Available online at www.sciencedirect.com
ScienceDirect
Computer Aided Geometric Design 25 (2008) 157–161
COMPUTER AIDED GEOMETRIC DESIGN
www.elsevier.com/locate/cagd

All triangulations are reachable via sequences of edge-flips: an elementary proof

E. Osherovich *, A.M. Bruckstein

Computer Science Department, Technion Haifa, 32000 Israel

Received 3 April 2007; received in revised form 23 July 2007; accepted 23 July 2007

Available online 28 July 2007

Abstract

A simple proof is provided for the fact that the set of all possible triangulations of a planar point set in a polygonal domain is closed under the basic diagonal flip operation.

Keywords: Triangulations; Edge flip; Swapping diagonals; Triangulations of polygons; Triangulations of point sets

1. Introduction

Triangulation of point sets is an important task in many areas including computer graphics, computational geometry, and finite element computations. There are many ways we can triangulate a given set of points, and we might prefer one triangulation over another. Often a functional that maps every triangulation to a "quality measure" is introduced and we want to maximize this measure by an algorithm that starts with an arbitrary triangulation and transforms triangulations into a better ones by applying some simple operation. In this context the question whether we can reach the optimal triangulation from an arbitrary, initial one, naturally arises. More generally we may ask whether all triangulations are reachable by applying a sequence of the transformation operations?

For a well-known and basic *edge-flipping* operation the answer is yes. *Edge-flipping* can be performed for any two adjacent triangles of a given triangulation that jointly form a convex quadrilateral: we replace their shared edge with the other diagonal, as shown in Fig. 1.

The fact that the world of triangulations is closed under edge flips is not new. It was first proved, for convex polygons, by Lawson in 1972 (Lawson, 1972). All available proofs today are based on various geometric properties that often, despite their intuitive "obviousness", need lengthy and case-based proofs.

We present here a simple mathematical argument which reduces the use of geometric properties to several obvious and easily proved facts. Our proof is based on mathematical induction and is logically divided into two parts: first, we show flip-closure of the set of triangulations of an *empty* polygonal region, then we extend the proof to a set of points in a polygonal domain. Note that triangulation of an arbitrary set of points is a particular case of points inside

* Corresponding author.
E-mail addresses: oeli@cs.technion.ac.il (E. Osherovich), freddy@cs.technion.ac.il (A.M. Bruckstein).

0167-8396/$ – see front matter
doi:10.1016/j.cagd.2007.07.002

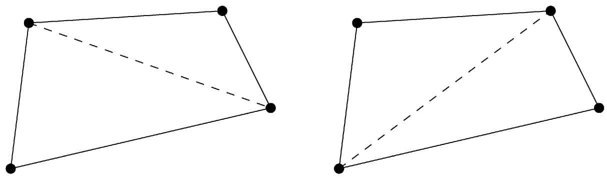

Fig. 1. The edge-flipping operation.

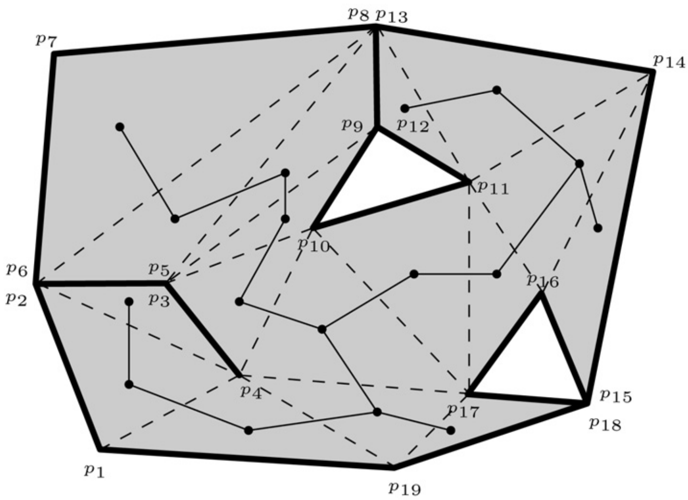

Fig. 2. Dual-tree of a simply connected polygonal region. Note that some triangles have overlapping, nevertheless different, edges, e.g., triangles $p_2p_3p_4$ and $p_5p_6p_8$ and consequently there is no edge in the dual tree connecting corresponding vertices.

a polygonal domain, since every such triangulation will include all edges of the convex hull, unambiguously defined by the point set.

