



Orienting triangulations

Boris Albar, Daniel Gonçalves, Kolja Knauer

► To cite this version:

Boris Albar, Daniel Gonçalves, Kolja Knauer. Orienting triangulations. EuroCG: European Workshop on Computational Geometry, Mar 2015, Ljubljana, Slovenia. lirmm-01264309

HAL Id: lirmm-01264309

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-01264309>

Submitted on 25 Jun 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Orienting triangulations^{*}

Boris Albar¹, Daniel Gonçalves², and Kolja Knauer¹

¹ LIF, UMR 7279, CNRS & Aix-Marseille Université

² LIRMM, UMR 5506, CNRS & Université de Montpellier

Abstract. We prove that any triangulation of a surface different from the sphere and the projective plane admits an orientation without sinks such that every vertex has outdegree divisible by three. This confirms a conjecture of Barát and Thomassen and is a step towards a generalization of Schnyder woods to higher genus surfaces.

1 Introduction

The notation and results we use for graphs and surfaces can be found in [10]. We start with some basic definitions:

A *map* (or *2-cell embedding*) of a multigraph into a surface, is an embedding such that deleting the graph from the surface leaves a collection of open disks, called the *faces* of the map. A *triangulation* is a map of a simple graph (i.e. without loops or multiple edges) where every face is triangular (i.e. incident to three edges). A fundamental result in the topology of surfaces is that every surface admits a map. The (*orientable*) *genus* of a map on an orientable surface is $\frac{1}{2}(2 - n + m - f)$ and the (*non-orientable*) *genus* of a map on a non-orientable surface is $2 - n + m - f$, where n, m, f denote the number of vertices, edges, and faces of the map, respectively. The *Euler genus* of a map is $2 - n + m - f$, i.e., the non-orientable genus or twice the orientable genus. All the maps on a fixed surface have the same genus, which justifies to define the (*Euler*) *genus of a surface* as the (Euler) genus of any of the maps it admits. In [1] Barát and Thomassen conjectured the following:

Conjecture 1. Let T be a triangulation of a surface of Euler genus $k \geq 2$. Then T has an orientation such that each outdegree is at least 3, and divisible by 3.

One easily computes that the number of edges m of a triangulation T of a surface of Euler genus k is $3n - 6 + 3k$. So while triangulations of Euler genus less than 2 simply have too few edges to satisfy the conjecture, in [1] the conjecture is proved for the case $k = 2$, i.e., the torus and the Klein bottle. Moreover, they show that any triangulation T of a surface has an orientation such that each outdegree is divisible by 3, i.e, in order to prove the full conjecture they miss the property that there are no sinks.

^{*} This work was supported by the project EGOS, ANR-12-JS02-002-01 and partially by PEPS grant EROS

Barát and Thomassen’s conjecture was originally motivated in the context of claw-decompositions of graphs, since given an orientation with the claimed properties the outgoing edges of each vertex can be divided into claws, such that every vertex is the center of at least one claw.

Another motivation for this conjecture is, that it can be seen as a step towards the generalization of planar Schnyder woods to higher genus surfaces. A *Schnyder wood* [11] of a planar triangulation is an orientation and a $\{0, 1, 2\}$ -coloring of the *inner* edges satisfying the following *local rule* on every *inner* vertex v : going counterclockwise around v one successively crosses an outgoing 0-arc, possibly some incoming 2-arcs, an outgoing 1-arc, possibly some incoming 0-arcs, an outgoing 2-arc, and possibly some incoming 1-arcs until coming back to the outgoing 0-arc.

Schnyder woods are one of the main tools in the area of planar graph representations and Graph Drawing. They provide a machinery to construct space-efficient straight-line drawings [12, 6], representations by touching \mathbb{T} shapes [5], they yield a characterization of planar graphs via the dimension of their vertex-edge incidence poset [11, 6], and are used to encode triangulations efficiently [3]. In particular, the local rule implies that every Schnyder wood gives an orientation of the inner edges such that every inner vertex has outdegree 3 and the outer vertices are sinks with respect to inner edges. Indeed, this is a one-to-one correspondance between Schnyder woods and orientations of this kind. As a consequence, the set of Schnyder woods of a planar triangulation inherits a natural distributive lattice structure, which in particular provides any triangulation with a unique *minimal* Schnyder wood [7]. These unique representatives are an important tool in proofs and lie at the heart of many enumerative results, see for instance [2].

When generalizing Schnyder woods to higher genus one has to choose which of the properties of planar Schnyder woods are desired to be carried over to the more general situation. Examples are: the efficient encoding of triangulations on arbitrary surfaces [4] and the relation to orthogonal surfaces and small grid drawings for toroidal triangulations [9], which lead to different definitions of generalized Schnyder woods. In [9], the generalized Schnyder woods indeed satisfy the local rule with respect to *all* edges and vertices of a toroidal triangulation and henceforth lead to orientations having outdegree 3 at every vertex. An interesting open problem is to generalize the local rule to triangulations with higher Euler genus in such a way that for some vertices the sequence mentioned in the local rule occurs several times around the vertex. Here, the mere existence of such objects is an open question. Clearly, such a generalized Schnyder wood would yield an orientation as claimed by the conjecture. Thus, proving the conjecture of Barát and Thomassen is a first step into that direction.

2 Preliminaries

A map M on a surface \mathbb{S} is characterized by a triple $(V(M), E(M), F(M))$, formed by the vertex, edge and face sets of M . In the following we will restrict

to triangulations $T = (V(T), E(T), F(T))$, i.e. the pair $(V(T), E(T))$ is a simple embedded graph such that every face is incident to exactly three edges.

A *submap* M' of T , is a triplet (V', E', F') where $V' \subseteq V(T)$, $E' \subseteq E(T)$, $F' \subseteq F(T)$, and such that:

- $uv \in E'$ implies $\{u, v\} \subseteq V'$, and
- $f \in F'$ implies $e' \in E'$ for any edge e incident to f .

The *boundary* $\partial M'$ of a submap $M' = (V', E', F')$ is the set of edges in E' that are incident to at most one face in M' .

