TRANSFORMING TRIANGULATIONS ON NONPLANAR SURFACES[*]

C. CORTÉS[†], C. I. GRIMA[†], F. HURTADO[‡], A. MÁRQUEZ[†], F. SANTOS[§], AND J. VALENZUELA[†]

Abstract. We consider whether any two triangulations of a polygon or a point set on a nonplanar surface with a given metric can be transformed into each other by a sequence of edge flips. The answer is negative in general with some remarkable exceptions, such as polygons on the cylinder, and on the flat torus, and certain configurations of points on the cylinder.

Key words. graph of triangulations, triangulations on surfaces, triangulations of polygons, edge flip

AMS subject classifications. 68U05, 52C99, 65D18, 68R10

DOI. 10.1137/070697987

1. Introduction. Most of the problems considered so far in computational geometry are restricted to the plane, or to the Euclidean 3-space. However, in many applications it is necessary to deal with input data that lies on a surface rather than in the plane. Recently, some works have been focused on solving some of the problems arising in those cases (cf. [8, 13, 16]). This paper is in this category, studying the graph of triangulations of a polygon on a surface.

Partitioning geometric domains into simpler pieces, such as triangles, is a common strategy to several fields, the finite element method being a most relevant example. In particular, the triangulation of polygons is an intermediate step in many algorithms in the area of computational geometry.

In many cases, we need to obtain not only a triangulation of a given region but also a "good" one. Some examples of this assertion can be found when it is desired to improve the quality of a graphic representation or to find a "nice" mesh on a given surface in order to apply finite element methods. When the *quality* of the triangulation with respect to some criterion is considered, and no direct method for obtaining the optimal triangulation is known, it is natural to perform operations that allow local improvements. The best-known method is the edge flip: when two triangles form a convex quadrilateral, their common edge is replaced by the other diagonal of the quadrilateral [2, 6]. This local transformation, introduced by Lawson in [12], can be combined if necessary with methods such as simulated annealing to escape local optima [7, 11] and has also been used for the purposes of enumeration [1]. It also admits several variations [17, 18]. Regarding the local operation we have just described, a basic issue is whether any two triangulations of a domain D can

be transformed into each other by means of a sequence of flips. If we define the *triangulation graph* of D as that graph $TG(D)$ having as nodes the triangulations of D, with adjacencies corresponding to edge flips, then the above question becomes obviously whether $TG(D)$ is a connected graph or not.

It is known that the graph of triangulations of a planar simple polygon or a point set with n vertices is connected and its diameter is $O(n^2)$, which is tight [5]. It is worth mentioning that even the case of a convex n-gon P has been thoroughly studied because $TG(P)$ is isomorphic to the rotation graph of binary trees with $n-2$ internal nodes [10, 19]. On the sphere, the situation is essentially the same as in the plane.

In this work we study the connectivity of the triangulation graph for simple polygons and point sets lying on surfaces. Regarding polygons, we prove that for the cylinder and the torus with their flat metrics the graph is always connected (if nonempty). For general surfaces and metrics the situation is usually the opposite. Even worse, for point sets only certain configurations on the cylinder have a connected graph of triangulations.

At this point it is convenient to clarify that with the general purpose of extending computational geometry to surfaces, it is necessary to "translate" some of the elements that usually appear in the plane to the surfaces; in our case, we need to know how to join a pair of points (in other words, how to translate the concept of segment); it is known that, in general, there are infinitely many geodesics joining two points but usually only one with the minimal length (see [4]). Thus, following the cited works [8, 13, 16] and others, this unique minimal geodesic joining a pair of points on a surface will be called *the segment* defined by that pair of points. In what follows, only segments between pairs of points will be considered.

Equally some words must be said about the surfaces, or, more concretely, about the metric, that we are considering here. In general, we will study the case of the locally Euclidean surfaces (those surfaces isometric to the plane in sufficiently small regions). These surfaces have two advantages; on one hand, they are general enough in order to model many practical cases or approximate some other metrics, and, on the other, they have an easy representation, as we will see in the next section. Nevertheless, in section 3 the results are presented in a more general context because we do not need the flat representation of the locally Euclidean surfaces (although an alternative proof of the main result of this section is presented later in the context of locally Euclidean surfaces).

The paper is organized as follows. In section 2 we give definitions and preliminary results, and we establish the notation that will be used throughout this paper. Section 3 shows one of the main results of this paper, which is that in every compact connected surface it is always possible to find a metric that admits polygons and point sets with nonconnected graphs of triangulations. Section 4 focuses on the connectivity of the graph of triangulations for both polygons and point sets on the locally Euclidean surfaces. We conclude in section 5 with some comments and open problems.

2. Preliminaries. As is known, many practical problems cannot be modeled by planar situations, and other surfaces are required. When we meet phenomena in which the same configuration of generating points appears in cycles, we may analyze them with the aid of a point configuration on the cylinder or the torus. These are two well-known surfaces since, together with the twisted cylinder (or infinite Möbius strip) and the Klein bottle, they easily admit quotient metrics that make them locally Euclidean. With these metrics, the graph of triangulations of a polygon both on the

cylinder and on the torus is connected, although this fact does not hold on the other two nonorientable surfaces.

We start this section summarizing the basic properties of the locally Euclidean surfaces, via their planar representation. A more complete study of them can be found in [15].

2.1. Locally Euclidean surfaces. A 2-dimensional locally Euclidean surface is a surface which is isometric with the plane in sufficiently small regions.

A *motion* in the plane is a map that preserves distances between points. The group of motions in the plane is denoted by $\mathrm{Mo}(\mathbb{R}^2)$ and consists of translations, rotations, reflections, and glide reflections.

A group $\Gamma \subseteq \mathrm{Mo}(\mathbb{R}^2)$ is said to be *uniformly discontinuous* if there exists a positive number d such that if γ is a motion in Γ and P any point in the plane being $\gamma(P) \neq P$, then the distance between P and $\gamma(P)$ is greater than or equal to d.

There are five different types of uniformly discontinuous groups of motions of the plane, up to isomorphisms: Types I, II.a, II.b, III.a, and III.b [15]. They can be generated as follows:

- Type I is generated by the identity motion.
- Type II.a is generated by a translation.
- Type II.b is generated by a glide reflection.
- Type III.a is generated by two noncollinear translation vectors.
- Type III.b is generated by a translation and a glide reflection, the direction of the translation vector being orthogonal to the axis of the glide reflection.

Given a group $\Gamma \subseteq \mathrm{Mo}(\mathbb{R}^2)$ and a point P in the plane, the *orbit* of P via Γ, denoted $\Gamma(P)$, is the set of the successive images of P under the action of the elements of Γ, that is, $\Gamma(P) = \{\gamma(P) : \gamma \in \Gamma\}$. For any uniformly discontinuous group of motions $\Gamma \subseteq \mathrm{Mo}(\mathbb{R}^2)$ the following notion of equivalence on points in the plane can be defined: points A and B are equivalent if they belong to the same orbit; namely, there exists a motion $\gamma \in \Gamma$ such that $\gamma(A) = B$. The orbits are then the equivalence classes under this relation. The set of all orbits of $\mathbb{R}^2$ under the action of Γ is written as $\mathbb{R}^2/\Gamma$ and is called the *quotient space*. The distance between two points (orbits) $\mathbf{A} = \Gamma(A)$ and $\mathbf{B} = \Gamma(B)$ in $\mathbb{R}^2/\Gamma$ is defined to be the shortest of the distances $|AB|$, where A and B are points of the plane with A belonging to $\mathbf{A}$ and B to $\mathbf{B}$.