2. Flip-closure of triangulations of polygonal domains

First we re-state the famous "two-ears" theorem (Meisters, 1975) in a slightly stronger version for simply connected polygonal regions. Recall that a polygonal region has a *ear* at vertex V_i if in the triangle formed by the three consecutive polygon vertices $V_{i-1}V_iV_{i+1}$ the (open) chord connecting V_{i-1} and V_{i+1} lies entirely inside the polygon.

Theorem 1. *Every nontrivial polygon with simply connected interior* (*that can be triangulated into more than one triangle*) *has two disjoint ears.*

The original theorem is usually proved for simple (Jordan) polygons only. The theorem remains true for a slightly wider family of polygons, the possibly *self-touching* polygons, having a simply connected interior. The original proof by Meisters works but instead we rely on the short proof found in O'Rourke (1987).

We first recall the definition of the *Dual-graph*: given a triangulated polygon with simply connected interior, the dual-graph is a graph generated by placing a vertex in each triangle and joining by edges vertices corresponding to adjacent triangles (triangle which share a side), as shown in Fig. 2. Note that this graph must be a tree, i.e., a graph without loops, since a loop would necessarily imply the existence of an internal point in the polygon and our polygon is, by assumption, empty.

Proof. Leaves in the dual-tree of the triangulated polygon correspond to ears and every tree of two or more vertices has at least two leaves. □

Next we show that one can transform any given triangulation of a polygon having simply connected interior into any other triangulation of the polygon.

Theorem 2. *Given two triangulations T_1 and T_2 of a polygon with simply connected interior, one can transform T_1 into T_2 by means of edge flips.*

Proof. We shall prove the theorem by induction on the number n of the polygon vertices. For $n = 3$ there is at most one possible triangulation, thus, inevitably, $T_1 = T_2$. Now let us assume the theorem holds for all n less than or equal to some k we shall show that it also holds for $n = k + 1$. According to the Theorem 1 every polygon with simply connected interior has two disjoint ears E_i and E_j, say located at vertices V_i and V_j, respectively. If E_i appears in both triangulations we can cut the ear resulting in two polygons with 3 and k vertices, respectively (as shown in Fig. 3) and thus, according to the induction hypothesis one can transform one triangulation into the other via edge flips.

What if the ear E_i does not appear in either one or even in both of the given triangulations? We shall show that we always can transform any given triangulation into a triangulation that has the ear E_i as one of its triangles. Let us first look at the polygon P_i induced by the vertex V_i and all its neighbors defined by the edges of the triangulation, see Fig. 4.

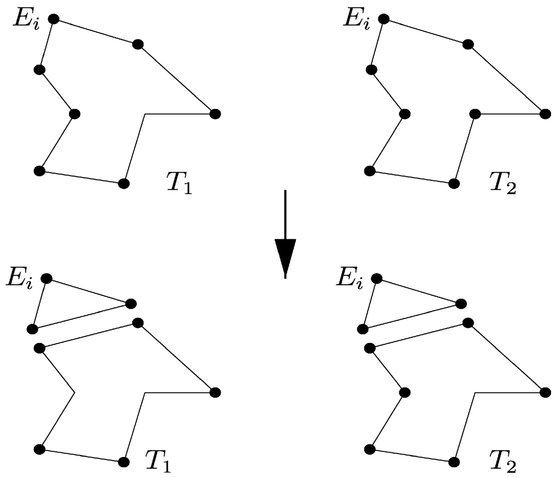

Fig. 3. Cutting ear E_i.

