Discrete Comput Geom 22:333–346 (1999)

Discrete & Computational
Geometry

Flipping Edges in Triangulations*

F. Hurtado,[1] M. Noy,[1] and J. Urrutia[2]

[1]Departament de Matemàtica Aplicada II,
Universitat Politècnica de Catalunya,
Barcelona, Spain
{hurtado,noy}@ma2.upc.es

[2]Department of Computer Science,
University of Ottawa,
Ottawa, Ontario, Canada
jorge@csi.uottawa.ca

Abstract. In this paper we study the problem of flipping edges in triangulations of polygons and point sets. One of the main results is that any triangulation of a set of n points in general position contains at least $\lceil (n-4)/2 \rceil$ edges that can be flipped. We also prove that $O(n+k^2)$ flips are sufficient to transform any triangulation of an n-gon with k reflex vertices into any other triangulation. We produce examples of n-gons with triangulations T and T' such that to transform T into T' requires $\Omega(n^2)$ flips. Finally we show that if a set of n points has k convex layers, then any triangulation of the point set can be transformed into any other triangulation using at most $O(kn)$ flips.

1. Introduction

Given a triangulation T of a set P of points on the plane, an edge e of T is *flippable* if it is adjacent to two triangles whose union is a convex quadrilateral C. By *flipping* e we mean the operation of removing e from T and replacing it by the other diagonal of C. In this way we obtain a new triangulation T' of P, and we say that T' has been obtained from T by means of a *flip*.

There are several reasons that make the study of flips in triangulations interesting. The first one is the existence of a simple greedy algorithm that constructs the Delaunay triangulation of a point set in general position by successive flips, starting from an

* F. Hurtado and M. Noy were partially supported by Proyecto MEC-DGES-SEUID PB96-0005-C02-02. J. Urrutia was supported by NSERC of Canada.

arbitrary triangulation of the point set (see [5], [9], and also [3] for a generalization of this result to higher dimensions). In this algorithm the iteration of local improvements produces a global optimum. A consequence of this success has been the use of flips in algorithms for finding triangulations that are at least approximate optima for criteria such as maximum angle, maximum vertex degree, total edge length, and minimum ratio between the area of the incircle and the area of the triangle [2].

Another motivation for studying flips comes from the existence of a bijection between triangulations of a convex $(n+2)$-gon and binary trees with n internal nodes. Under this bijection, flipping an edge in a triangulation corresponds precisely to a *rotation* in the corresponding binary tree [12] (see also [15] and [8]). Finally, we mention that the flip operation also appears in other kinds of triangulations; for instance, in the work of Avis [1] for enumerating rooted triangulations up to isomorphism, or in that of Pocchiola and Vegter [13] for computing the visibility graph of a set of objects in the plane.

Given a collection of points $P_n = \{v_1, \ldots, v_n\}$ we define the graph $G_T(P_n)$, the graph of triangulations of P_n, to be the graph such that the vertices of $G_T(P_n)$ are the triangulations of P_n, two triangulations being adjacent if one can be obtained from the other by flipping an edge. Triangulations of polygons, flipping edges in them, and their corresponding graphs of triangulations are defined in an analogous way.

In [15] Sleator et al. showed that the diameter of the graph G_n of triangulations of a convex n-gon is equal to $2n-10$ for n sufficiently large (the upper bound $2n-6$ for all n is easy). In [12] Lucas proved that G_n is Hamiltonian, and Lee [10] realized G_n as the skeleton of an $(n-3)$-polytope. Some properties of these graphs are also studied in [8], where it is also shown that the graph of triangulations of a simple polygon is connected and has diameter $O(n^2)$.

One of the main results in this paper is that any triangulation of an n point set contains at least $\lceil (n-4)/2 \rceil$ edges that can be flipped, and that this bound is tight. The remaining results concern mainly the diameter of graphs of triangulations. In Section 2 we give some preliminaries. In Section 3 we give a new simple proof that the graph of triangulations of a polygon is connected, that its diameter is bounded by twice the size of the visibility graph of the polygon, and that the diameter is $O(n^2)$, where n is the number of vertices. The proof of the latter result follows the analysis of the greedy flip algorithm to compute the Constrained Delaunay Triangulation of a planar straight line graph [2]. We use in the proof the upper bound of $O(n+k)$ for spiral polygons with n vertices, k of them being reflex, proved by Hanke [6]. We give examples showing that the diameter can be $\Omega(n^2)$. We also show that the diameter is sensitive to geometric features of the polygon: we prove an upper bound of $O(n+k^2)$, where k is the number of *reflex* vertices of the polygon.

In Section 4 we study triangulations of point sets on the plane. After proving the aforementioned result on the number of edges that can be flipped, we show that there are sets of n points such that the diameter of their graph of triangulations is $\Omega(n^2)$. A similar result appears in [5], but there only a restricted class of flips is allowed to transform one triangulation into another one. On the other hand, the quadratic upper bound for the diameter in Section 2 can be easily modified to the case of point sets. In fact, we prove a stronger result, namely that the diameter is $O(kn)$, where k is the number of convex layers of the point set. We conclude in Section 5 with some remarks and open problems.