For any vertex v of T its *surrounding* is the circular sequence of edges and faces successively met while going around v . This sequence as no particular direction as T can be non-orientable. In a submap M' of T a (*boundary*) *angle* at vertex v is a sub-sequence $(e_0, f_1, e_1, \dots, f_t, e_t)$ of its surrounding such that the edges e_0 and e_t are the only elements of this sequence belonging to M' . Those edges are the *sides* of this angle. This angle can be denoted by $\widehat{e_0 v e_t}$ or simply by \widehat{v} . It can occur that $e_0 = e_t$. Consider for example a submap consisting of a single edge. Let us mention, that this definition could be modified in order to include the angle around a vertex with respect to a submap without edges. Since we will not consider this situation we prefer avoiding further technicalities.

Note that an edge is in $\partial M'$ if and only if it is a side of (at least) one angle of M' . Actually, the notion of angles endows the boundary $\partial M'$ of M' with some further structure. As each angle has two sides (possibly two occurrences of the same edge) and as each occurrence of an edge of $\partial M'$ is a side for two angles, one can define the *boundary sequence* of M' , that is a collection of circular sequences, alternating between angles and edges, $(\widehat{a_0}, e_0, \widehat{a_1}, e_1, \dots, \widehat{a_t}, e_t)$ (sometimes simply denoted by $(\widehat{a_0}, \widehat{a_1}, \dots, \widehat{a_t})$ or (e_0, e_1, \dots, e_t)), where e_i is the common edge of $\widehat{a_i}$ and $\widehat{a_{i+1}}$. Note that an edge e may appear twice in the boundary sequence, e.g. if e is a bridge of M' . Thus, if necessary we will refer to a specific *occurrence* of e in $\partial M'$. For simplicity, we denote the boundary sequence of M' by $\partial M'$. This naturally leads to the notion of *consecutive angles*. Note that two angles $(e_0, f_1, e_1, \dots, f_t, e_t)$ and $(e'_0, f'_1, e'_1, \dots, f'_t, e'_t)$ are *consecutive* on the boundary sequence if $e_t = e'_0$ and $f_t = f'_1$.

In the following, a *disk* is a submap M' of T if it is homeomorphic to a (closed) topological disk. Furthermore, a disk is a *k-disk* if its boundary is a cycle with k edges. A 3-disk is called *trivial* if it contains only one face. A disk is called *chordless* if its outer vertices (i.e. on its boundary) induce a graph that is a (chordless) cycle. A cycle is *contractible* if it is the boundary of a disk otherwise it is called *non-contractible*.

Given a triangulation T and a set of vertices $X \subseteq V(T)$, the *induced submap* $T[X]$ is the maximal submap with vertex set X . In other words this submap has edge set $\{uv \in E(T) \mid u \in X \text{ and } v \in X\}$, and face set $\{uvw \in F(T) \mid u \in X, v \in X, \text{ and } w \in X\}$.

Given an induced submap $M' = T[X]$ of a triangulation T , and any occurrence of an edge ab in $\partial M'$ (corresponding to angles \widehat{a} and \widehat{b}) there exists a unique vertex c such that there is a face abc in $F(T) \setminus F(M')$ that belongs to both angles \widehat{a} and \widehat{b} . For any such vertex c (and $ab \in \partial M'$) we define the

operation of *stacking* c on M' , as adding c to X , i.e., going from $M' = T[X]$ to $M'' = T[X + c]$. In such stacking the *neighborhood* of c in M' is the graph with vertices x such that $cx \in E(M'')$ and edges xy such that $cxy \in F(M'')$. As T is simple, note that this neighborhood is either a cycle or a union of paths, one of which with at least one edge (the edge allowing the stacking), and let us call them the *neighboring cycle* and the *neighboring paths* of c in M' , respectively. See Figure 1 for an illustration.

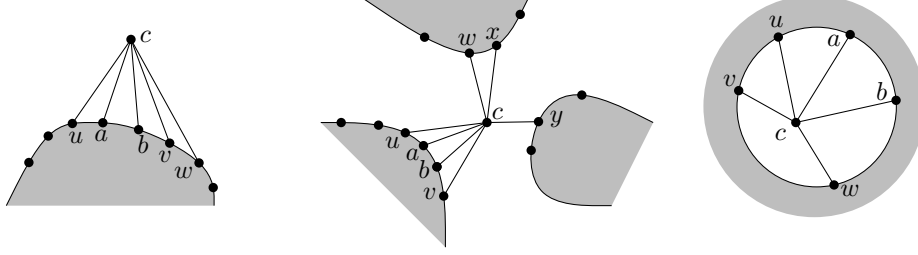


Fig. 1. Different scenarios of stacking c on M' . Left: one neighboring path $P_1 = (u, a, b, v, w)$. Middle: three neighboring paths $P_1 = (u, a, b, v)$, $P_2 = (w, x)$, $P_3 = (y)$. Right: A boundary cycle $C = (u, v, w, b, a)$.

3 Proof of Conjecture 1

Let us consider for contradiction a minimal counterexample T . Note that T does not contain any non-trivial 3-disk D . Otherwise we would remove the interior of D and would replace it by a face. By minimality of T , this new triangulation would admit an orientation such that every vertex has non-zero outdegree divisible by 3. As D is a planar triangulation, there exists an orientation of its interior edges so that inner and outer vertices have out-degree 3 and 0, respectively. This is the case for orientations induced by a Schnyder wood on these triangulations [11]. Then the union of these two orientations would give us an orientation of T with non-zero outdegrees divisible by three. Let us now proceed by providing an outline of the proof.

3.1 Outline

We first prove that one can partition the edges of the triangulation T into the following graphs:

- The *initial graph* I , which is an induced submap containing a non-contractible cycle. Furthermore, I contains an edge uv such that the map $I \setminus uv$ is a disk \tilde{D} whose underlying graph is a maximal outerplanar graph with only two degree two vertices, u and v . See Figure 2 for an illustration.