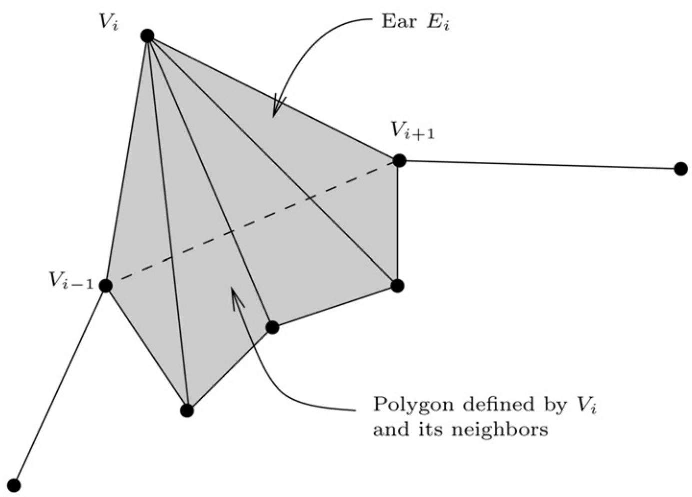

Fig. 4. The ear E_i does not appear in the triangulation.

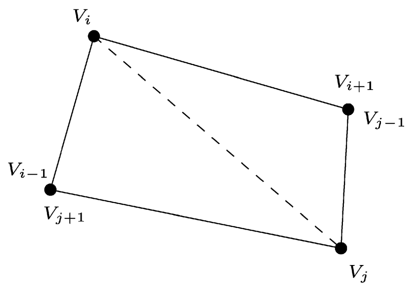

Fig. 5. Both V_i and V_j are connected to all vertices.

There are three possible scenarios:

1. The total number of vertices in P_i is less than $k+1$ (i.e., V_i is not connected to all other vertices). Then, according to the induction hypothesis the polygon P_i can be transformed into any other triangulation in particular into one with the ear E_i (and, obviously, there exists such a triangulation).
2. P_i has exactly $k+1$ vertices which means that V_i is connected to every other vertex of the polygon, while the polygon P_j defined by V_j (the vertex defining ear disjoint from E_i) and all its neighbors has less than $k+1$ vertices. According to the induction hypothesis P_j can be transformed into any other triangulation in particular to one with ear E_j. At this moment V_i cannot be connected to V_j anymore and thus the new P_i will have less than $k+1$ points and, therefore, can be transformed into a triangulation in which the ear E_i appears.
3. Both P_i and P_j have exactly $k+1$ vertices, that is both V_i and V_j are connected to every other vertex. It is easy to see that this case is only possible if we have a convex quadrilateral. Indeed, let us show it by construction: connecting V_i to every other vertex would generate a triangulation T, since all faces are triangles. In this triangulation V_j is connected to V_i and to its two neighbors: V_{j+1} and V_{j-1}, V_j cannot be connected to any other vertex, since this would break the planarity of the triangulation (triangulations are maximal planar graphs). But we assumed that V_j was connected to all vertices, thus we, inevitably have a convex quadrilateral. This case is depicted in Fig. 5. Here we can simply flip the edge $V_i V_j$. □

3. Flip-closure of triangulations of point sets in polygonal domains

Next we shall extend our proof to triangulations of points sets in polygonal domains. We start with the proof of a very useful property of triangulations:

Lemma 1. *Let T_1 be a triangulation of a set S of points lying inside a polygon P and let L be a line segment connecting two points a and b from the set S or from the vertices of P, such that L lies inside P, then T_1 can be transformed into a triangulation T_2 that exhibits L as one of the edges, by edge-flipping operations. Moreover, only edges that intersect with L need to be flipped.*

Proof. Recall that any triangulation is a maximal planar graph, thus, if T_1 does not contain L as an edge there exist edges $e_1, e_2, \ldots, e_m$ of T_1 that are intersected by L. We number the edges in the order they intersect L, say from a to b. Let us also denote by l_i (r_i) the endpoint of e_i that is to the left (right) of the directed line from a to b. Finally we define a set of points $u_1, u_2, \ldots, u_n$, each one corresponding to a set of consecutive l_i's that refer to the same point, and a set of points $v_1, v_2, \ldots, v_k$, each one corresponding to a set of consecutive r_i's. The ordered set of points $a, u_1, u_2, \ldots, u_n, b, v_k, v_{k-1}, \ldots, v_1$ defines a closed polygon (see Fig. 6 for possible examples). Note that the above polygon is not necessarily simple, but its interior is simply connected and it obviously is triangulated and there are no interior points in it. All this follows from the fact that the polygon is a union of triangles with edges that are crossed by L. Since, this polygon is properly triangulated, let us denote this triangulation by t_1. According to Theorem 2, t_1 can be transformed into any other triangulation, including a triangulation t_2, which has L as an edge. By transforming t_1 into t_2 we also transform T_1 into T_2. □

