Geometry

Geometric triangulations and flips

Triangulations géométriques et flips

Guillaume Tahar

Faculty of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot, 7610001, Israel

A R T I C L E I N F O	A B S T R A C T
	We prove that for a given flat surface with conical singularities, any pair of geometric triangulations can be connected by a chain of flips. © 2019 Académie des sciences. Published by Elsevier Masson SAS. All rights reserved.

R É S U M É

Nous démontrons que, dans chaque surface plate à singularités coniques, deux triangulations géométriques peuvent être reliées par une séquence de *flips*.
© 2019 Académie des sciences. Published by Elsevier Masson SAS. All rights reserved.

1. Introduction

Topological surfaces can be studied by combinatorial means through simplicial decompositions.

Definition 1.1. A topological triangulation of a topological surface with marked points and boundary (possibly empty) is a maximal family of topological arcs connecting marked points in such a way that they do not intersect themselves or each other in their interior. The arcs cut out the surface into ideal triangles (vertices and edges may be not distinct).

The set of topological triangulations of a given topological surface has a rich combinatorial structure. Elementary transformations of triangulations are called flips.

Definition 1.2. A flip is a transformation of an ideal triangulation that removes an edge that is a diagonal of a quadrilateral and replaces it by the other diagonal.

In the topological setting, it was proved in [2,3] that, for a topological surface with marked points, any pair of ideal triangulations can be joined by a chain of flips. In other words, the flip graph of topological triangulations of a topological surface is connected.

The problem is quite different in flat geometry. If we require edges to be geodesic segments, some topological triangulations may not appear as geometric triangulations of a given flat surface. Similarly, an edge of a flat triangulation cannot be flipped when the associated quadrilateral is nonconvex. We prove that a similar theorem holds in the flat setting as well as in the topological setting. The flip graph of geometric triangulations of a flat surface (which is a subgraph of the flip graph of topological triangulations) is also connected.

In the following, a flat surface is a topological compact surface with an everywhere flat metric outside of a finite set of conical singularities (of arbitrary angle) and boundary formed by a finite (possibly empty) union of geodesic segments connecting conical singularities. The boundary does not need to be connected.

In particular, such a definition includes translation surfaces and more generally $1/k$-translation surfaces, see [4,1]. We can also count marked points as conical singularities of angle 2π.

Definition 1.3. For a given flat surface, a geometric triangulation is a topological triangulation whose edges (including the boundary) are geodesic segments and whose vertices are conical singularities (every conical singularity should be a vertex of the triangulation).

Remark 1.4. There are no bigons in flat surfaces of finite area. Consequently, there is at most one geodesic segment in each isotopy class of topological arcs of the surface punctured at the conical singularities.

Theorem 1.5 *(Main theorem). For a given flat surface, any pair of geodesic triangulations can be connected by a chain of flips.*

2. Results

Definition 2.1. In a flat surface X, we define the *singular locus* $\text{Sing}(X)$ of X as the union of the conical singularities and the boundary.

Lemma 2.2. *In any flat surface, there is at least one geometric triangulation.*

Proof. We consider a flat surface X. If $\text{Sing}(X)$ is not connected, there is always at least one topological arc without self-intersection that connects two different components of $\text{Sing}(X)$. We take the shortest arc α with these requirements. It is a geodesic segment. Iterating this process provides a system of geodesic segments that join any connected components of $\text{Sing}(X)$. Cutting along these geodesic segments provides a new connected flat surface X^1 whose singular locus is connected. Existence of a geometric triangulation in X^1 implies existence of a geometric triangulation in X. Therefore, we reduced the problem to the case where the singular locus is connected.

Using a similar argument, we can once again reduce the problem to the case of flat surfaces with genus zero. If $\text{Sing}(X)$ is connected and the genus of X is nonzero, then X contains a topological arc β without self-intersection, whose both ends are the same conical singularity and which does not divide the surface. The geodesic representative of β is a chain of geodesic segments that do not intersect each other and that do not divide X. Cutting along these geodesic segments provides a new flat surface with smaller genus, whose singular locus is still connected. Iterating this process reduces the existence of a geodesic triangulation in X to the existence of a flat triangulation in some flat surface of genus zero.