A tool we use repeatedly is the insertion of a set S of suitable edges into a given triangulation, i.e., to perform a sequence of edge flips until we reach a new triangulation containing S. This operation is related to, but different from, the *edge-insertion technique* described in [4], which consists in adding a new edge e to the current triangulation, deleting edges that cross e, and retriangulating the resulting polygonal region in both sides of e.

2. Preliminaries

Let $P_n = \{v_1, \ldots, v_n\}$ be a collection of points on the plane. A *triangulation of* P_n is a partitioning of the convex hull $Conv(P_n)$ of P_n into a set of triangles $T = \{t_1, \ldots, t_m\}$ with disjoint interiors such that the vertices of each triangle t_i of T are points of P_n. The elements of P_n are called the vertices of T and the edges of the triangles $t_1, \ldots, t_m$ of T are called the edges of T. The degree $d(v_i)$ of a vertex v_i of T is the number of edges of T that have v_i as an endpoint. We say that an edge e of T is *flippable* if e is contained in the boundary of two triangles t_i and t_j of T and $C = t_i \cup t_j$ is a convex quadrilateral. By *flipping* e we mean the operation of removing e from T and replacing it by the other diagonal of C. See Fig. 1.

Given two triangulations T' and T'' of P_n, we say that they are at distance k if their distance in $G_T(P_n)$ is k, i.e., there is a set of triangulations $T_0 = T', \ldots, T_k = T''$ such that T_{i+1} can be obtained from T_i by flipping an edge of it, $i = 0, \ldots, k-1$. We also say that T' can be *transformed* into T'' by flipping k edges.

Throughout this paper, P_n denotes point sets and Q_n polygons. The vertices of Q_n are always assumed to be labeled $v_1, \ldots, v_n$ in clockwise order. We assume that no three consecutive vertices are collinear. When the internal angle $v_{i+1}v_iv_{i-1}$ is less than π we say that v_i is convex, otherwise we say it is reflex.

Let T be a triangulation of a polygon Q_n, and let v_i and v_j be vertices of Q_n such that the line segment v_iv_j connecting them is not an edge of T. We say that v_iv_j can be inserted in T by flipping $k-1$ edges if there is a sequence of triangulations $T = T_1, \ldots, T_k$ such that v_iv_j is an edge of T_k, and T_{i+1} can be obtained from T_i by flipping an edge of it, $i = 1, \ldots, k-1$. We say that a vertex v_i of Q_n is *external* if it lies in the convex hull of Q_n.

The *visibility graph* of Q_n is the graph with vertex set $\{v_1, \ldots, v_n\}$. Two vertices v_i and v_j of Q_n are adjacent in the visibility graph of Q_n if the line segment joining them has no point exterior to Q_n.

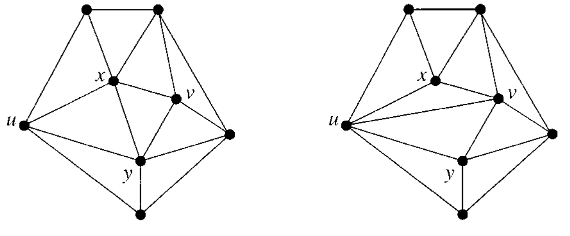

Fig. 1. Two triangulations of a point set. The second one is obtained from the first one by flipping edge xy.

3. Triangulations of Polygons

The main result of this section is Theorem 3.3, an upper bound on the diameter of the graph of triangulations $G_T(Q_n)$ of a simple polygon Q_n, sensitive to the number of reflex vertices of Q_n. The ingredients of this proof are (a) an upper bound on the diameter proportional to the size of the visibility graph, for which we provide a new simple proof; (b) the insertion in a linear number of flips of the *inner convex hull* of the polygon, i.e., the shortest polygonal chain that goes through all reflex vertices of the n-gon; and (c) a linear upper bound on the flip-distance between triangulations of spiral polygons due to Hanke [6]. We also prove in this section that there exist n-gons H_n such that the diameter of $G_T(H_n)$ is $\Omega(n^2)$.

We first show that the graph of triangulations $G_T(Q_n)$ of a simple polygon Q_n is connected and that the diameter is at most twice the number of edges of the visibility graph of Q_n. This fact can also be derived from the analysis of the greedy algorithm for computing the Constrained Delaunay Triangulation [2].

Consider the two vertices v_{i-1} and v_{i+1} of Q_n adjacent to a vertex v_i. The shortest polygonal chain joining v_{i-1} to v_{i+1} totally contained in Q_n will be denoted by $P_{i-1,i+1}$. We now prove:

Lemma 3.1. *Let Q_n be a simple polygon, let v_i be a convex vertex of Q_n, and let T be a triangulation of Q_n. Then we can insert all the edges of $P_{i-1,i+1}$ into T using exactly as many flips as the number of edges of T, not in $P_{i-1,i+1}$, intersecting $P_{i-1,i+1}$.*

Proof. Suppose that at least one edge e of $P_{i-1,i+1}$ is not in T. Consider the polygon P_e formed by the union of all triangles of T intersected by e, and consider the chain of vertices of P_e joining the endpoints of e. At least one of these vertices, say w, is a convex vertex of P_e, and thus the edge joining v_i to w can be flipped, decreasing the number of edges of T that intersect e by one. Our result now follows (see Fig. 2). □

Theorem 3.2. *The graph of triangulations $G_T(Q_n)$ of a simple polygon Q_n is connected. Moreover, the diameter of $G_T(Q_n)$ is at most twice the number of edges of the visibility graph of Q_n.*

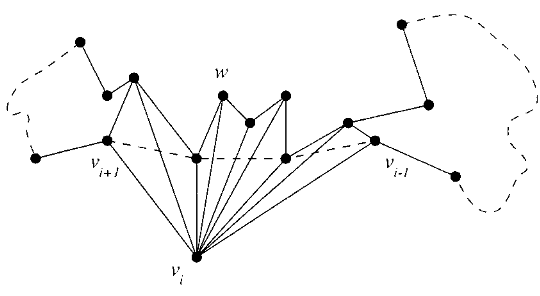

Fig. 2. Inserting the chain $P_{i-1}P_{i+1}$.