Every locally Euclidean surface Σ corresponds to a uniformly discontinuous group Γ of motions of the plane so that Σ can be obtained from Γ as the quotient space $\mathbb{R}^2/\Gamma$. Hence there are exactly five types of locally Euclidean surfaces [15]: the plane (Type I), the cylinder (Type II.a), the twisted cylinder (Type II.b), the (flat) torus (Type III.a), and the Klein bottle (Type III.b). Although the term *flat torus* applies to surfaces generated by any group of motions of Type III.b, we will follow the convention that considers the translations to be orthogonal. If the translations are not orthogonal, then we call the surface so obtained a *skew torus*. As we will see in section 4.2, this distinction is not trivial and has important consequences on the connectivity of the graph of triangulations.

According to the above definitions and results, a point a of the surface defined by a uniformly discontinuous group Γ is specified by an orbit $\mathbf{A}$ of Γ. However, in order to specify a, there is no need to know all points of $\mathbf{A}$; we need only know one point A of $\mathbf{A}$, and then all the others are obtained from A by applying motions in the given group Γ. Therefore, in order to determine the set of all points of the surface, we need only specify some region of the plane, for example, a polygon, satisfying the following properties:

1. The region contains one point from every set of equivalent points of the plane.

2. No interior point of the region is equivalent to any other point of the region; that is, equivalent points of the region can lie only on the boundary.

A region in the plane satisfying 1 and 2 is called a *fundamental domain*, and the set of points on the surface is obtained from this region by identifying or gluing together equivalent points of its boundary. In general, we will use the fundamental domains that are more common in the literature, that is, an infinite band for both the cylinder and the twisted cylinder and a rectangle on the torus and the Klein bottle. In the skew torus it is also usual to consider as a fundamental domain a parallelogram whose sides are parallel to the direction of the translations.

In order to fix the points in the examples given in section 4, we will consider an orthogonal reference system in these surfaces which will be centered, for simplicity, in the leftmost side of the band or in the lowest leftmost corner of the rectangle (or parallelogram) considered as the fundamental domain. In the nonorientable case the OX axis will be taken to coincide with one glide reflection axis of Γ. The tesselations of the plane generated by the previous fundamental domains of each surface together with the orbit of a polygon are depicted in Figures 1 and 2.
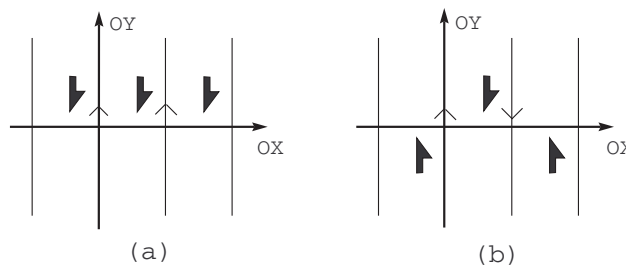

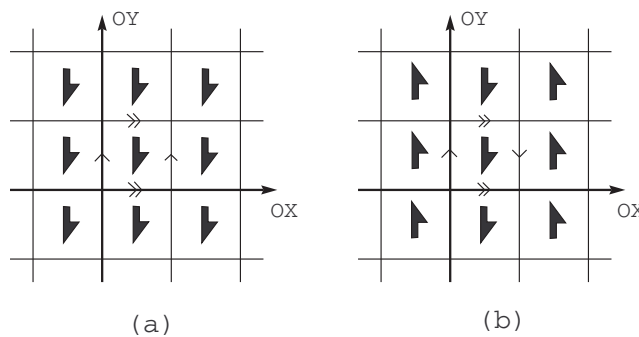FIG. 1. *The orbit of a polygon in* (a) *the cylinder and* (b) *the twisted cylinder.*

FIG. 2. *The orbit of a polygon in* (a) *the torus and* (b) *the Klein bottle.*

2.2. Triangulations of Euclidean polygons. Flips. A *Euclidean polygon* in a locally Euclidean surface is a region homeomorphic to a closed disc and whose boundary consists of finitely many geodesic arcs. A Euclidean polygon may be represented

as a simple planar polygon, although, depending on the election of the fundamental domain, it might not be completely contained in only one of them.

From now on, Euclidean polygons will be assumed to be already drawn in the plane.

The segment (that is, the minimum geodesic) between two nonconsecutive vertices of a Euclidean polygon is called a *diagonal* of the polygon. The diagonal $\overline{uv}$ is said to be *admissible* if it is contained inside the polygon (Figure 3).
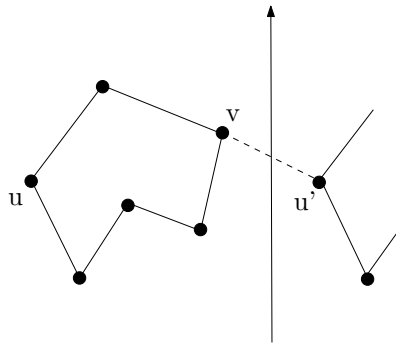

FIG. 3. *Since the nearest copy of u from v is u′, u and v cannot be matched inside the polygon and the diagonal $\overline{uv}$ is not admissible.*

A *(metrical) triangulation* of a Euclidean polygon is a partition of the polygon into triangular regions (that is, regions homeomorphic to a disc bounded by three segments) by means of *admissible* diagonals with no intersections except for their ends. Note that we force every face of a triangulation to be triangular instead of considering a maximal set of segments since, despite being equivalent definitions in the plane, this is no longer true in other surfaces, as will be apparent in section 4.2. In the same way, we define triangulations of point sets as a maximal set of noncrossing segments such that each bounded region is triangular. On the contrary, what happens in Euclidean polygons, given a point set the shape of the region triangulated, depends on the position of the points on the surface, and it can be a Euclidean polygon, or a strip bounded by two geodesics, or the whole surface (see [3, 8]).

Let $\{v_i, v_j, v_k\}$ and $\{v_i, v_j, v_l\}$ be two triangles in a triangulation sharing the diagonal $\overline{v_i v_j}$. By *flipping* $\overline{v_i v_j}$ we mean the operation of removing $\overline{v_i v_j}$ and replacing it by the other diagonal $\overline{v_k v_l}$ if it is admissible in the quadrangle $\{v_i, v_k, v_j, v_l\}$. The *graph of triangulations* of a polygon or a point set P is the graph $TG(P)$ having as nodes the triangulations of P, with adjacencies corresponding to diagonal flips (Figure 4).

3. Graph of triangulations of a polygon on nonplanar surfaces. One expects that metrical triangulations depend strongly on the metric considered since small changes in the metric might turn admissible diagonals into nonadmissible ones and flip performance would be affected. In this section, we define a metric on the sphere that produces polygons and point sets with nonconnected graphs of triangulations. The same idea will be used to extend this result to a general closed connected surface.