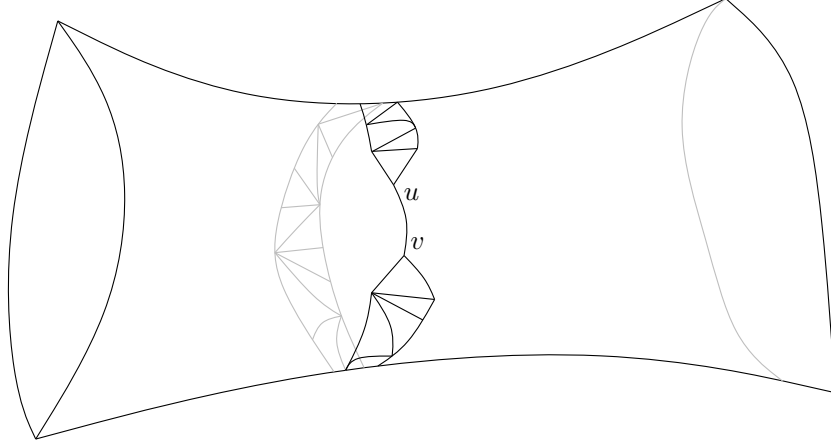


Fig. 2. Example of a submap I .

- The *correction graph* B (with blue edges in the figures), which is oriented acyclically in such a way that each vertex of $V(T) \setminus V(I)$ has outdegree 2, while the other vertices have outdegree 0,
- The *last correction path* G (with green edges in the figures), which is a $\{u, v\}$ -path.
- The *non-zero graph* R (with red edges in the figures), which is oriented in such a way that all vertices in $(V(T) \setminus V(G)) \cup \{u, v\}$ have out-degree at least 1.

The existence of such graph I is proven in Section 3.2, then in Section 3.3 we prove the existence of graphs B , G and R (with the mentioned orientations). To do the latter we start from I and we incrementally conquer the whole triangulation T by stacking the vertices one by one (this procedure is inspired by [4]).

Finally, the edges of I , B and G are (re)oriented, to obtain the desired orientation. The orientation of edges in R does not change, as they ensure that many vertices (all vertices of T except the interior vertices of the path G) have non-zero outdegree. The $\{u, v\}$ -path G is either oriented from u to v or from v to u , but this will be decided later. However in both cases its interior vertices are ensured to have non-zero outdegree. Hence all vertices are ensured to have non-zero outdegree and it remains to prove that they have outdegree divisible by 3.

We start in Section 3.4 by reorienting the B -arcs in order to ensure that vertices of $V(T) \setminus V(I)$ have outdegree divisible by 3 (this part is inspired by the proof of Theorem 4.5 in [1]). In the last step, in Section 3.5, we choose the orientation of the $\{u, v\}$ -path G , and we orient I in order to achieve the desired orientation.

3.2 Existence of I

To prove the existence of I , we first need the following lemma.

Lemma 1. *Any triangulation T with Euler genus at least 2, has an induced submap I obtained from a disk D by stacking a vertex v , such that for any two neighbors a, b of v belonging to distinct neighboring paths (of v w.r.t. D), every cycle C in I going through edges av and vb is non-contractible.*

Proof. Any face of T is an induced disk. Consider a maximal induced disk D of T . For any edge xy of ∂D , stack a vertex v on xy . Let us denote by I the map obtained by stacking v on D . As T has Euler genus at least 2 the neighborhood of v is not a cycle. Also, as D is maximal, v has at least two neighboring paths. Assume for contradiction, that there is a contractible cycle C of I going through av , vb (where a and b belong to distinct neighboring paths of v w.r.t. D) and through some $\{a, b\}$ -path P of ∂D . Denote by D' the disk bounded by C and note that (as I is induced) D' contains vertices not in I . As $\partial D'$ intersects D on a path it is clear that $V(D) \cup V(D')$ induces disk with v on its boundary. Furthermore, as only two neighbors of v in this disk are on the border, we have that $(V(D) \cup V(D')) \setminus \{v\}$ also induces a disk. This disk is larger than D , contradicting its maximality. \square

Lemma 2. *Any triangulation T' with Euler genus at least 2, has an induced submap I containing a non-contractible cycle, and an edge uv such that $I \setminus uv$ is a disk \tilde{D} , and for each of the two $\{u, v\}$ -paths of $\partial \tilde{D}$, all its interior vertices have a neighbor in the interior of the other $\{u, v\}$ -path.*

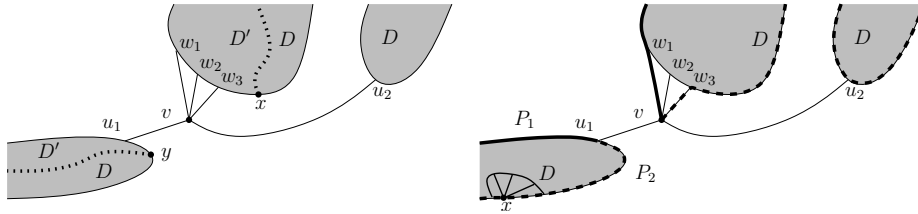


Fig. 3. The situation in Claim 1 (left) and Claim 2 (right) in the proof of Lemma 2.

Proof. Among the induced subgraphs of T' that satisfy Lemma 1 let I be a minimal one. Let v and D be the vertex of I and the disk $I \setminus \{v\}$ described in Lemma 1, respectively. As v is stacked on D let us denote by (w_1, \dots, w_s) , with $s \geq 2$, some neighboring path of v , and let us denote by u_1, \dots, u_t , with $t \geq 1$, the other neighbors of v in D . Finally, let us denote by \tilde{D} the disk obtained from D by adding vertex v , edges vw_i for $1 \leq i \leq s$, and faces vw_iw_{i+1} for $1 \leq i < s$. The minimality of I implies all the needed properties:

Claim 1 ∂D induces no chord xy inside D such that some $\{x, y\}$ -path of ∂D contains both an edge $w_i w_{i+1}$, for some $1 \leq i < s$, and a vertex u_j , for some $1 \leq j \leq t$.