Following discrete Gauss–Bonnet theorem, there are no flat surfaces of genus zero without boundary and with a unique conical singularity. Therefore, a flat surface of genus zero whose singular locus is connected is a polygon. We still have to prove the lemma in the case of polygons.

In any polygon there is a corner whose magnitude of angle is smaller than π. This corner A has two neighbors B and C (in the cyclical ordering of the corners of the polygon). We consider the topological arc γ that connects B and C staying close to two boundary geodesic segments. Concerning the geodesic representative of γ, there are three cases:

- it is a unique geodesic segment of the boundary; the flat surface is a triangle and there is a flat triangulation;
- it is a unique geodesic segment that does not belong to the boundary (it is a diagonal of the polygon); cutting along this segment provides two flat surfaces with strictly fewer corners;
- it is a chain of geodesic segments. Let D be one of the intermediary conical singularity of the chain. There is a geodesic segment between A and D (a diagonal of the polygon); cutting along this segment provides two flat surfaces with strictly fewer corners.

We have a systematic process to cut out a polygon into polygons with fewer corners. Besides, the lemma is trivially true for triangles. Therefore, the lemma holds for any polygon. This ends the proof. □

Definition 2.3. For a pair of distinct geodesic segments (α, β) of a given flat surface, we define $i(\alpha, \beta)$ as the number of intersection points of α and β outside their ends. This number is equal to the topological intersection.

Lemma 2.4. *Let X be a flat surface with conical singularities that admits two geodesic triangulations S and T such that:*

 (i) *for any pair $(\alpha, \beta) \in S \times T$, we have $i(\alpha, \beta) \geq 1$;*
(ii) *there is an edge $\alpha' \in S$ such that we have $i(\alpha', \beta) = 1$ for any $\beta \in T$.*

Then X is a convex quadrilateral.

Proof. We first choose an orientation of α'. Since α' intersects each internal edge of T once, the edges of T can be linearly ordered according to the position of their intersection with α'. Among triangles of T, we distinguish the two terminal triangles the ends of α' belong to. If the two terminal triangles coincide, then the two ends of α' belong to the same of corner of the terminal triangle because otherwise α' would intersect itself. Thus, α' crosses at least two times the edge opposite to this corner. Therefore, the two terminal triangles are distinct.

In each of the two terminal triangles, α' intersects once the edge opposite to the end corner. The two other edges of the triangle belong to the boundary of the surface. Indeed, if they were internal edges, then α' would cross them at least one time and α' would either intersect itself or cross an internal edge more than once.

Concerning triangles other than terminal, since the number of crossings of α' with the boundary of each of these triangles is even, exactly one edge of these triangles belongs to the boundary of the surface. Therefore, there is also a linear order for the triangles of the triangulation. The geometry of the flat surface is very restricted. It is a polygon, that is, a topological disk with a boundary formed by a chain of geodesic segments. Since every internal edge of S crosses each internal edge of T once, its ends belong to the same corners of the same terminal triangles. They pass through each internal edge of T in the same order.

Consequently, any other edge of S is isotopic to α'. There are no bigons in flat surfaces so there is a unique internal edge in S. Therefore, X is a quadrilateral. Since there is another triangulation T whose unique internal edge crosses α', X is a convex quadrilateral. □

Proof of the main theorem. The number of triangles in a flat triangulation of a given flat surface is determined by the sum of the total angles of the conical singularities. In addition, flat triangles have a total angle of π. Therefore, each geometric triangulation of a given flat surface has the same number of triangles. As a consequence, they have the same number of internal edges (edges that do not belong to the boundary). We suppose there exists a flat surface X with a minimal number of internal edges and such that there are two flat triangulations S and T of X that cannot be joined by a chain of flips.

Triangulations S and T cannot have a common internal edge because we could cut along this edge and get a flat surface with two additional boundary edges and one less internal edge. The two induced geometric triangulations of this new surface cannot be joined by a chain of flips. Thus, the minimality hypothesis is negated. Similarly, we prove that an edge x of S and an edge y of T always have a nontrivial intersection. Indeed, if there were such a pair of edges, we could complete x and y to get a new geometric triangulation Z. Indeed, any flat surface with conical singularities admits a geometric triangulation, see Lemma 2.2. Using the previous argument of minimality, S and Z would have a common edge and could be joined by a chain of flips. It is the same for T and Z. Therefore, S and T could be joined by a chain of flips.