Proof. Let v_i be an external vertex of Q_n and let T_1 and T_2 be two triangulations of Q_n. Since v_i is convex, by Lemma 3.1 we can insert in each of T_1 and T_2 all the edges of $P_{i-1,i+1}$ to obtain two new triangulations T_1' and T_2' of Q_n. Delete from Q_n the subpolygon bounded by the vertices of $P_{i-1,i+1}$ and v_i. This will result in a collection of simple polygons with disjoint interiors. Each of these polygons has two triangulations induced by T_1' and T_2', respectively, and fewer vertices than Q_n. The connectivity of $G_T(Q_n)$ follows by induction on the number of vertices of Q_n.

To prove the second part of our result, we simply notice that each edge of the visibility graph of Q_n incident to v_i may be used twice; the first time while inserting $P_{i-1,i+1}$ into T_1 and the second time when we insert $P_{i-1,i+1}$ into T_2. Once we delete v_i from Q_n, these edges are not used again, and our result follows. Our argument gives a diameter of twice the number of edges of the visibility graph of Q_n. □

The bound on the diameter of $G_T(Q_n)$ given in Theorem 3.2 can, in general, be bad. For example, when Q_n is a convex polygon, the visibility graph of Q_n has $O(n^2)$ edges, while the diameter of $G_T(Q_n)$ is at most $2(n-3)$. On the positive side, if the visibility graph of Q_n has few edges, Theorem 3.2 gives us a method to transform one triangulation into another using few flips. Notice that if the visibility graph of Q_n has few edges, it has many reflex vertices. Thus the question of studying the tradeoffs in the diameter of $G_T(Q_n)$ and the number of reflex vertices of Q_n becomes relevant. We prove here that if Q_n is a polygon with k reflex vertices, then the diameter of $G_T(Q_n)$ is $O(n+k^2)$, i.e., the diameter of the graph of trangulations of Q_n depends mainly on its reflex vertices; convex vertices hardly matter.

Theorem 3.3. *Let Q_n be a simple polygon with k reflex vertices. Then the diameter of $G_T(Q_n)$ is $O(n+k^2)$.*

Several definitions and lemmas will be needed before we can prove Theorem 3.3. Let T be a triangulation of a polygon Q_n and let $v_i v_j$ be an internal visibility edge not belonging to T; $v_i v_j$ splits the polygon into two subpolygons Q' and Q''. We call V' (resp. V'') the set of vertices in Q' (resp. Q'') which are end-points of edges of T that cross $v_i v_j$. When all the vertices in V' or in V'' are convex we say that $v_i v_j$ is a *proper diagonal* with respect to T.

The following lemma will prove useful:

Lemma 3.4. *Let $v_i v_j$ be a proper diagonal of a triangulation T of a polygon Q_n. Then if $v_i v_j$ is intersected by t edges of T, $v_i v_j$ can be inserted in T using at most $2t$ flips.*

Proof. Let $v_i v_j$ be a proper diagonal with respect to a triangulation T. Assume without loss of generality that for each edge e of T intersecting $v_i v_j$, the end vertex of e below $v_i v_j$ is a convex vertex of Q_n. See Fig. 3.

Let $Q_{i,j}$ be the subpolygon of Q_n obtained by joining all the triangles of T intersected by $v_i v_j$ and consider the triangulation T' of $Q_{i,j}$ induced by T in $Q_{i,j}$. Suppose that $v_i v_j$ is intersected by t edges of T', $t \geq 1$. We now show that by flipping at most two edges of T' we can obtain a new triangulation in which $v_i v_j$ is intersected by $t-1$ edges.

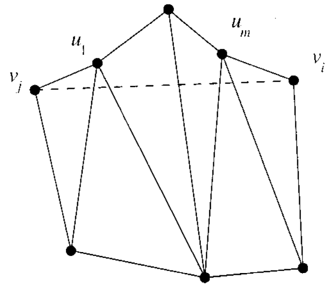

Fig. 3. $v_i v_j$ is a proper diagonal.