Proof. If such chord xy exists, let $D' \subsetneq D$ be the disk with boundary in $\partial D + xy$ which contains both $w_i w_{i+1}$ and u_j . Then the graph induced by $V(D') \cup \{v\}$ contradicts the minimality of I . See the left of Figure 3. \square

This implies that $\partial \tilde{D}$ has no chord at u_j , for all $1 \leq j \leq t$.

Claim 2 For all $1 \leq j \leq t$, every interior vertex x of a $\{v, u_j\}$ -path of $\partial \tilde{D}$ is adjacent to an interior vertex of the other $\{v, u_j\}$ -path.

Proof. Let P_1 and P_2 be the $\{v, u_j\}$ -path of $\partial \tilde{D}$ containing w_1 and w_s , respectively. Assume for contradiction, there exists an inner vertex x in P_1 having no neighbor in the interior of P_2 . By Claim 1 this implies that D (the disk induced by $V(I) \setminus \{v\}$) has no chord at x . Thus the map induced by $V(D) \setminus \{x\}$ is a disk containing P_2 on its border, hence containing the vertex u_j and the edge $w_{s-1} w_s$. Hence the map induced by $V(I) \setminus \{x\}$ contradicts the minimality of I . See the right of Figure 3. \square

As $\partial \tilde{D}$ has no chord at u_j , for all $1 \leq j \leq t$, this implies that $t = 1$. This concludes the proof of the lemma. \square

Consider now our counterexample T , and let I , uv and $\tilde{D} (= I \setminus uv)$ be an induced submap, an edge and a disk, verifying Lemma 2. In the beginning of the section we have seen that by minimality, T does not contain non-trivial 3-disks. Hence by the properties of I , if \tilde{D} would contain an inner vertex, this vertex would be in a chordless 4-disk of \tilde{D} . By the following lemma this is not possible, hence \tilde{D} is a maximal outerplanar graph. Finally the adjacency property between vertices of $\partial \tilde{D} \setminus \{u, v\}$ imply that u and v are the only degree two vertices of \tilde{D} .

Lemma 3. The submap \tilde{D} does not contain chordless 4-disks.

Proof. If \tilde{D} would contain such a disk D_4 , with boundary (v_1, v_2, v_3, v_4) , we would remove the interior of D_4 and we would add one of the two possible diagonals, say $v_2 v_4$ (if $v_2 v_4$ are not uv 's ends), and the corresponding two triangular faces, $v_1 v_2 v_4$ and $v_2 v_3 v_4$. The obtained map T' is defined on the same surface as T and is smaller. Furthermore as I is an induced submap without non-trivial 3-disk and as $v_2 v_4 \neq uv$, there is no edge $v_2 v_4$ in T . Hence T' is simple and it is a triangulation. Now by minimality of T , this new triangulation T' has an orientation such that every vertex has non-zero outdegree divisible by 3. Let us suppose without loss of generality that in this orientation the edge $v_2 v_4$ is oriented from v_2 to v_4 .

Using the fact that for any planar triangulation, there exists an orientation of the interior edges such that inner and outer vertices have out-degree 3 and 0 [11] respectively, one can orient the inner edges of D_4 in such a way that inner vertices, vertex v_2 , and vertices v_1 , v_3 and v_4 have out-degree 3, 1 and 0,

respectively. For this consider the orientation given by a Schnyder wood of the triangulation $D_4 + v_1v_3$ (with outer face $v_1v_3v_4$) and notice that the edges v_2v_1 and v_2v_3 are necessarily oriented from v_2 to v_1 and v_3 respectively (as v_1, v_3 and v_4 have outdegree 0).

Then the union of these orientations, of $T' \setminus v_2v_4$ and of D_4 's inner edges, would give us an orientation of T with non-zero outdegrees divisible by three. \square

3.3 Existence of B , G , and R

As mentioned in the outline, we will start from I and we incrementally explore the whole triangulation T by stacking the vertices one by one. At each step, we will assign the newly explored edges to B , G or R , and we will orient those assigned to B or R . At each step the *explored region* is a submap of T induced by some vertex set X . Such explored region is denoted by $T[X]$ and its boundary $\partial T[X]$. The connected pieces of the surface obtained after removing $T[X]$ are called the *unexplored regions*, and if one of them is homeomorphic to an open disk it is called an *unexplored disk*. Given an unexplored disk D (by abuse of notation) we denote by ∂D the cycle of $\partial T[X]$ bordering D . During the exploration we maintain the following invariants:

- (I) The graphs I , B , G , and R partition the edges of $T[X]$.
- (II) All interior vertices of $T[X]$ (i.e. in $X \setminus V(\partial T[X])$) have at least one outgoing R -arc, or two incident G -edges. Furthermore G either is an $\{u, v\}$ -path, or is the union of two vertex disjoint paths G_u and G_v , going from u to u_* , and from v to v_* , respectively, for some vertices u_* and v_* on $\partial T[X]$.

Here the vertices u_* and v_* may coincide with vertices u and v , respectively, if G_u or G_v is a trivial path with only one vertex.

- (III) The graph B is acyclically oriented in such a way that the vertices of I have outdegree 0, while the other vertices of $T[X]$ have outdegree 2.

Furthermore, to help us in properly finishing the construction of the graphs B , G and R in the further steps, we introduce the notion of *requests* on the angles of $\partial T[X]$. There are two types of requests, G -requests and R -request. An angle is allowed to have at most one request, and an angle having no request is called *free*. Informally, a G -request (resp. an R -request) for an angle \hat{a} means that in a further step an edge inside this angle will be added in G (resp. in R and oriented from a to the other end). In the figures, a G -request (resp. an R -request) is depicted by a green (resp. red) arrow.

- (IV) Every vertex of $(\partial T[X] \setminus \{u_*, v_*\}) \cup \{u, v\}$ having (still) no outgoing R -arc, has an incident angle with an R -request.
- (V) If G is not a $\{u, v\}$ -path (yet), the vertices u_* and v_* (at the end of G_u and G_v , respectively), have one incident angle each, say $\widehat{u_*}$ and $\widehat{v_*}$, that are consecutive on $\partial T[X]$, and that have a G -request. Furthermore, there are no other G -requests.