On the sphere, with its natural metric, geodesics correspond to great circles and the distance between two points is the length of the shortest arc of the great circle joining them (Figure 5), which is unique with the exception of *antipodal* (or diametri-
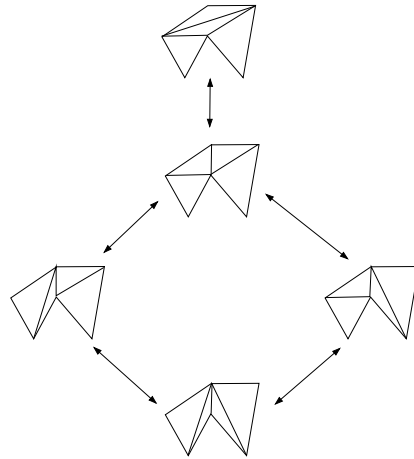
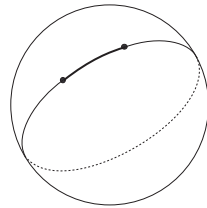FIG. 4. *The graph of triangulations of a polygon in the plane.*

FIG. 5. *The distance between two points in the sphere is given by the shortest arc of the great circle joining the points.*

cally opposite) points. A (Euclidean) polygon on the sphere, as in a locally Euclidean surface, is a region homeomorphic to a closed disc and whose boundary consists of finitely many geodesic arcs. Triangulations, flips, and graphs of triangulations of polygons on the sphere are also defined in the same way as they were in the previous section.

By using arguments similar to those in [12], it can be established that the graph of triangulations of any polygon on the sphere is connected with this metric. But it is possible to slightly disturb the metric so that this assertion will no longer be true.

LEMMA 1. *There exists a surface M homeomorphic to the sphere (in other words, M is a sphere with a metric other than the Euclidean distance) such that in M there exists a Euclidean polygon with a nonconnected graph of triangulations and a point set also with a nonconnected graph of triangulations.*

Proof. Consider a great circle C that divides the sphere into two open hemispheres H_1 and H_2. Let $p_1, p_2, \ldots, p_6$ be a sequence of vertices uniformly distributed on C such that the great circles joining (p_1, p_4), (p_2, p_5), and (p_3, p_6) intersect only in two antipodal points n and s in H_1 and H_2, respectively. We move the vertices $p_1, p_2, \ldots, p_6$ slightly toward n until the arc joining them inside H_1 is slightly shorter than the one that crosses through H_2.

Let $L = \langle p_1, p_2, \ldots, p_6 \rangle$ be a closed polygonal chain strictly contained in H_1, and let P be the polygon bounded by L whose interior is the region with a smaller area of the two into which the surface is divided by the polygonal chain. Now, M

is obtained from the sphere by lifting up a small region around n until the distances (considering the metric inherit from $\mathbb{R}^3$) between (p_1, p_4), (p_2, p_5), and (p_3, p_6), are enlarged enough to ensure that the diagonals joining them are nonadmissible in P (so those admissible diagonals are exterior to P), but without changing the length of the other diagonals of P (Figure 6(a)).
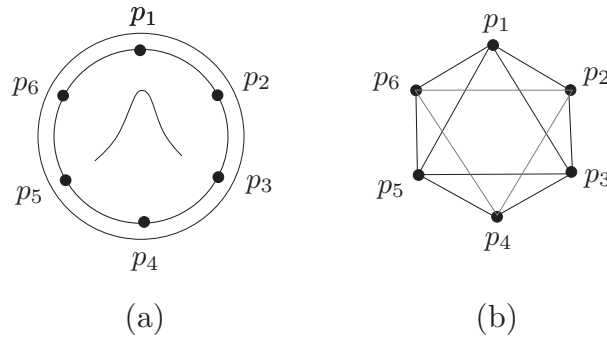

(a)   (b)

FIG. 6. *A hexagon with two disjoint triangulations in a "mountainous" sphere.*

After the lifting of the region around n, the length of any geodesic inside H_1 on M either is increased or remains the same as its length before the lifting. Moreover, the segments (shortest geodesic arcs) joining (p_1, p_4), (p_2, p_5), and (p_3, p_6) are the arcs of the great circles that join those points in H_2. Therefore, P admits only two different triangulations, shown in Figure 6(b), which cannot be transformed into each other by a sequence of flips, and hence, the graph of triangulations of P is nonconnected.

Basically, the same example can be used for point sets by adding a new vertex p_7 on s. To complete a triangulation, join p_7 to all the other vertices to obtain a set S. By construction, it is not possible to perform flips in any of the quadrilaterals having p_7 as a vertex (the new diagonals are outside the quadrilaterals). So, S has the two different triangulations of the original polygon P, and no flip is possible in any of those triangulations.   □

Using the previous lemma, the same reasoning can be extended to the remaining closed connected surfaces by using the fact that every closed connected surface is topologically equivalent to a sphere, or a connected sum of tori (handles), or a connected sum of projective planes.

THEOREM 1. *Any closed connected surface S admits a metric that allows polygons and point sets whose graphs of (metrical) triangulations are nonconnected.*

Proof. We can modify the surface M described in the proof of Lemma 1 by adding to it as many handles or projective planes as needed in order to obtain a surface homeomorphic to S.

By virtue of this fact, and mimicking the argument we followed on the sphere, it is possible to find a metric on each closed and connected surface that allows polygons and point sets with nonconnected graphs of triangulations (see Figure 7).   □

Note that the reasoning used in the proof of Theorem 1 can be easily extended to any kind of surface.

4. Connectivity of the graph of triangulations on locally Euclidean surfaces. As has been said in the introduction, some of the most common and useful surfaces are the locally Euclidean surfaces because of the advantage of their planar
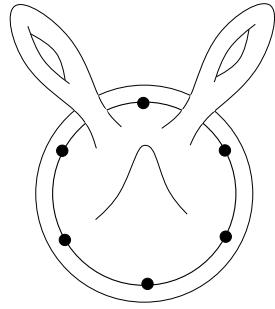
FIG. 7. *The construction of Figure* 6 *on the sphere with two handles.*

representations. It is interesting to emphasize the different behavior that these surfaces show when we study the graph of triangulations of a polygon: while the graph of triangulations is connected both in the cylinder and in the flat torus, polygons with nonconnected graphs can easily be constructed in the two nonorientable surfaces. The behavior of the graph of triangulations in the torus is remarkable, since that graph is connected for polygons with the metric of the flat torus, but this is not true for the skew torus. On the other hand, the graph of a point set is nonconnected in general, but, as we shall see next, we can describe all the connected components in the case of the cylinder.

4.1. The cylinder. Let $\vec{a}$ be the vector that generates the cylinder. Given that an orthogonal reference system is the OX axis parallel to $\vec{a}$, a geodesic arc is a segment if and only if its vertical projection is smaller than $|\vec{a}|/2$. In order to add a new diagonal to a triangulation, a procedure to determine if the geodesic arc joining two vertices it is a segment is to check if its vertical projection is contained inside the vertical projection of a previously existing diagonal (and, therefore, a segment).

4.1.1. Polygons. If the planar copies of a polygon P on the cylinder are (each of them) strictly contained in vertical bands of length $|\vec{a}|/2$, then any internal diagonal is admissible and planar arguments can straightforwardly be used to establish the connectivity of the graph of triangulations [8].

However, although many different proofs are known for planar polygons in the plane, the authors are not aware of any proof that can be adapted for the general case. Actually, it is not even obvious that in this general situation a polygon can always be triangulated; although, in this case, essentially the same ideas as in the plane provide a proof of this fact.

LEMMA 2. *Any Euclidean polygon of $n \geq 4$ vertices on the cylinder has an admissible diagonal. Hence, any Euclidean polygon on the cylinder is triangulable.*

Proof. This proof is based on the proof of Meister's lemma [14], which establishes the same result for simple polygons in the plane.