Let $u_1, \ldots, u_m$ be the vertices of $Q_{i,j}$ between v_j and v_i in the clockwise direction. At least one of these vertices, say u_t, is a convex vertex of $Q_{i,j}$; otherwise v_i and v_j would not be visible in Q_n. If in T', vertex u_t is adjacent to exactly one element in the chain $v_{i+1}, \ldots, v_{j-1}$, then the edge connecting them in T' can be flipped, reducing by one the number of edges that intersect $v_i v_j$. If u_t is adjacent to at least three vertices of $Q_{i,j}$ in $v_{i+1}, \ldots, v_{j-1}$, say v_{s-1}, v_s, and v_{s+1}, then we can flip the edge $u_t v_s$ inserting $v_{s-1} v_{s+1}$ and our result follows. Suppose then that u_t is adjacent to exactly two vertices, say v_s and v_{s+1}, in $v_{i+1}, \ldots, v_{j-1}$. See Fig. 4. Notice that since u_t is convex, we can flip $u_t v_{s+1}$. Next flip $u_t v_s$, and the number of edges intersecting $v_i v_j$ has gone down by one. Our result now follows. □

A polygon Q_n is called a spiral polygon if the vertices of Q_n can be labeled $v_1, \ldots, v_s, v_{s+1}, \ldots, v_n$ such that $v_1, \ldots, v_s$ are reflex vertices of Q_n and $v_{s+1}, \ldots, v_n$ are convex vertices of Q_n. We quote the following result obtained by Hanke [6].

Lemma 3.5. *Let Q_n be a spiral polygon with k reflex vertices. Then the diameter of $G_T(Q_n)$ is at most $2(n + k - 4)$.*

This result was proved by selecting a target triangulation T_0 defined recursively, and showing that it can be reached from any other triangulation with at most $n + k - 4$ flips.

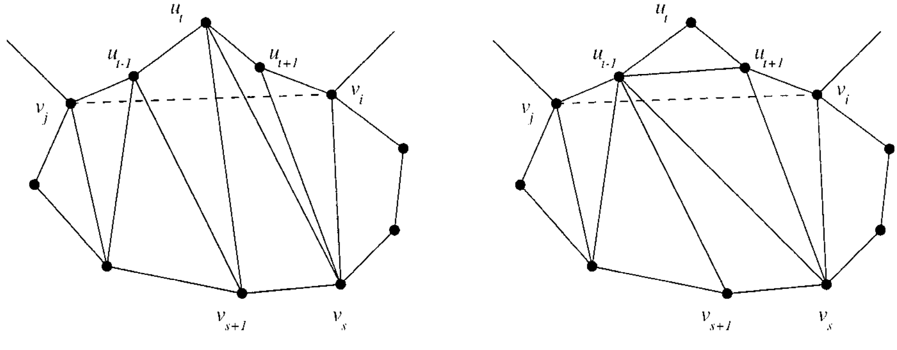

Fig. 4. Proof of Lemma 3.5.

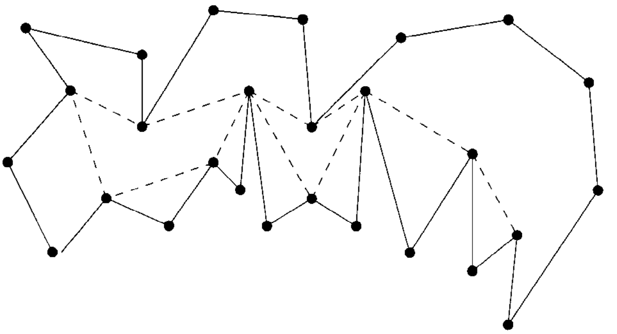

Fig. 5. Polygon R is dashed.

The method is based on locally progressing from any given triangulation toward T_0, exploiting the recursive definition of T_0.

Suppose next that Q_n has k reflex vertices labeled $v_{i_1}, \ldots, v_{i_k}$ such that $i_1 < \cdots < i_k$. For each $j = 1, \ldots, k$ let R_j be the shortest polygonal chain contained in Q_n joining v_{i_j} to $v_{i_{j+1}}$, addition taken mod k. Finally, let $R = R_1 \cup \cdots \cup R_k$. See Fig. 5. Then we have the following lemma:

Lemma 3.6. *Any internal visibility edge of Q_n crosses at most two edges of R. Moreover, if e is an edge of R and T is any triangulation of Q_n either e is an edge of T or e is a proper diagonal with respect to T.*

Proof. Every segment in R belongs to a shortest path R_j between reflex vertices v_{i_j} and $v_{i_{j+1}}$ which are "consecutive," i.e., a counterclockwise traversal of the boundary from v_{i_j} to $v_{i_{j+1}}$ encounters only convex vertices. Hence if we consider R_j as a path oriented from v_{i_j} to $v_{i_{j+1}}$, all the vertices strictly to the right side of the path are convex, and the second part of the statement is obvious.

For the first claim, let ab be an internal visibility edge crossing three segments u_1v_1, u_2v_2, u_3v_3 from R with the u_i on one side of ab and the v_i on other side, u_2v_2 crossing between the other two (see Fig. 6). We consider the case with six different extreme points (all the possibilities are handled similarly); they partition the boundary into six portions which we denote $[v_1, u_1), [u_1, u_2), \ldots, [v_2, v_1)$, closed at the origin and open at the end, in counterclockwise order. Now we prove that u_2v_2 cannot belong to any R_j: if v_{i_j} is in $[v_1, u_1)$, then R_j is inside the subpolygon $v_1 \cdots u_1v_1$; if v_{i_j} is in $[u_1, u_2) \cup [u_2, u_3)$, then R_j is inside $ba \cdots u_1 \cdots u_2 \cdots u_3 \cdots b$; if v_{i_j} is in $[u_3, v_3)$, then R_j is inside $u_3 \cdots v_3u_3$; and finally if v_{i_j} is in $[v_3, v_2) \cup [v_2, v_1)$, then R_j is inside $b \cdots v_3 \cdots v_2 \cdots v_1 \cdots ab$. □