- (VI) If there is an unexplored disk D' , then there are at least three free angles (of $\partial T[X]$) around D' .

Before starting this exploration, let us observe that if these invariants are maintained until the end of the exploration, we obtain the desired partition of the edges. Note that at the end of the exploration $T[X]$ has no border, hence no requests, and by (V) G is thus an $\{u, v\}$ -path. As u and v have degree 1 in G , by (II) every vertex in $(V(T) \setminus V(G)) \cup \{u, v\}$ has out-degree at least 1 in R . Finally, by (III) B is oriented acyclically in such a way that each vertex of $V(T) \setminus V(I)$ has outdegree 2, while the other vertices have outdegree 0. We can now proceed to the exploration itself.

This exploration starts with $T[X] = I$. In this case as all the edges of $T[X]$ are in I and as there are no interior vertices yet, (I), (II) and (III) are trivially satisfied. Since I contains a non-contractible cycle and since the Euler genus of T is at least 2 there is no unexplored disk, hence (VI) is satisfied. Since uv appears twice in $\partial T[X]$, the vertices u, v appear twice consecutively in $\partial T[X]$. To achieve (V), choose the angles of one consecutive appearance of u, v as G -requests. To achieve (IV), all the other angles are assigned R -requests. See Figure 4 for an illustration.

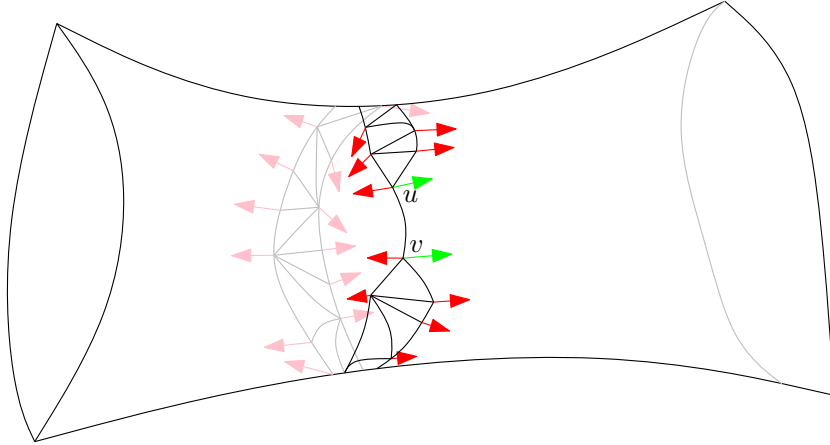


Fig. 4. Assigning requests to I in order to satisfy the invariants.

For the rest of the construction in each step we enlarge the explored submap $T[X]$ by stacking a vertex x on $T[X]$. The vertex x is chosen according to the following rules:

- (i) If there is only one edge in the neighborhood of x in $T[X]$, this edge is not $\{u_*, v_*\}$.

- (ii) If x belongs to an unexplored disk D , either x is adjacent to all the vertices of ∂D or x has exactly one neighboring path P on ∂D such that P does not contain all the free angles of ∂D .
- (iii) In the case x does not belong to an unexplored disk, if possible we choose x such that no unexplored disk is created. Furthermore, if unexplored disks are created we choose x in order to minimize the total surface of these unexplored disks (measured by the number of faces in these regions).

Let us explain why choosing such a vertex x is always possible. If there is an unexplored disk D , let us choose x inside D . If there is a vertex adjacent to all the vertices of ∂D we are fine ((i) follows). Otherwise, one can show that there are at least two vertices inside D , say x_1 and x_2 , having exactly one neighboring path $P \neq (u_*, v_*)$ on ∂D , say P_1 and P_2 respectively. These two paths intersect on at most two vertices, so one of them, say P_1 , avoids one of the (at least) three free angles around D . In that case choosing x_1 as the next vertex to stack fulfills (i) and (ii). Now if there is no unexplored disk, as there are at least three edges on $\partial T[X]$ there are candidates fulfilling (i). As (iii) is not constraining we are done.

In the following we show how to extend B, G, R on the newly introduced edges and how to deal with the newly created angles to maintain all invariants valid. We will describe the construction and we will check the validity of invariants only for the non-trivial ones. We distinguish cases according to the topology of the unexplored region containing x .

1) The vertex x is contained in an unexplored disk D and has a neighboring cycle. By (VI) the unexplored disk containing x has at least 3 free angles. We orient the corresponding edges from x to its neighbors, put two into B and the rest into R . All non-free angles satisfy their request with the edge incident to x . See Figure 5 for an illustration.

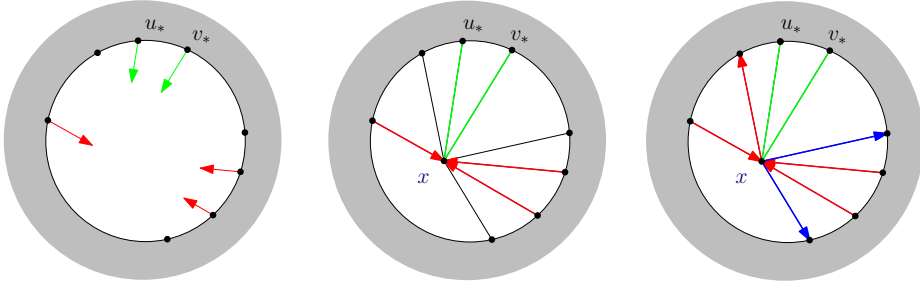


Fig. 5. Case where x is in an unexplored disk D and has a neighboring cycle.