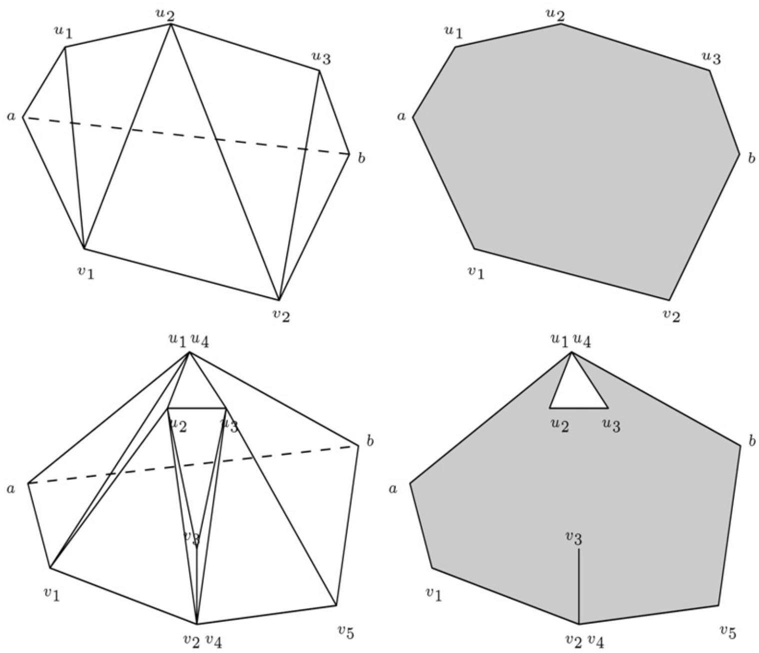

Fig. 6. Examples of possible polygons around L.

A similar lemma is used in other papers dealing with triangulations, for example (Dyn and Goren, 1993), proving closeness of the triangulations under the flipping operation, based on additional geometric properties which we do not need.

Finally, we are ready to prove the main statement of this paper:

Theorem 3. *Given any two triangulations T_1 and T_2 of a set S of points lying inside a polygonal domain P, one can transform T_1 into T_2 by means of edge flips.*

Proof. Let us enumerate edges of the triangulation T_1 in an arbitrary order $e_1, e_2, \ldots, e_r$. We run over all edges of T_1, e_i for $i = 1, 2, \ldots, r$ and check whether current edge e_i appears in a "transient triangulation" $T_{2\to 1}$ (initially $T_{2\to 1} = T_2$). If it does we go to the next edge, if it does not then according to the Lemma 1 we can make it appear in $T_{2\to 1}$. Note that during this process we only flip edges that properly intersect e_i from T_1 in the triangulation $T_{2\to 1}$, thus we do not flip edges $e_1, e_2, \ldots, e_{i-1}$ since they cannot properly intersect with e_i (as they all belong at this stage to both T_1 and $T_{2\to 1}$). Moreover, flipping edges that are intersecting with e_i in $T_{2\to 1}$ we do not create new intersections with e_j for $j < i$ because all these edges at this moment appear in $T_{2\to 1}$ and the edge-flipping operation does not create edge intersections (proper intersection of edges is not possible in triangulations). After we finish ($i = r$) all edges of T_1 appear in $T_{2\to 1}$ and since all triangulations of given points set have the same number of edges we conclude that $T_1 = T_{2\to 1}$. □

Acknowledgements

We would like to thank Professor N. Dyn for her valuable comments and our reviewers for their suggestions.

References

Dyn, N., Goren, I., Rippa, S., 1993. Transforming triangulations in polygonal domains. Computer Aided Geometric Design 10 (6), 531–536.
Lawson, C.L., 1972. Transforming triangulations. Discrete Mathematics 3, 365.
Meisters, G.H., 1975. Polygons have ears. The American Mathematical Monthly 82 (6), 648–651.
O'Rourke, J., 1987. Art Gallery Theorems and Algorithms. Oxford University Press, Inc.