We now prove:

Lemma 3.7. *Let T be any triangulation of Q_n. Then all the edges of R can be inserted in T using $O(n)$ flips.*

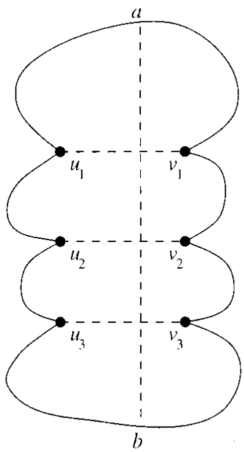

Fig. 6. Proof of Lemma 3.6.

Proof. According to Lemma 3.6, the diagonals of T cross R at most twice. Let s be the number of diagonals that cross T twice. We show that there always exists a flip that removes one of them. Hence, they can all be removed using s flips, and clearly $s < n$. Let then uv be a diagonal crossing R twice, and observe that since u and v are convex vertices, uv is flippable. Let w be a vertex such that uvw is a triangle of T. If w is reflex, then we can flip uv obtaining a new diagonal ww' that crosses R at most once. Otherwise, either uw or vw crosses R twice. Suppose it is uw, and let $x \neq v$ be such that uwx is a triangle of T. If x is reflex we flip uw and we are done as before. Iterating the process we eventually find a reflex vertex and we flip the corresponding diagonal.

We can now assume that no diagonal of T crosses R twice. For every edge e of R consider the polygon P_e formed by the union of all the triangles of T that cross e. Since e is the only diagonal of P_e that belongs to R and is a proper diagonal with respect to the triangulation induced on P_e, we can apply Lemma 3.4 to the polygon P_e to insert e without creating new crossings with R. Summing up, at most $2n$ flips are needed to insert R in T. □

We can now finish the proof of Theorem 3.3.

Proof of Theorem 3.3. Let T and T' be any two triangulations of Q_n. By Lemma 3.7, by flipping $O(n)$ edges, we can transform each of them into triangulations T_1 and T_1', respectively, of Q_n such that each of them contains all the edges of R. The edges of R induce a partition of Q_n into a set of polygons of two types:

(a) A set of at most k convex or spiral polygons $Q^1, \ldots, Q^m$, $m \leq k$, bounded by edges of Q_n and edges of R.
(b) A set of polygons $R_1', \ldots, R_s'$ bounded by the edges of R such that the total number of edges of these polygons is at most k.

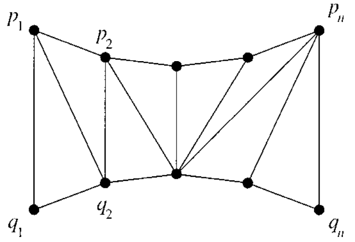

Fig. 7. A triangulation of the polygon H_5.

Notice that the total number of edges bounding $Q^1, \ldots, Q^m$ is at most $n + k$. Both T_1 and T_1' induce triangulations of $Q^1, \ldots, Q^m$ which may be different. Since each $Q^1, \ldots, Q^m$ is a spiral or a convex polygon, by Lemma 3.5 the triangulations induced by T_1 in $Q^1, \ldots, Q^m$ can be transformed into those induced by T_1' in $Q^1, \ldots, Q^m$ using at most $O(n)$ flips. Since the total number of edges bounding all the polygons in $R_1', \ldots, R_s'$ is at most k, then by Theorem 3.2 or [8] the triangulations induced in them by T_1 and T_1' can be transformed into each other with at most $O(k^2)$ flips. The proof is now finished. □

We close this section by producing a polygon H_n with $2n$ vertices such that the diameter of $GT(H_n)$ is exactly $(n-1)^2$.

Consider the polygon with $2n$ vertices $H_n = \{p_1, \ldots, p_n, q_1, \ldots, q_n\}$ such that $\{p_1, \ldots, p_n\}$ lie on a convex curve, $\{q_1, \ldots, q_n\}$ lie on a concave curve, and the line joining p_i to p_j, $1 \le i < j \le n$, leaves all the elements of $\{q_1, \ldots, q_n\}$ below it, and all the elements of $\{p_1, \ldots, p_n\}$ lie above any line joining q_i to q_j, $1 \le i < j \le n$; see Fig. 7.

Any triangulation T of H_n can be encoded as follows: Each triangle t_i of T has either two vertices in $\{p_1, \ldots, p_n\}$ or two vertices in $\{q_1, \ldots, q_n\}$. In the first case, assign a 1 to t_i; in the second case, t_i is assigned a 0. See Fig. 7.

If we read the numbers assigned to the triangles of T from left to right, we obtain an ordered sequence of zeros and ones; this sequence is the code assigned to our triangulation.

The triangulation of H_n presented in Fig. 7 receives the code 01011100. It is clear that each triangulation of H_n is thus assigned a unique sequence containing $n-1$ zeros and $n-1$ ones. Clearly, each sequence of $n-1$ zeros and $n-1$ ones also defines a unique triangulation of H_n, and thus we have a one-to-one correspondence between the set of triangulations of H_n and the set of binary sequences containing $n-1$ zeros and $n-1$ ones. Flippings can be easily identified within this encoding. An internal edge of a triangulation T can be flipped if the triangles of T containing it have been assigned a 1 and 0. Moreover, a flip of T simply corresponds to a transposition of a 0 with an adjacent 1 in the code of T.