We have assigned all the newly explored edges, hence (I) remains valid. As (IV) and (V) were valid in $T[X]$, all the neighbors of x (i.e. the vertices around

D) have now (in $T[X + x]$) an outgoing R -arc or two incident G -edges. The vertex x also does, hence (II) is valid. In the acyclic graph B , adding the vertex x with only outgoing B -arcs cannot create any circuit, hence (III) remains valid. As in this case, as $\partial T[X + x]$ is included in $\partial T[X]$, (IV) remains valid. If u_* and v_* were around D in $T[X]$, the two parts of G are now connected by the adjunction of xu_* and xv_* in G . Otherwise, G was already an $\{u, v\}$ -path, or u_* and v_* were elsewhere in $\partial T[X]$ fulfilling (V). Hence in any case (V) holds. Finally, as no unexplored disk has been created and as the requests around existing unexplored disks have not changed, (VI) remains valid.

For the remaining cases we introduce some further notation. Given a neighboring path $P = (p_1, \dots, p_s)$ of x , with corresponding angles $\widehat{p}_1, \dots, \widehat{p}_s$, the *inner angles* are the angles \widehat{p}_i with $1 < i < s$. The other ones are the *outer angles*. An inner angle with an R - or G -requests, has to satisfy its constraint (this cannot be further delayed). Hence for any inner angle \widehat{p}_i with a G -request (resp. an R -request) we add the edge xp_i to G (resp. to R oriented towards x). This is a preprocessing step valid for both the remaining two cases.

2) The unexplored region containing x is not a disk. For simplicity assume, that there are no free angles. Otherwise we assign an R -request to all these angles. Here after the preprocessing step described above, there is an intermediate step 2.1) and a final step 2.2). See Figure 6 for an illustration of how this case is handled.

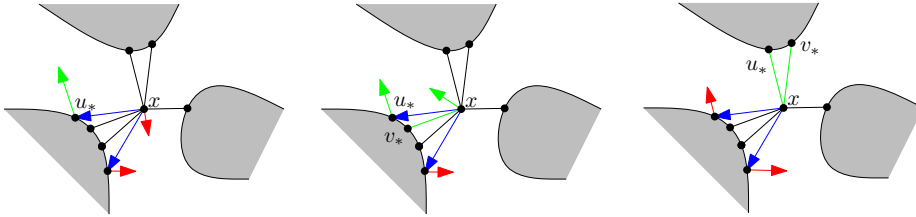


Fig. 6. Case where x is not in an unexplored disk. (left) One G -request is on a neighboring path of x . (center) One G -request is on an outer angle and one is on an inner angle. (right) Both of the G -requests are on outer angles.

2.1) The intermediate step. This step depends on the position of the G -requests, if any.

If there is no G -request on the neighboring paths of x , then we assign an R -request to some angle \widehat{x} incident to x .

If only one G -request, say on $\widehat{u_*}$, is on a neighboring path of x , then by (V) $\widehat{v_*}$ is next to it, hence $\widehat{u_*}$ is an end of this neighboring path. Here the new angle at u_* (inside the former angle $\widehat{u_*}$) that is created by stacking x inherits

\widehat{u}_* 's G -request. If two angles are created inside the former angle \widehat{u}_* , that is if u_* is alone in its neighboring path, we choose the angle next to v_* in order to fulfill (V). Then we assign an R -request to some angle \widehat{x} incident to x .

If one G -request, say \widehat{u}_* , is on an outer angle and the other one, \widehat{v}_* , is on an inner one, we have added the edge v_*x to G in the preprocessing. Here the new angle at u_* inherits \widehat{u}_* 's G -request and the next angle on $\partial T[X + x]$, that is incident to x gets a G -request too.

If both G -requests are on inner angles, the edges v_*x and u_*x have been added to G in the preprocessing. Hence x has already two incident G -edges and does not need any request around. We thus leave all angles incident to x free.

If both of the G -requests are on outer angles, then by (i) x has one neighboring path of length one, (u_*, v_*) , and at least one other neighboring path of length at least one. In that case, we add edges v_*x and u_*x to G and we leave the new angles at u_* and v_* , as well as all angles incident to x , free.

2.2) The final step. We now assign two outgoing B -arcs to x , depending on the G -requests. If there is an outer angle \widehat{u}_* (in $T[X + x]$) with a G -request add the arc xu_* directed towards u_* to B . The remaining one or two needed B -arcs are chosen arbitrarily among the edges from x to outer vertices. All other edges, between x and outer vertices will be put into R and directed towards x , and the corresponding angles will be left free. Note that among the newly created outer angles and the angles associated to x there are at most 3 requests: two at the angles receiving a B -arc from x and one at an angle incident to x .

If adding x creates an unexplored disk D' , we still have to argue, that (VI) is satisfied with respect to D' . We make use of the following:

Claim 3 *For any unexplored disk D' created by stacking a vertex x on $T[X]$, the vertex x appears several times on the boundary of D' .*

Proof. Suppose we create an unexplored disk D' such that x appears only once on its boundary. Assume x is chosen such that the number of faces in D' is minimized. Since there are no non-trivial 3-disks, the boundary of D' is of length at least 4. Therefore D' contains an unexplored vertex x' which could have been stacked on a subpath of $\partial D' \setminus x$. Furthermore, x' can be chosen such that the path does not only contain the G -requests. Hence, stacking x' would either not create any unexplored disk, or would create some included in D' , hence smaller. Both cases contradict the choice of x with respect to (iii). \square

This claim and the fact that T is simple imply that there are at least 6 angles on the boundary of D' incident to outer vertices of the neighborings paths of x (4 of them) or incident to x (2 of them). As argued above at most 3 of these angles have a request. Thus, there are at least 3 free angles on the boundary of D' and (VI) is satisfied.

3) The unexplored region containing x is a disk, but x 's neighborhood is not a cycle. By (ii) the vertex x has only one neighboring path. Let us denote

this path by $P = (p_1, \dots, p_s)$ for some $s \geq 2$ and $\widehat{p}_1, \dots, \widehat{p}_s$ the corresponding angles. Denote by t the number of free angles on P .

We start with the preprocessing described above, that deals with non-free interior angles (by fulfilling the requests). To fulfill (VI) we have to maintain the number of free angles in this unexplored disk above three. Since by (ii) there is at least one free angle not on P , to achieve this we need to have at least $\min\{t, 2\}$ free angles among the new angles \widehat{p}_1 , \widehat{x} , and \widehat{p}_s .