Consider a Euclidean polygon P already developed in the plane. Let v be a convex vertex such that the two edges incident on it go upward (recall that a vertex is *convex* if its interior angle is less than π radians; otherwise, the vertex is *reflex*). Let a and b be the vertices adjacent to v (Figure 8).

If $\overline{ab}$ is an admissible diagonal (a segment contained in P), then we have finished. Otherwise, either $\overline{ab}$ intersects ∂P or it is exterior to P.

If $\overline{ab}$ intersects ∂P, the argument given in [14] can be mimicked: Start sweeping a line from v, keeping it parallel to the line through $\overline{ab}$, until it reaches another vertex
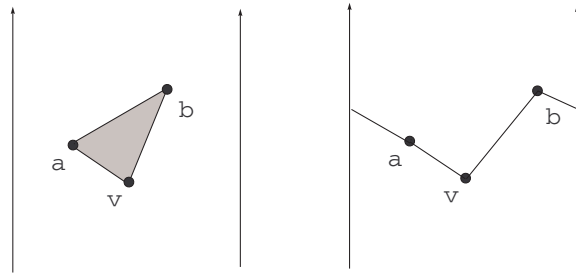
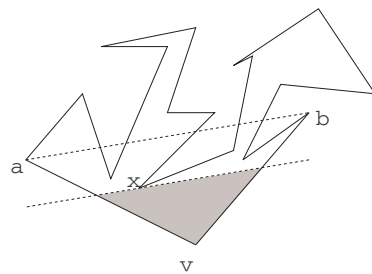FIG. 8. *The segment matching a and b may or may not determine a bounded triangle.*
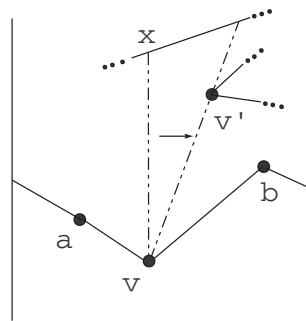

FIG. 9. *vx is an inner diagonal of the polygon.*

FIG. 10. *vv′ is a diagonal of the polygon.*

x of P (it must exist since P has at least four vertices). Then, $\overline{vx}$ is an admissible diagonal (Figure 9).

If $\overline{ab}$ is exterior to P, consider the vertical *ray* (half-line) with v as endpoint, and let x be the first point of the boundary of P that it reaches. If x is a vertex, then $\overline{vx}$ is an admissible diagonal. Otherwise, rotate the ray either to the right or to the left until it intersects another vertex v' of P (Figure 10). The vertical projection of $\overline{vv'}$ is contained inside the vertical projection of the diagonal containing x, so $\overline{vv'}$ is an admissible diagonal.   ☐

The connectivity of the graph of triangulations of a Euclidean polygon on the cylinder is established by the next theorem. As in the plane, three consecutive vertices

of the polygon a, v, b are said to form an *ear* if $\overline{ab}$ is an admissible diagonal: v is called the ear *tip*. Two ears are *nonoverlapping* if their triangle interiors are disjoint.

THEOREM 2. *Any two triangulations of a Euclidean polygon of $n \geq 4$ vertices on the cylinder can be transformed one into each other by a sequence of flips.*

Proof. Observe first that any triangulation of a Euclidean polygon of $n \geq 4$ vertices has at least two nonoverlapping ears. The proof of this fact in the plane [14] is essentially topological and works the same in the planar representation of the polygon, regardless of the metric.

Consider two triangulations T_1 and T_2 of a Euclidean polygon P with $n \geq 4$ vertices. We proceed by induction on the number of vertices (the base case is obvious). The inductive hypothesis straightforwardly leads to the result if T_1 and T_2 have a diagonal in common, since this diagonal divides P into two smaller polygons.

Suppose then that T_1 and T_2 do not share any diagonal. Let v_1 (resp., v_2) be the tip of an ear in T_1 (resp., in T_2). Since both T_1 and T_2 have at least two nonoverlapping ears, v_1 and v_2 can be assumed to be nonadjacent.

Consider $T_2 - \{v_2\}$, that is, the triangulation of a polygon of $n - 1$ vertices. By induction, $T_2 - \{v_2\}$ can be transformed by flips into another triangulation $T_2' - \{v_2\}$ having an ear in v_1 (Lemma 2 ensures that such a triangulation exists). As a consequence, T_2 can be transformed into another triangulation T_2' having an ear in v_1. Now v_1 is the tip of an ear both in T_1 and in T_2', which means that they share a diagonal.

It follows from the induction hypothesis that it is possible to go from T_1 to T_2 by flips through T_2': $T_1 \hookrightarrow T_2' \hookrightarrow T_2$.   □

It is easily seen that the proof of Theorem 2 implies an $O(2^n)$ upper bound for the number of flips. This bound is not tight at all on the cylinder since we have the same example as in the plane [5] for an $\Omega(n^2)$ bound, and any simple planar polygon, conveniently reduced, can be embedded in half a cylinder with the same (global) metric [8].

4.1.2. Point sets. Regarding the connectivity of the graph of triangulation of a point set on the cylinder, three different situations can be presented.

Consider the circle determined by cutting the cylinder with a plane orthogonal to its axis. If the smallest arc covering the orthogonal projection of the points on that circle is smaller than π, then the set is in *Euclidean position* and it has a planar behavior [3]. So, their graph of triangulation is connected.

If the set is not in Euclidean position, then there exist three of its points such that the polygonal line joining them wraps around the cylinder. We call this an *essential polygonal line*. The triangulated region of a set that is not in Euclidean position is bounded by two closed essential polygonal lines. The polygonal lines bounding the triangulated region are not uniquely determined by the points, and different triangulations of the same set may be bounded by different polygonal lines, describing different regions, as shown in Figure 11. In this case, as it is not possible to perform flips over the segments of the boundary, the graph of triangulations of the set is nonconnected.

Finally, the third case occurs when we consider triangulations of a point set that share the same boundaries. Then it is possible to carry one into each other by a sequence of flips, as is established in the following theorem.

THEOREM 3. *Given two triangulations of a point set on the cylinder, with the same boundaries, it is possible to transform one into the other by a sequence of flips.*

Proof. Let T_1 and T_2 be two triangulations of a point set S, and let U and L be the (upper and lower, respectively) essential polygonal lines bounding S. Denote D
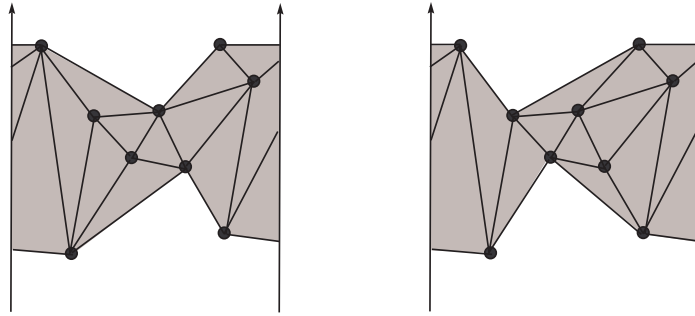
FIG. 11. *The points are the same, but the triangulated region is different. As no flip can be performed on the segments of the boundary, the graph of triangulations is disconnected.*

the region bounded by U and L.