Consider the triangulations T_1 and T_2 of H_n that receive the encodings $11\ldots100\ldots0$ and $00\ldots011\ldots1$. It is now clear that to transform T_1 to T_2 we need $(n-1)^2$ flips. We have just obtained:

Theorem 3.8. *The diameter of $G_T(H_n)$ is exactly $(n-1)^2$.*

We remark that a similar construction can be found in [5]. However, in our case the points are in general position and we are not restricted to Delaunay flips.

4. Triangulations of Point Sets

The main results in this section are a tight bound on the number of edges that can be flipped in any triangulation of a point set in general position, and an upper bound on the diameter of the graph of triangulations of a point set P_n that is sensitive to the number of convex layers of P_n. The basic ingredient for the second result is a lemma on how to insert the second convex layer of P_n into any triangulation.

We start by answering the following question: Given a triangulation T of a collection $P_n = \{v_1, \ldots, v_n\}$ of n points on the plane, how many edges of T are flippable? We show:

Theorem 4.1. *Any triangulation of a collection P_n of n points on the plane in general position contains at least $\lceil (n-4)/2 \rceil$ flippable edges. Moreover, the bound is tight.*

Observe that the general position assumption is clearly necessary, since otherwise it is easy to construct triangulations where no edge can be flipped, for example by inserting $n-3$ points in one of the sides of a triangle.

Some definitions are needed before we can prove Theorem 4.1. Let T be a triangulation of P_n. We divide the set of edges of T into two subsets, $F(T)$ containing all flippable edges of T, and $NF(T)$, the set of not flippable edges of T. Clearly all the edges of T contained in the boundary of $Conv(P_n)$ are not flippable. We orient the edges e of $NF(T)$ according to the following rules:

(R1) If e is an edge of $Conv(P_n)$, orient it in the clockwise direction around $Conv(P_n)$ of P_n.

(R2) If $e = uv$ is not in $Conv(P_n)$, consider the quadrilateral C formed by the union of the two triangles of T containing e. Since C is not convex, it follows that one of the end vertices of e, say v, is a reflex vertex of C while u is a convex vertex of C. Orient e from u to v.

The following observation will be useful:

Observation 1. The angle of C at v is greater than π.

Consider any vertex v_i of T. Let $d^-(v_i)$ be the number of edges $v_i v_j$ in $NF(T)$ oriented from v_j to v_i. Notice that $d(v_i)$ is the total number of edges of T incident with v_i, whereas $d^-(v_i)$ involves only edges of T in $NF(T)$. We now prove:

Lemma 4.2. *Let v_i be any vertex of T. Then $d^-(v_i) \leq 3$. Moreover, if $d(v_i) \geq 4$ in T, then $d^-(v_i) \leq 2$. Furthermore, if $d^-(v_i) = 2$, the two edges oriented toward v_i are consecutive edges around v_i.*

Proof. It is clear that if v_i is in $Conv(P_n)$, then $d^-(v_i) = 1$. Suppose then that v_i is in the interior of $Conv(P_n)$. Two cases arise:

(a) $d(v_i) = 3$ in T. In this case, all the edges of T incident with v_i are nonflippable and are oriented toward v_i. It follows that $d^-(v_i) = 3$.
(b) $d(v_i) > 3$ in T. In this case it is trivial to verify using Observation 1 that no more than two edges of G can be oriented toward v_i and that if $d^-(v_i) = 2$, then the two edges oriented toward v_i must be consecutive edges incident to v_i. □

We are now ready to prove Theorem 4.1.

Proof of Theorem 4.1. Let P_n be a point set on the plane, let T be a triangulation of P_n, and let S be the set of elements of P_n with degree 3 in T that are not in the convex hull of P_n.

By adding a point w in the exterior of $Conv(P_n)$ and joining it with a set of *disjoint curves* to all the vertices of $Conv(P_n)$, we obtain a triangulation of the plane with $n + 1$ points which by Euler's theorem contains $3n - 3$ edges. (Notice that in this triangulation, the triangular regions outside of $Conv(P_n)$ are bounded by two curves and a line segment.) We classify the edges incident to w as nonflippable edges and orient them from w to their other vertex in $Conv(P_n)$. Next orient all nonflippable edges of T according to (R1) and (R2). Notice that with these orientations, $d^-(v_i) = 2$ for all the elements of P_n of $Conv(P_n)$.

Remove from T all the elements of S. Notice that we remove exactly $3|S|$ edges of T which are not flippable. Notice that what remains is still a triangulation T' of $P_n - S + \{w\}$, which by Euler's formula contains $2(|P_n - S| + 1) - 4 = 2(n - |S|) - 2$ triangles. Moreover, any element v_i of $P_n - S + \{w\}$ that is not on the convex hull of P_n, has degree at least 4 in T, and by Lemma 4.2 has $d^-(v_i) \leq 2$. Let Q be the set of vertices of $P_n - S + \{w\}$ that have $d^-(v_i) = 2$. Then by Lemma 4.2 we can associate to every vertex v_i of Q in the interior of $Conv(P_n)$ a triangle $t(v_i)$ of T' which is also a triangle in T, bounded by two oriented edges, and the triangles $t(v_i)$ are all different. See Fig. 8. To each vertex v_i of T' in the convex hull of P_n we can also associate a different "triangle" of T' among those having w as one of their vertices.