For any pair of geodesic segments (α, β) of X with nontrivial intersection, we can define a pair of topological arcs (γ, δ) that is a desingularization of the pair (α, β). These topological arcs (without self-intersection) follow the geodesic segments outside a neighborhood of the intersection points of α and β. Near these intersection points, they are drawn in such a way that γ and δ do not intersect with each other. In addition, they do not intersect either α or β. In fact, there are two ways to provide such a desingularization depending on which endpoints of α and β are connected by desingularized arcs.

Homotopy classes of topological arcs γ and δ have geodesic representatives that are a chain of geodesic segments. These geodesic representatives minimize the self-intersection number and minimize the geometric intersection with α and β in their homotopy class.

We are going to prove that, in our minimal hypothetic counter-example, the singular locus is connected. We suppose that it is not. In each of the two triangulations S and T that cannot be joined by a chain of flips, we consider internal edges $\alpha \in S$ and $\beta \in T$ that join two connected components of the singular locus. Indeed, in any triangulation there is always enough internal edges to relate the connected components of the singular locus. The two edges α and β intersect each other. Desingularization provides two chains of geodesic segments that relate several connected components of the singular locus. No geodetic segment of these chains intersects α and β. Therefore, this segment u belongs to the boundary of X. Otherwise, we could provide a triangulation S' that contains both u and α. In the same way, there would be another triangulation T' with u and β. Cutting along u provides a simpler flat surface (with fewer internal edges) where any two geometric triangulations are joined by a chain of flips. Therefore, S' and T' are joined by a chain of flips. Since S and S' (similarly for T and T') share an internal edge, they are joined by a chain of flips (using the same argument). Consequently, any such segment u belongs to the boundary of the surface. So, the singular locus of X is connected.

Similarly, we prove that our minimal hypothetic counter-example has genus zero. We suppose its genus is nonzero. Then, in any triangulation of X, there is at least one internal edge such that cutting along it does not disconnect the surface. In each of the two triangulations S and T that cannot be joined by a chain of flips, we consider internal edges $\alpha \in S$ and $\beta \in T$ that do not disconnect the surface. These two geodesic segments have a nontrivial intersection. After desingularization, we

get a topological arc γ that does not intersect either α or β. Arc γ does not disconnect the surface, so its geodesic representative does not disconnect the surface either. It is true for any geodesic segment of the chain. In the same way, as we proved connectedness of the singular locus, from the existence of such a geodesic segment u that does not intersect α or β, we can construct a chain of flips that join S and T. Therefore, our minimal counter-example is a flat surface of genus zero whose singular locus is connected.

A flat surface of genus zero whose singular locus is connected is a polygon, that is, a topological disk whose boundary is a cyclic chain of geodesic segments. There should be at least three segments in the boundary, because otherwise, the angles in the inner corners would be degenerated.

In polygons, two geodesic segments have at most one intersection; otherwise, they would form a flat bigon. If there were two geometric triangulations that could not be joined by a chain of flips, then they would satisfy the hypothesis of Lemma 2.4, and the polygon would be a convex quadrilateral. This ends the proof. □

Acknowledgements

References

[1] M. Bainbridge, D. Chen, Q. Gendron, S. Grushevsky, M. Möller, Strata of k-differentials, Algebraic Geom. 6 (2) (2019) 196–233.
[2] S. Fomin, M. Shapiro, D. Thurston, Cluster algebras and triangulated surfaces. Part I: Cluster complexes, Acta Math. 201 (1) (2008) 83–146.
[3] A. Hatcher, On triangulations of surfaces, Topol. Appl. 40 (2) (1991) 189–194.
[4] A. Zorich, Flat surfaces, in: P. Cartier, B. Julia, P. Moussa, P. Vanhove (Eds.), Frontiers in Number Theory, Physics, and Geometry, Springer, 2006, pp. 439–586.