The first step in the proof is to triangulate D using only admissible diagonals from U to L, or between two vertices of U or two vertices of L, without considering interior points of S. In order to meet this goal, we use the same arguments as in Lemma 2. Denote T_D the triangulation so obtained.

Now, let e be a diagonal in T_D, and let P_e be the polygon defined by the union of all the triangles in T_1 that intersect e. We can extend e to a triangulation of P_e by using Lemma 2. Obviously, by Theorem 2, we can transform one of those triangulations into the other, and so we can transform T_1 into another triangulation containing the diagonal e. We can perform the same process for all diagonals in T_D to obtain a triangulation T_D^1, and again starting from T_2 to obtain a new triangulation T_D^2 such that it is possible to transform T_i into T_D^i using admissible flips $(i = 1, 2)$ (Figure 12). The only differences between T_D^1 and T_D^2 are in diagonals that are contained in triangles of T_D. But the vertical projection of any diagonal inside a triangle of T_D is contained on the vertical projection of one of the sides of the triangle, so all the diagonals are admissible; then we have the same situation that we had in the plane, and, therefore, we can transform T_D^1 into T_D^2 using flips.   □
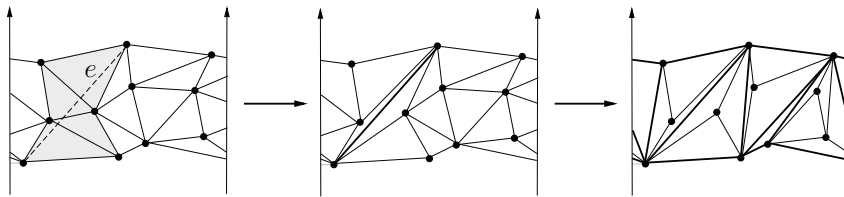

FIG. 12. *Construction of T_D^1.*

Of course, Theorem 3 provides a method of characterizing the connected components of the graph of triangulations of a point set on the cylinder in terms of the number of upper and lower polygonal lines; in [9] a tool is given (the polar diagram) that allows us to count those chains.

4.2. The torus. Throughout this section, unless otherwise stated, we assume we have a flat torus generated by a pair $\vec{a}$ and $\vec{b}$ of orthogonal vectors. Those vectors

are considered to be horizontal and vertical, respectively. A fundamental domain for this surface is an isothetic rectangle of dimension $|\vec{a}| \times |\vec{b}|$. Recall that an arc of a geodesic is a segment if and only if its vertical projection is smaller than $|\vec{a}|/2$ and its horizontal projection is smaller than $|\vec{b}|/2$. Note that this implies that if a planar copy of a set on the torus is contained on an isothetic rectangle of dimension $|\vec{a}|/2 \times |\vec{b}|/2$, then any diagonal is admissible and the set has a planar behavior. Analogously, if a planar copy of the set is inside a vertical (resp., horizontal) strip of width $|\vec{a}|/2$ (resp., $|\vec{b}|/2$), then the behavior is equivalent to being on the cylinder [3, 8].

We will center our efforts on proving the connectivity of the graph of triangulations of polygons on the flat torus. It is still an open problem whether the graph of triangulations is connected or not when we consider point sets instead of polygons. It is worthy to note that if the torus is generated by nonorthogonal vectors (skew torus), then sets and polygons with disconnected graphs of triangulations appear.

To establish the connectivity of the graph of triangulations of a polygon on the flat torus is more complicated than on the cylinder. The most common proofs of the connectivity of the graph of triangulations (either in the plane or on the cylinder) are based, in some sense, on the fact that every polygon is triangulable, but this is no longer true on the torus (as Figure 13 shows), and inductive reasonings fail. To be more precise, a maximal set of admissible diagonals does not necessarily divide the interior of the polygon into triangular regions. That is why the term *triangulation of a polygon* was defined in section 2.2 as a partition of the interior of the polygon into triangular regions rather than maximal sets of admissible diagonals.
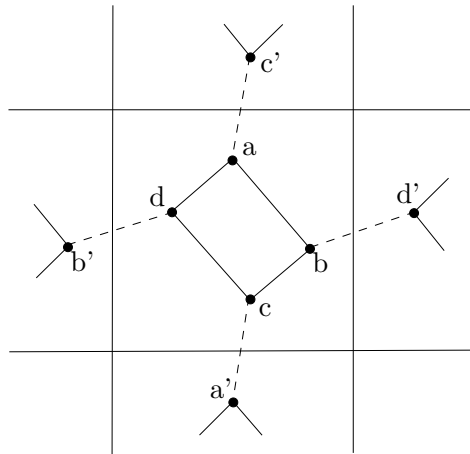

FIG. 13. *There is no admissible diagonal in this quadrilateral.*

Moreover, even if the polygon admits a triangulation, some admissible diagonals may not take part in any of them, as is shown in Figure 14, where the polygon is triangulable but the admissible diagonal $\overline{cd}$ does not participate in any triangulation. In Figure 15 we can see two different triangulations of a polygon where the method used on the cylinder fails. Since there is no admissible diagonal other than those shown in the figure, it is not possible to transform one of the triangulations into the other by keeping one ear in common with one of the two triangulations throughout the process.

This anomalous behavior of the torus restricts our attention to the study of the
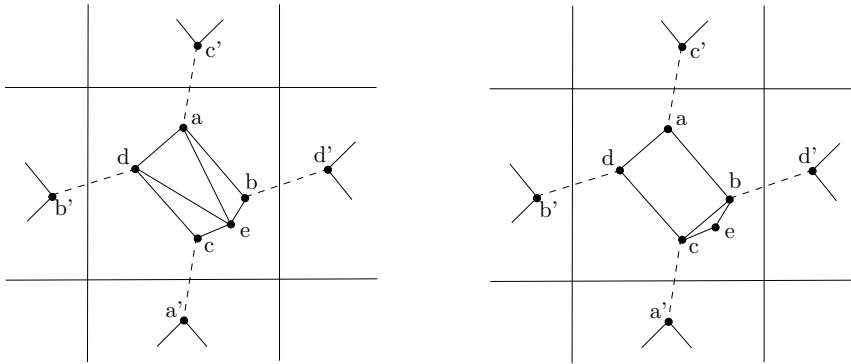
FIG. 14. *This pentagon is triangulable, but there exists no triangulation containing the diagonal* $\overline{cd}$.
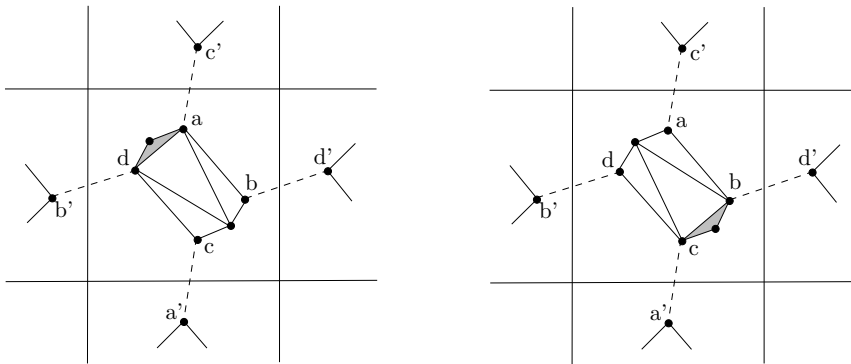

FIG. 15. *If we fix the selected ear in the first triangulation, as we did on the cylinder, it is not possible to perform a flip in the rest of the triangulation to carry it to one that contains the selected ear in the second triangulation.*

graph of triangulations of triangulable polygons and forces the search for new techniques to establish the connectivity of the graph of triangulations.