To achieve that we need to exploit free angles as follows. For any free angle \widehat{p}_i (inner or not), the edge $x p_i$ is added either to B or to R , in both cases oriented towards p_i . Among these t angles, $\min\{t, 2\}$ lead to a B -arc, and $\max\{0, t - 2\}$ lead to an R -arc. It remains to deal with the (at least $2 - t$) angles that are neither inner nor free. We proceed by distinguishing cases according to the position of G -requests.

If there is no G -request on P , we proceed as follows. Let us first deal with the new angle \widehat{x} . If $t \leq 2$, the vertex x has no outgoing R -arc and we hence assign an R -request to the angle \widehat{x} . Otherwise (i.e. if $t \geq 3$) the vertex x has an outgoing R -arc, we hence leave \widehat{x} free. Then we use $\max\{0, 2 - t\}$ of the non-free outer angles to add B -arcs leaving x . We satisfy the possibly remaining non-free outer angles (that are $\min\{2, t\}$), by adding R -arcs towards x , and leave their new incident angle free. If $t \leq 2$ (resp. $t \geq 3$), there are hence $\min\{2, t\} = t$ (resp. $1 + \min\{2, t\} = 3$) free angles among the new angles \widehat{p}_1 , \widehat{x} , and \widehat{p}_s . We hence have the expected (at least) $\min\{t, 2\}$ free angles.

If only one G -request (say on \widehat{u}_*) is on P , then \widehat{u}_* is an end of P , say $p_1 = u_*$ (see Figure 7). Here the new angle at u_* inherits \widehat{u}_* 's G -request, and we add the edge $x p_1$ in B if $t \leq 1$, or in R otherwise (if $t \geq 2$). In both cases $x p_1$ is oriented towards p_1 . Hence, if $t \leq 1$ we assign an R -request to angle \widehat{x} and otherwise we leave \widehat{x} free. If $t = 0$ then \widehat{p}_s is not free, then as it cannot have a G -request, \widehat{p}_s has an R -request. In that case we add $x p_s$ in B oriented from x to p_s and the new angle \widehat{p}_s inherits the R -request. If $t \geq 1$, we satisfy the R -request of \widehat{p}_s (if it has one) with edge $x p_s$. In any case, \widehat{p}_s having a request or not in $T[X]$, the new angle \widehat{p}_s is left free. Hence if $t \geq 2$ the angle \widehat{x} is free, and if $t \geq 1$ the angle \widehat{p}_s is free. We hence have the expected (at least) $\min\{t, 2\}$ free angles.

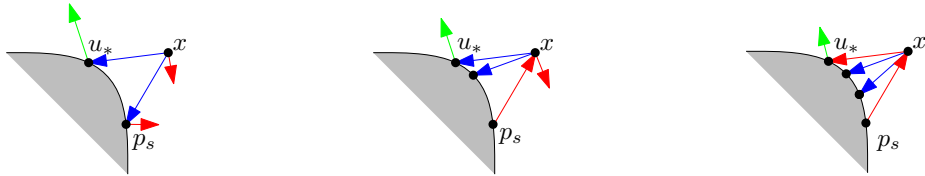


Fig. 7. Case where there is only one G -request (on \widehat{u}_*) and where \widehat{p}_s has an R -request. The 3 subcases from left to right correspond to $t = 0$, $t = 1$, and $t = 2$.

If one G -request say \widehat{u}_* is on an outer angle and the other one \widehat{v}_* on an inner one, say $u_* = p_1$ and $v_* = p_2$ with $s > 2$, we have added the edge p_2x to G (see Figure 8). Around p_1 , if $t \leq 1$ we assign the new angles \widehat{p}_1 and \widehat{x} a G -request, and we add the edge xp_1 in B oriented from x to p_1 . Otherwise (i.e. $t \geq 2$) we add the edge p_1x to G , and we leave both new angles \widehat{p}_1 and \widehat{x} as free. Around p_s , if $t = 0$ (hence \widehat{p}_s has an R -request) we add xp_s in B oriented from x to p_s , and the new angle \widehat{p}_s keeps its R -request. Otherwise (i.e. $t \geq 1$), if \widehat{p}_s has an R -request we add xp_s in R and orient it from p_s to x , and in any case (\widehat{p}_s having an R -request or not) we leave the new angle \widehat{p}_s as free. Hence if $t \geq 1$ the angle \widehat{p}_s is free, and if $t \geq 2$ both \widehat{p}_1 and \widehat{x} are free. We hence have the expected (at least) $\min\{t, 2\}$ free angles.

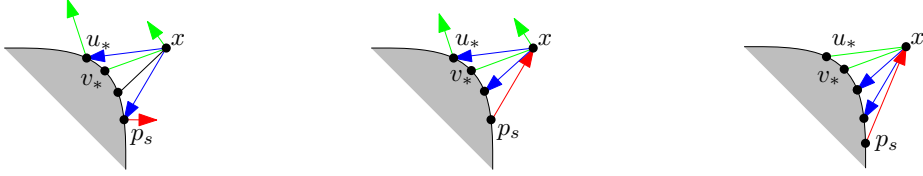


Fig. 8. Case where there is one G -request on an outer angle, and one in an inner angle, and where \widehat{p}_s has an R -request. The 3 subcases from left to right correspond to $t = 0$, $t = 1$, and $t = 2$.

If both G -requests are on inner angles, edges v_*x and u_*x have been added to G . Now x is an inner vertex of G , and we thus leave \widehat{x} free. Then we use $\max\{0, 2 - t\}$ of the outer angles for B -arcs from x , and the remaining non-free outer angles have their R -requests satisfied, and are left free. In any case, $\min\{2, t\}$ of the outer angles are free (as is the angle \widehat{x}).

Finally by definition of stacking P , x 's unique neighboring path is distinct from (u_*, v_*) and hence all cases have been addressed.