That is, to each vertex of T', except w and the vertices of T with $d^-(v_i) < 2$, we can associate a different triangle of T' that contains no element of S. Clearly, the number of edges of T that can be flipped is minimized when all the vertices v_i not in S have $d^-(v_i) = 2$. In this case, we have associated exactly $n - |S|$ triangles to vertices of Q. As T' has $2n - 2|S| - 2$ triangles, at most $n - |S| - 2$ of them contained points of S in

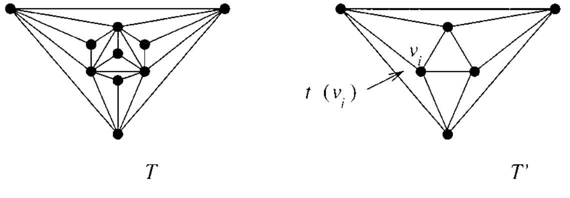

Fig. 8. T' is obtained from T by removing vertices of degree 3.

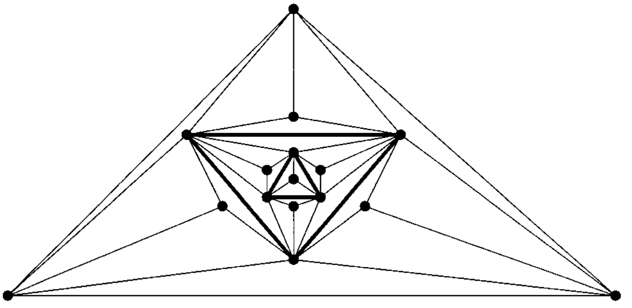

Fig. 9. Only $(n-2)/4$ edges, the thick ones, are flippable.

T, that is, $|S| \leq n - |S| - 2$, i.e., $|S| \leq (n-2)/2$. Using this inequality, the number of flippable edges is

$$(3n-3) - 3|S| - 2(n - |S|) = n - |S| - 3 \geq (n-4)/2,$$

and this concludes the first part of our proof.

We now show that our bound is tight. We give two different examples. Our first example is obtained as follows: Take any collection of m points that are the vertices of a convex polygon Q_m on the plane and any triangulation T of it. Next, add to the interior of each triangle of T an extra vertex adjacent to its three vertices. We obtain a triangulation of a set with $2m-3$ points such that the only edges that can be flipped are the $m-3$ internal edges of T. If $n = 2m-2$, then $m-3 = (n-4)/2$.

For an example with only three points on the convex hull, see Fig. 9. Clearly the example can be generalized to any n of the form $6k+4$. □

The same problem can be posed for triangulations of polygons, taking the number of reflex vertices as a parameter. Observe that the indegree of a convex vertex is zero, and the indegree of a reflex vertex is at most two. Hence, the same method as in the previous proof gives the following result (actually, it is easy to construct examples where the bound below is tight):

Theorem 4.3. *Any triangulation of a polygon Q_n with n vertices, k of them being reflex, contains at least $n-3-2k$ diagonals that can be flipped.*

Now we turn to the problem of bounding the diameter of $G_T(P_n)$ for a point set P_n. First we give a lower bound.

Theorem 4.4. *There are collections P_{2n} of $2n$ points on the plane such that the diameter of $G_T(P_{2n})$ is greater than $(n-1)^2$.*

Proof. Let P_{2n} be the set of vertices of the polygon Q_{2n} presented in Section 3. Notice that any triangulation of P_{2n} will necessarily include the edges of Q_{2n}. Our result now follows by extending the triangulations of Q_{2n} at distance $(n-1)^2$ to triangulations of $Conv(P_{2n})$. □

The proof of Theorem 3.2 extends easily showing that the graphs $G_T(P_n)$ are connected and have diameter $O(n^2)$. However, we present a finer bound on the diameter, in the spirit of Theorem 3.3. Given a point set, the *convex layers* are obtained as follows: remove the convex hull (the first layer) and repeat the operation with the remaining point set, until no point is left. We next show that the diameter of $G_T(P_n)$ is sensitive to the number of convex layers.

Theorem 4.5. *Let P_n be a collection of n points on the plane, and let k be the number of convex layers in P_n. Then the diameter of $G_T(P_n)$ is $O(kn)$.*

The proof of the theorem needs the following lemma.

Lemma 4.6. *Let T be any triangulation of P_n. Then the edges of the second convex layer can be inserted in T using $O(n)$ flips.*

Proof. Let $C_1 = Conv(P_n)$ and let C_2 be the first and second convex layers, respectively. Observe that an edge of T can cross C_2 at most twice. If it crosses C_2 twice, then the two endpoints belong to C_1, and if it crosses C_2 only once, then exactly one of the endpoints belong to C_1. Let s be the number of edges that cross C_2 twice. We first show that they can all be removed with s flips (clearly $s < n$). Indeed, let uv be an edge of T crossing C_2 twice, and let w be a vertex such that uvw is a triangle of T. If w is not in C_1, the we can flip uv obtaining a new edge that crosses C_2 at most once. Otherwise, either uw or vw crosses C_2 twice. As in the proof of Lemma 3.7, we iterate the process until we find a suitable edge to flip.