However, it is possible to give necessary conditions for a polygon to be triangulable on the flat torus. We define a *quadrant* in the torus as an isothetic rectangle of dimension $|\vec{a}|/2 \times |\vec{b}|/2$, with $|\vec{a}|$ and $|\vec{b}|$ being the generating vectors of the flat torus. It is easily seen that a diagonal of a polygon on the flat torus is admissible if and only if it fits into a quadrant. Then the following result is straightforward.

PROPOSITION 1. *If a polygon P on the torus contains the center of an empty (with no vertex of the polygon inside it) quadrant, then P is not triangulable.*

Proof. Suppose, contrary to our claim, that P is a triangulable polygon containing the center O of an empty quadrant. Then there must exist a triangle T_O of the triangulation of P in which O lies. Now it can easily be checked that some of the edges of T_O are not admissible, a conclusion contrary to our assumption.   ☐

We now introduce some definitions and preliminary results that will lead to a proof of the connectivity of the graph of triangulations of a triangulable polygon on the flat torus.

Let P be a Euclidean polygon on the flat torus. We consider a copy of P in the plane with the usual reference system. A convex vertex v of P is said to be a *top*

vertex (resp., *bottom*, *right*, and *left*) if the two edges incident on it go downward (resp., upward, leftward, and rightward). Top, bottom, right, and left vertices will be called *extreme* vertices; see Figure 16.

A vertex u of P is said to be *earable* if the segment joining its two adjacent vertices is an admissible diagonal.
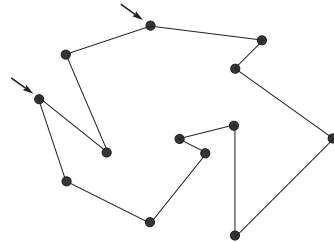

FIG. 16. *A polygon with the* top *vertices marked.*

LEMMA 3. *Any triangulable quadrilateral on the flat torus admits a triangulation having an ear in one of its extreme vertices.*

Proof. Since the two ears in a triangulation of a quadrilateral are at opposite vertices, and since every polygon has at least two extreme vertices, the only nontrivial case is that of a quadrilateral with two nonextreme vertices opposite one another. Such a quadrilateral is contained in the isothetic rectangle having the two extreme vertices as corners (Figure 17). This implies that if the diagonal joining the extreme vertices is admissible, then the diagonal joining the nonextreme vertices is admissible as well. One of them must be admissible because the quadrilateral is assumed to be triangulable. ⊓⊔
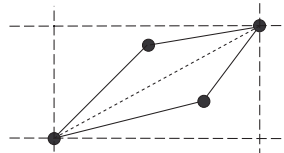

FIG. 17. *Any triangulable quadrilateral has an extreme earable vertex.*

Now, we extend the result given in Lemma 3 to a general polygon.

LEMMA 4. *Any triangulable polygon on the flat torus admits a triangulation having an ear in one of its extreme vertices.*

Proof. We proceed by induction in the number of vertices. The base case is given by the previous lemma, so consider a triangulable polygon P_{n+1} with $n > 4$ vertices and let T be a triangulation of P_{n+1}. Let v be the tip of an ear of T and let a and b be its adjacent vertices. If v is an extreme vertex in P_{n+1}, we finish. Otherwise, remove v and its incident segments $\overline{av}$ and $\overline{vb}$ from P_{n+1} to obtain a polygon P_n with n vertices. By the induction hypothesis, P_n admits a triangulation T' having an ear in one extreme vertex v'.

If v' is other than a and b, then v' is also an extreme vertex in P_{n+1} and $T'' = T' \cup \{\triangle avb\}$ is a triangulation of P_{n+1} having an ear in v', so the result holds. Therefore, assume T' has an extreme ear $\triangle cab$ at a.

Without loss of generality, suppose a is a right vertex of P_n. Let R be the lower triangle defined by the edge $\overline{ab}$ in the isothetic rectangle having $\overline{ab}$ as one of its diagonals (R is the shaded area in Figure 18). Since v is not extreme, it must be inside R. But this implies that a is also a right vertex in P_{n+1}, as is shown in Figure 18.
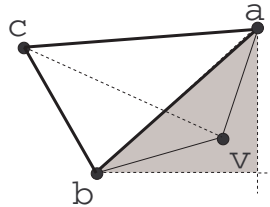

FIG. 18. *If v is outside R, then it is an extreme vertex. Otherwise, a flip can be performed to obtain an extreme ear in a.*

The vertical (resp., horizontal) projection of $\overline{vc}$ is inside the vertical (resp., horizontal) projection of $\overline{ac}$ (resp., $\overline{ab}$), and since both $\overline{ac}$ and $\overline{ab}$ are admissible diagonals, $\overline{vc}$ is admissible, and $\overline{ab}$ can be flipped in $T'' = T' \cup \{\triangle avb\}$ in order to get the extreme ear $\triangle cav$ in the triangulation T'' of P_{n+1}.   □

Now, we can prove the key result in order to obtain the connectivity of the graph of triangulations of a polygon on the flat torus.

LEMMA 5. *Let P be a triangulable polygon on the flat torus having an extreme earable vertex u. Then, any triangulation of P can be transformed by a sequence of flips into a triangulation having an ear in u.*

Proof. For the sake of simplicity we assume that u is a top vertex. Let T be a triangulation of P such that u is not an ear in T. Consider the subpolygon P' of P covered by the triangles of T incident to u (Figure 19). It is clear that u is the topmost vertex of this subpolygon, and since all the vertices of P' are joined to u by admissible diagonals, P' is contained in a horizontal strip of width $|\vec{b}|/2$. Thus, we can consider P' as it is embedded on a cylinder generated by $\vec{a}$, and, by our results in section 4.1, P' admits a triangulation which has an ear at the earable vertex u, and this triangulation is connected by flips to the restriction of T to P'. This finishes the proof.   □
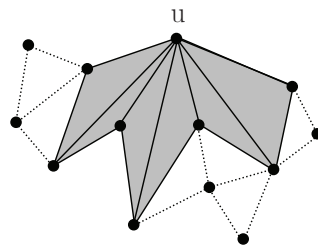

FIG. 19. *P' is made by the triangles incident to u.*

Now, we are in the position to enunciate the main result of this section.

THEOREM 4. *The graph of triangulations of a polygon on the flat torus is either empty or connected.*

Proof. Let T_1 and T_2 be two triangulations of a polygon P on the flat torus. Let u be an extreme earable vertex in P, which exists by Lemma 4. By virtue of Lemma 5, T_1 (resp., T_2) can be transformed by a sequence of flips into another triangulation T_1' (resp., T_2') having an ear in u. Therefore, T_1' and T_2' are connected by the inductive hypothesis using flips.   ☐

Regarding the connectivity of the graph of triangulations of a point set S in the flat torus there are three possible situations:

1. If S is inside a quadrant, then S is in Euclidean position, and it has a planar behavior [3], so the graph is connected.

2. If a planar copy of S is inside a vertical (resp., horizontal) strip of width $|\vec{a}|/2$ (resp., $|\vec{b}|/2$), the situation is equivalent to the cylinder. The graph is connected if and only if the borders of the triangulated region are fixed (section 4.1).