3.4 Reorienting B

Given a partial orientation O of T we define the *demand* of a vertex v as $\text{dem}_O(v) := -\delta_O^+(v) \bmod 3$, where $\delta_O^+(v)$ denotes the outdegree of v with respect to O . We want to find an orientation of T with all demands 0.

Recall we will not modify the orientation on R , which guarantees that all vertices in $(V(T) \setminus V(G)) \cup \{u, v\}$ have non-zero outdegrees. Furthermore, as G will be oriented either entirely forward or backwards (this will be chosen later), all its interior vertices will have non-zero outdegrees. Hence every vertex of $T[X]$ has non-zero outdegree. Suppose that G is entirely oriented forward.

Now we linearly order vertices in $V(T) \setminus V(I) = (v_1, \dots, v_\ell)$ such that with respect to B every vertex has its two outgoing B -neighbors among its predecessors

and I . Denote by B_i the subgraph of B induced by the arcs leaving v_i, \dots, v_ℓ (before the reorienting). We process $V(T) \setminus V(I)$ from the last to the first element. At a given vertex v_i we look at $\text{dem}_{G \cup R \cup B_i}(v_i)$ and reorient the two originally outgoing B -arcs of v_i in such a way that afterwards $\text{dem}_{G \cup R \cup B_i}(v_i) = 0$ (i.e. $\delta_{G \cup R \cup B_i}^+(v_i) \equiv 0 \pmod{3}$). As these B -arcs were heading at I or at a predecessor, the demand on the vertices v_j , with $j > i$, is not modified and hence remains 0.

3.5 Orienting G and I

Denote by O the partial orientation of T obtained after 3.4. Pick an orientation of G (either all forward or all backward) and of uv such that for the resulting partial orientation O' we have $\text{dem}_{O'}(v) \equiv 1 \pmod{3}$.

Now, take the triangle Δ of I containing v . Since $\tilde{D} = I \setminus uv$ is a maximal outerplanar graph with only two degree two vertices, \tilde{D} can be peeled by removing degree two vertices until reaching Δ . When a vertex x is removed orient its two incident edges so that $\text{dem}_{O'}(x) = 0$ (as for B -arcs). We obtain a partial orientation O'' , such that all vertices except the ones of Δ have non-zero outdegree divisible by 3.

Since the number of edges of T , and the number of edges of Δ are divisible by 3, the number of edges of $T \setminus \Delta$ is divisible by 3. As this number equals the sum of the outdegrees in O'' , and as every vertex out of Δ has outdegree divisible by 3, then the outdegree of Δ 's vertices sum up to a multiple of 3. Hence their demands sum up to 0, 3 or 6. As $\text{dem}_{O''}(v) = \text{dem}_{O'}(v) = 1$, the demands of the other two vertices of Δ are either both 1, or 0 and 2. It is easy to see that in either case Δ can be oriented to satisfy all three demands.

4 Towards Schnyder woods

We see our proof of Conjecture 1 as a step towards generalizing Schnyder woods to triangulations of arbitrary orientable surfaces (the notion does not have much sense for non-orientable ones). By results of [8] another step towards generalizing Schnyder woods to maps of arbitrary orientable surfaces can be formulated after introducing a couple of definitions:

A map G is said essentially k -connected, if its universal cover is k -connected. Given a map G , its *primal-dual-completion* \hat{G} is the map obtained from simultaneously embedding G and its dual, G^* , such that vertices of G^* are embedded inside faces of G and vice-versa. Moreover, each edge crosses its dual edge in exactly one point in the interior, which also becomes a vertex of \hat{G} . Hence, \hat{G} is a bipartite graph with one bipartition consisting of *primal-vertices* and *dual-vertices* and the other partition consisting of *edge-vertices* (of degree 4).

Conjecture 2. Given an essentially 3-connected map G , the map \hat{G} has an orientation where primal- and dual-vertices have non-zero outdegrees divisible by three, and where edge-vertices have indegrees divisible by three, that is indegree 0 or 3 (i.e. outdegree 4 or 1).

References

1. J. BARÁT AND C. THOMASSEN, *Claw-decompositions and Tutte-orientations*, J. Graph Theory, 52 (2006), pp. 135–146.
2. O. BERNARDI, *Bijective counting of tree-rooted maps and shuffles of parenthesis systems*, Electron. J. Combin., 14 (2007), pp. Research Paper 9, 36 pp.
3. O. BERNARDI AND N. BONICHON, *Intervals in Catalan lattices and realizers of triangulations*, J. Combin. Theory Ser. A, 116 (2009), pp. 55–75.
4. L. CASTELLI-ALEARDI, E. FUSY, AND T. LEWINER, *Schnyder woods for higher genus triangulated surfaces, with applications to encoding*, Discrete & Computational Geometry, 42 (2009), pp. 489–516.
5. H. DE FRAYSSEIX, P. O. DE MENDEZ, AND P. ROSENSTIEHL, *On triangle contact graphs*, Combin. Probab. Comput., 3 (1994), pp. 233–246.
6. S. FELSNER, *Convex drawings of planar graphs and the order dimension of 3-polytopes*, Order, 18 (2001), pp. 19–37.
7. S. FELSNER, *Lattice structures from planar graphs*, Electron. J. Combin., 11 (2004), pp. Research Paper 15, 24 pp.
8. D. GONÇALVES, K. B. KNAUER, AND B. LÉVÊQUE, *Structure of Schnyder labelings on orientable surfaces*, CoRR, abs/1501.05475 (2015).
9. D. GONÇALVES AND B. LÉVÊQUE, *Toroidal maps: Schnyder woods, orthogonal surfaces and straight-line representations*, Discrete & Computational Geometry, 51 (2014), pp. 67–131.
10. B. MOHAR AND C. THOMASSEN, *Graphs on surfaces*, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, 2001.
11. W. SCHNYDER, *Planar graphs and poset dimension*, Order, 5 (1989), pp. 323–343.
12. W. SCHNYDER, *Embedding planar graphs on the grid*, in Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '90, 1990, pp. 138–148.