We can now assume that no edge of T crosses C_2 twice. For every edge e of C_2, let P_e be the polygon formed by the union of all the triangles that cross e. Apply Lemma 3.4 to P_e, to insert e without creating new crossings with C_2. We conclude that C_2 can be inserted with a linear number of flips. □

Proof of Theorem 4.5. Let T and T' be any two triangulations of P_n. The above lemma says that we can insert the second convex layer of P_n into T and T' using $O(n)$ flips obtaining two new triangulations T_1 and T_1'. Let $C_1 = \{v_1, \ldots, v_q\}$ and $C_2 = \{u_1, \ldots, u_p\}$ be the first and second convex layers and assume, without loss of generality, that u_1v_1 does not cross C_2.

We can retriangulate the polygon between C_1 and C_2 as follows. Since u_1v_1 behaves as a proper diagonal, we can insert it both in T_1 and T_1' with $O(n)$ flips using Lemma 3.7. Then $Q = v_1v_2 \cdots v_q v_1 u_1 u_p \cdots u_2 u_1 v_1$ is a spiral polygon. By Lemma 3.5, the triangulations induced by T_1 and T_1' in Q can be transformed into each other with $O(n)$ flips. Finally repeat the process inside C_2 and since k is the number of convex layers, the result follows. □

5. Final Remarks

To conclude, we remark that it is possible to give a proof similar to that of Theorem 3.2 to show that the graph of triangulations of a polygon with holes, or more generally of *planar straight line graphs* as defined in [2], is connected and the diameter is at most quadratic. The details are different in several respects, but to avoid being repetitive we omit the proof.

We remark here for readers familiar with *regular* triangulations [11] that our results are for arbitrary triangulations of point sets, not for regular triangulations. We recall that regular triangulations are known to have at least $n - 3$ flips; moreover, some of the flips allowed for regular triangulations are not allowed in our case.

A problem that has received attention in the past is to compute or to approximate a shortest path between two triangulations of the same point set using flips [7], [14]. Our work in this paper is combinatorial in nature, however, our lower bound examples provide worst cases for such algorithms.

Finally, as an open problem, it would be interesting to improve the bound in Theorem 4.5 and to obtain, for instance, a bound like $O(n + k^2)$.

References

1. D. Avis, Generating rooted triangulations without repetitions, *Algorithmica* 16 (1996), 618–632.
2. M. Bern and D. Eppstein, Mesh generation and optimal triangulation, in *Computing in Euclidean Geometry*, 2nd edition, D.Z. Du and F.K. Hwang, eds., World Scientific, Singapore (1995), pp. 47–123.
3. H. Edelsbrunner and N.R. Shah, Incremental topological flipping works for regular triangulations, *Algorithmica* 15 (1996), 223–241.
4. H. Edelsbrunner, T. Seng Tan, and R. Waupotitsch, An $O(n^2 \log n)$ time algorithm for the minmax angle triangulation, *SIAM J. Sci. Statist. Comput.* 13 (1992), 994–1008.
5. S. Fortune, Voronoi diagrams and Delaunay triangulations, in *Computing in Euclidean Geometry*, 2nd edition, D.Z. Du and F.K. Hwang, eds., World Scientific, Singapore (1995), pp. 225–265.
6. S. Hanke, Ebene Triangulationen, Master's Thesis, Insitut für Informatik, University of Freiburg (1994).
7. S. Hanke, T. Ottmann, and S. Schuierer, The edge-flipping distance of triangulations, *J. Universal Comput. Sci.* 2 (1996), 570–579.
8. F. Hurtado and M. Noy, The graph of triangulations of a convex polygon, Report MA2-IR-94-13, Universitat Politècnica de Catalunya (1994) (an abstract has appeared in *Proc. Symp. on Computational Geomextry* (1996), pp. C7–C8). To appear in *Comput. Geom. Theory Appl.*
9. C.L. Lawson, Software for C_1 surface interpolation, in *Mathematical Software III*, J. Rice, ed., Academic Press, New York (1977), pp. 161–194.
10. C.W. Lee, The associahedron and triangulations of the n-gon. *European J. Combin.* 10 (1989), 551–560.
11. C.W. Lee, Regular triangulations of convex polytopes, in *Applied Geometry and Discrete Mathematics*, P. Gritzmann and B. Sturmfels, eds., DIMACS Series 4 (1991), pp. 443–456.
12. J. Lucas, The rotation graph of binary trees is Hamiltonian, *J. Algorithms* 9 (1988), 503-535.
13. M. Pocchiola and G. Vegter, Topologically sweeping visibility complexes via pseudotriangulations, *Discrete Comput. Geom.* 16(4) (1996), 419–453.
14. R.O. Rogers, Finding paths in the rotation graph of binary trees, Ph.D. Dissertation, Department of Computer Science, University of Central Florida (1996).
15. D.D. Sleator, R.E. Tarjan, and W.P. Thurston, Rotations distance, triangulations and hyperbolic geometry, *J. Amer. Math. Soc.* 1 (1988), 647–682.

Received May 13, 1997, *and in revised form July* 21, 1998, *and February* 1, 1999.