3. In the other case the connectivity of the graph of triangulations is still an open problem. Our conjecture is that this graph is connected.

Nevertheless, as we pointed out in section 3, the connectivity of the graph of triangulations is not preserved if the torus is generated by two nonorthogonal translations. In this way, consider the planar representation of a skew torus generated by two translations with vectors forming an angle of $\arccos\sqrt{\frac{1}{5}}$. In order to simplify the coordinates of the vertices, we choose a horizontal unitary vector and the other one with modulo $\frac{2\sqrt{5}}{5}$, and hence the height of a fundamental region is one unit. Using the usual reference system, we can draw a hexagon of vertices $a(\frac{1}{2}+\varepsilon, \frac{3}{4})$, $b(1-\varepsilon, \frac{3}{4})$, $c(1+\frac{\varepsilon}{3}, \frac{1}{2})$, $d(1-\varepsilon, \frac{1}{4})$, $e(\frac{1}{2}+\varepsilon, \frac{1}{4})$, and $f(\frac{1}{2}-\frac{\varepsilon}{3}, \frac{1}{2})$, with $\varepsilon < \frac{1}{8}$. Since the diagonals ad, be, and cf are not admissible, it is not possible to perform flips in either of the two triangulations depicted in Figure 20.
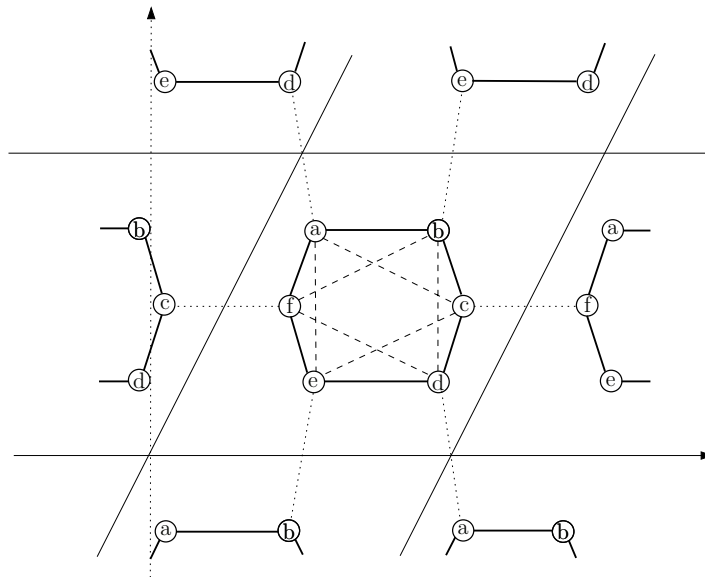

FIG. 20. *The previous hexagon with all the possible segments between its vertices.*

And, as we have done in section 3, new points can be added to the previous

construction to obtain a point set with a nonconnected graph of triangulations. We include points $g(0, \frac{3}{4})$, $h(\frac{1}{4}, \frac{3}{5})$, $i(\frac{1}{4}, \frac{2}{5})$, and $j(0, \frac{1}{4})$, as is shown in Figure 21. The central hexagon (bold lines) still admits only six diagonals, giving rise to only two different triangulations, and the segments of the boundary of the hexagon cannot be flipped. So the graph of triangulations of the set has two connected components.
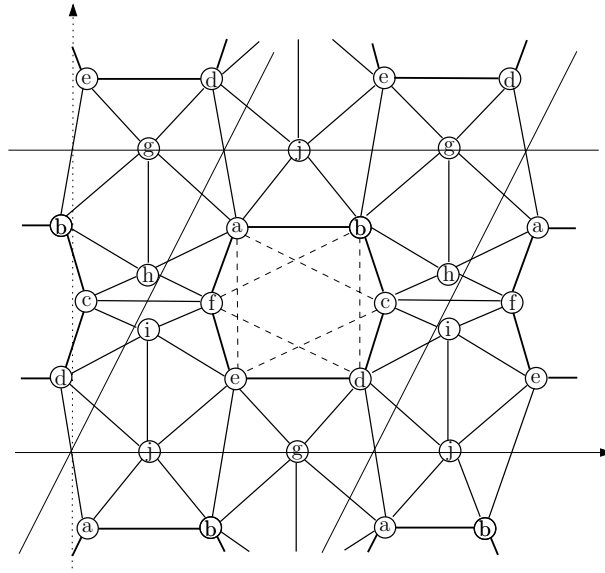

FIG. 21. *It is not possible to carry one of the triangulations of the central polygon into the other by flips.*

Therefore, the previous example shows (applying a suitable angle transformation if necessary) the following result.

THEOREM 5. *It is possible to find a polygon and a point set on a skew torus such that their graphs of (metrical) triangulations are nonconnected.*

It is worth pointing out that the previous result leads to another proof of Theorem 1.

4.3. Nonorientable locally Euclidean surfaces. It is easy to embed in the twisted cylinder and in the Klein bottle a polygon whose graph of triangulations is nonconnected. We look for a situation similar to the one used previously for the skew torus (Figure 20)—a hexagon of vertices (in clockwise order) $a, b, \ldots, f$ such that all the diagonals are admissible except for the diagonals (a, d), (b, e), and (c, f). This hexagon has only two triangulations, and it is not possible to perform a flip.

Consider the twisted cylinder with the usual coordinate system, and assume a glide reflection with unitary vector. The hexagon of vertices $a(\frac{1}{4} + \varepsilon, \frac{1}{4})$, $b(\frac{3}{4} - \varepsilon, \frac{1}{4})$, $c(\frac{3}{4} + \frac{\varepsilon}{3}, 0)$, $d(\frac{3}{4} - \varepsilon, -\frac{1}{4})$, $e(\frac{3}{4} + \varepsilon, -\frac{1}{4})$, and $f(\frac{1}{4} - \frac{\varepsilon}{3}, 0)$, with $\varepsilon < \frac{1}{16}$, has a nonconnected graph of triangulations (Figure 22).

The same construction can be easily obtained on the Klein bottle. A similar study can be extended to the other surfaces obtained as the quotient of the plane over a group of motions (*Euclidean* 2-*orbifolds*) if the group contains a glide reflection. In particular, this hexagon can also be embedded into the projective plane with the quotient metric.
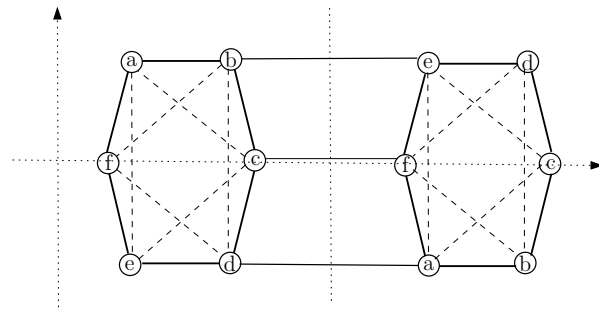
FIG. 22. *The segments ad, be, and cf are outside the polygon.*

By adding only two new points, as Figure 23 shows, the previous example can be extended to a point set with a nonconnected graph of triangulations. Again, as we saw in section 4.2 for the skew torus, the central hexagon has two possible triangulations and no flip is possible inside it. And, since the segments of the boundary of the hexagon cannot be flipped, the two triangulations belong to different connected components.
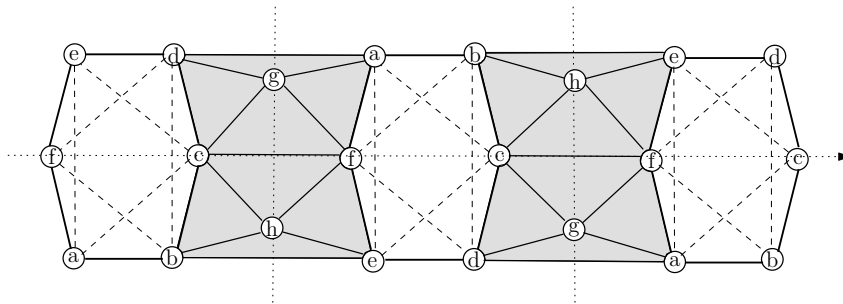

FIG. 23. *The only flips allowed are restricted to the shaded regions.*

Slightly more complicated is the example given for the Klein bottle (Figure 24), but the reasoning is the same; the segments that form the central hexagon cannot be flipped; hence the set has two disjoint triangulations.

