. S. Felsner. Lattice structures from planar graphs. *Electron. J. Combin.*, 11:Research Paper 15, 24 pp. (electronic), 2004.
19. P. Flajolet, P. Zimmermann, and B. Van Cutsem. A calculus for random generation of labelled combinatorial structures. *Theoret. Comput. Sci.*, 132(2):1–35, 1994.
20. E. Fusy, D. Poulalhon, and G. Schaeffer. Dissections and trees, with applications to optimal mesh encoding and to random sampling. In *Proceeding of ACM SODA*, 2005.
21. P.-M. Gandoin and O. Devillers. Progressive lossless compression of arbitrary simplicial complexes. *ACM Transactions on Graphics*, 21(3):372–379, 2002.
22. Z. Gao and J. Wang. Enumeration of rooted planar triangulations with respect to diagonal flips. *J. Combin. Theory Ser. A*, 88(2):276–296, 1999.
23. O. Gimenez and M. Noy. Asymptotic enumeration and limit law of labelled planar graphs. Preprint available on `arXiv.org`, 2005.

24. X. He, M.-Y. Kao, and H.-I Lu. Linear-time succinct encodings of planar graphs via canonical orderings. *SIAM J. on Discrete Mathematics*, 12(3):317–325, 1999.
25. X. He, M.-Y. Kao, and H.-I Lu. A fast general methodology for information-theoretically optimal encodings of graphs. *SIAM J. Comput*, 30(3):838–846, 2000.
26. G. Kant. Drawing planar graphs using the canonical ordering. *Algorithmica*, 16:4–32, 1996. (also *FOCS'92*).
27. D. King and J. Rossignac. Guaranteed 3.67v bit encoding of planar triangle graphs. In *CCCG*, 1999.
28. H.-I Lu. Linear-time compression of bounded-genus graphs into information-theoretically optimal number of bits. In *SODA*, pages 223–224, 2002.
29. J. I. Munro and V. Raman. Succint representation of balanced parantheses and static trees. *SIAM J. on Computing*, 31:762–776, 2001.
30. D. Osthus, H.J. Prömel, and A. Taraz. On random planar graphs, the number of planar graphs and their triangulations. *J. Comb. Theory, Ser. B*, 88(1):119–134, May 2003.
31. 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. Random generation of combinatorial objects and bijective combinatorics.
32. L. B. Richmond and N. C. Wormald. Almost all maps are asymmetric. *J. Combin. Theory Ser. B*, 63(1):1–7, 1995.
33. J. Rossignac. Edgebreaker: Connectivity compression for triangle meshes. *IEEE Transactions on Visualization and Computer Graphics*, 5(1):47–61, 1999.
34. G. Schaeffer. Bijective census and random generation of Eulerian planar maps with prescribed vertex degrees. *Electron. J. Combin.*, 4(1):# 20, 14 pp., 1997.
35. G. Schaeffer. *Conjugaison d'arbres et cartes combinatoires aléatoires.* PhD thesis, Université Bordeaux I, 1998.
36. G. Schaeffer. Random sampling of large planar maps and convex polyhedra. In *STOC*, pages 760–769, 1999.
37. W. Schnyder. Embedding planar graphs on the grid. In *SODA*, pages 138–148, 1990.
38. C. Touma and C. Gotsman. Triangle mesh compression. In *Proceedings of Graphics Interface*, Vancouver, 1998.
39. W. T. Tutte. A census of planar triangulations. *Canad. J. Math.*, 14:21–38, 1962.
40. D. B. Wilson. Annotated bibliography of perfectly random sampling with markov chains. `http://dimacs.rutgers.edu/~dbwilson/exact`.