5. Conclusions and open problems. In this paper we have studied the connectivity of the graph of triangulations of polygons and point sets on a surface. We have seen that, in general, this graph is nonconnected. More precisely, we have proven that any surface admits a metric such that there exist a polygon and a point set on that surface (and with that metric) with a nonconnected graph of triangulations. There exist some remarkable exceptions. The first among these, of course, are the plane and the sphere, and then polygons on the cylinder and the flat torus, these surfaces being the only ones to which Lawson's method [12] to obtain an optimal triangulation could be applied. Nevertheless the practical applications of this method are not clear since we have not found reasonable bounds for the diameter of the graph of triangulations in these surfaces.

Some problems are still unsolved. The main question that remains after Theorem 1 is whether it is possible to define a metric in a surface forcing the graph of triangulations of any polygon to be connected, as has been shown for the torus.
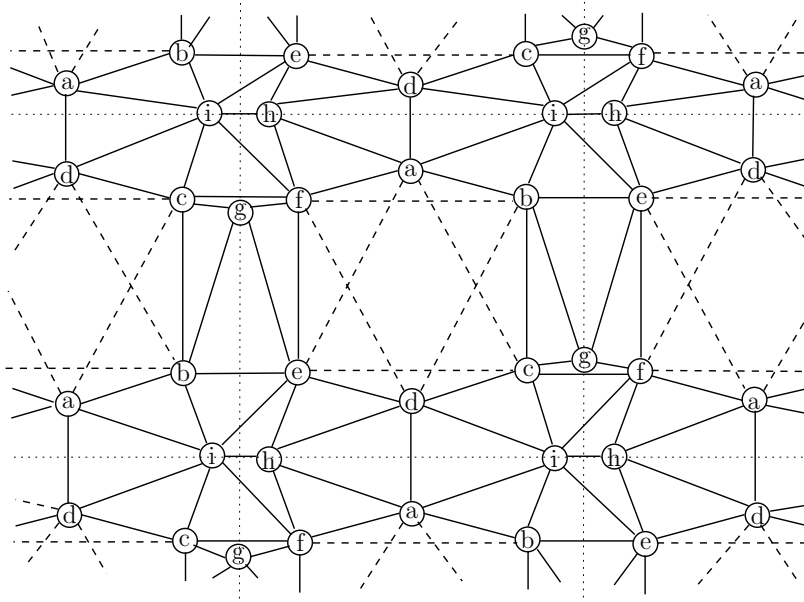
FIG. 24. *It is not possible to carry one of the triangulations of the central polygon into the other by flips.*

And, concerning the flat torus, the connectivity of the graph is not established if it is associated to triangulations of point sets instead of polygons.

REFERENCES

[1] D. AVIS AND K. FUKUDA, *Reverse search for enumeration*, Discrete Appl. Math., 6 (1996), pp. 21–46.

[2] M. BERN AND D. EPPSTEIN, *Mesh generation and optimal triangulation*, in Computing in Euclidean Geometry, D. Z. Du and F. K. Hwang, eds., World Scientific, River Edge, NJ, 1992, pp. 23–90.

[3] C. CORTES, A. MARQUEZ, AND J. VALENZUELA, *Euclidean position in euclidean 2-orbifolds*, Comput. Geom., 27 (2004), pp. 27–41.

[4] M. P. DO CARMO, *Differential Geometry of Curves and Surfaces*, Prentice–Hall, Englewood Cliffs, NJ, 1976.

[5] M. NOY, F. HURTADO, AND J. URRUTIA, *Flipping edges in triangulations*, Discrete Comput. Geom., 22 (1999), pp. 333–346.

[6] S. FORTUNE, *Voronoi diagrams and Delaunay triangulations*, in Computing in Euclidean Geometry, D. Z. Du and F. K. Hwang, eds., World Scientific, River Edge, NJ, 1992, pp. 193–234.

[7] C. D. GELATT, S. KIRKPATRICK, AND M. P. VECCHI, *Optimization by simulated annealing*, Science, 220 (1983), pp. 671–680.

[8] C. I. GRIMA AND A. MÁRQUEZ, *Computational Geometry on Surfaces*, Kluwer Academic Publishers, Dordrecht, The Netherlands, 2001.

[9] C. I. GRIMA, A. MÁRQUEZ, AND L. ORTEGA, *A new 2d tessellation for angle problems: The polar diagram*, Comput. Geom., 34 (2006), pp. 58–74.

[10] F. HURTADO AND M. NOY, *Graph of triangulations of a convex polygon and tree of triangulations*, Comput. Geom., 13 (1999), pp. 179–188.

[11] P. J. M. VAN LAARHOVEN AND E. H. L. AARTS, *Simulated Annealing: Theory and Practice*, Kluwer Academic Publishers, Dordrecht, The Netherlands, 1987.

[12] C. L. LAWSON, *Transforming triangulations*, Discrete Math., 3 (1972), pp. 365–372.

[13] M. MAZÓN AND T. RECIO, *Voronoi diagrams on orbifolds*, Comput. Geom., 8 (1997), pp. 219–230.

[14] G. H. MEISTER, *Polygons have ears*, Amer. Math. Monthly, 82 (1975), pp. 648–651.
[15] V. V. NIKULIN AND I. R. SHAFAREVICH, *Geometries and Groups*, Springer Series in Soviet Mathematics, Springer, Berlin, 1987.
[16] A. OKABE, B. BOOTS, AND K. SUGIHARA, *Spatial Tesselations. Concepts and Applications of Voronoi Diagrams*, John Wiley & Sons, New York, 1992.
[17] M. POCCHIOLA AND G. VEGTER, *Computing the visibility graph via pseudo-triangulations*, in Proceedings of the 11th Annual ACM Symposium on Computational Geometry, ACM, New York, 1996, pp. 248–257.
[18] F. SANTOS, *Geometric bistellar flips. The setting, the context and a construction*, in International Congress of Mathematicians, Vol. III, M. Sanz-Solé, J. Soria, J. L. Varona, and J. Verdera, eds., European Mathematical Society, Helsinki, Finland, 2006, pp. 931–962.
[19] D. D. SLEATOR, R. E. TARJAN, AND W. P. THURSTON, *Rotations distance, triangulations and hyperbolic geometry*, J. Amer. Math. Soc., 1 (1988), pp. 647–682.