

Triangulating Simple Polygons: Pseudo-Triangulations

Andranik Mirzaian

Technical Report No. CS-88-12

August 1988

Department of Computer Science
York University
North York, Ontario
Canada M3J 1P3

ABSTRACT

Triangulating a given n -vertex simple polygon means to partition the interior of the polygon into $n - 2$ triangles by adding $n - 3$ nonintersecting diagonals. Significant theoretical advances have recently been made in finding efficient polygon triangulation algorithms. However, there is substantial effort being made in finding a *simple* and *practical* triangulation algorithm. We propose the concept of *pseudo-triangulation* (a generalized version of triangulation in which the member triangles need not all have the same orientation), and explore some of its combinatorial and topological properties. Some of the main results of this paper are: (1) We prove the *triangulation-flip-graph* of a simple polygon is connected. Using this theorem we obtain a very simple linear-time algorithm to recognize whether a given triangulation of a simple polygon is its unique triangulation. (2) We prove the maximum diameter of the triangulation-flip-graph is $\Theta(n^2)$. (3) We prove the *Spin-Number Theorem* on simple polygons; an interesting topological result. (4) We propose a triangulation heuristic that uses the *angular (deficit) indices*, and the chord-flip operation, in a local search to transform an initial pseudo-triangulation (which is easy to construct) into a triangulation. The significant open problem with this regard is finding an effective criterion in further refinement of the heuristic regarding the selection of the chord in the chord-flip operation.

key words. simple polygon, triangulation, pseudo-triangulation, chord-flip operation, spin number, angular (deficit) index.

AMS(MOS) subject classifications. 51M15, 68Q25.

1. Introduction

Let P be a simple n -vertex polygon ($n \geq 3$) defined by the list of its vertices v_0, v_1, \dots, v_{n-1} in *positive orientation* around the boundary. Let us assume the boundary of P is denoted ∂P in positive orientation as well. (To be concrete, we may assume

This work was partly supported by Natural Sciences and Engineering Research Council of Canada grant OGP0005516.

An earlier version of this paper appeared in [20].

positive orientation is counter-clockwise, and negative orientation is clockwise. If we walk on ∂P in positive orientation the interior of P in the immediate neighborhood will be on our left hand side.) We define $v_n := v_0$. For convenience we will assume that no three vertices of P are collinear. We call this the *general position property*. (After we describe our methodology, it will become clear how to remove this assumption without significant adjustments.) The *edges* of P are the open line segments with endpoints v_i, v_{i+1} for $0 \leq i < n$. A *chord* of P is an open line segment whose endpoints are two nonadjacent vertices of P . A *diagonal* of P is a chord of P that has empty intersection with the exterior of P . Note that by the general position property, a diagonal of P does not intersect ∂P either. The *triangulation problem* is to find $n-3$ nonintersecting diagonals of P , which partition the interior of P into $n-2$ triangles. It is a known fact that any simple polygon has at least one triangulation. (See Figure 1.)




Figure 1. A triangulated simple polygon.

For a simple polygon, in general, determining whether an arbitrary given chord is also a diagonal is a nontrivial task. However, if P is convex, then any chord is also a diagonal. This fact allows us to find a triangulation of a convex polygon in $O(n)$ time quite easily. Linear time triangulation algorithms are known for other special cases as well, such as for monotone polygons [7], and star-shaped polygons [12, 25]. The first nontrivial triangulation algorithm for the general case was proposed in 1978 by Garey, Johnson, Preparata, and Tarjan [7]. This algorithm takes $O(n \log n)$ time. Since then, one of the most outstanding open problems in computational geometry has been whether the triangulation problem can be solved in linear time. In the recent years substantial amount of research effort has gone into resolving this open problem. Some researchers have devised triangulation algorithms with running time $O(n \log k)$, for a parameter k that measures the complexity of the polygon, such as the number of reflex angles [12], or the sinuosity [2]. Since all these measures admit classes of polygons with $k = \Theta(n)$, the worst case running time of these algorithms is only known to be $O(n \log n)$. Fournier and Montuno [6] showed the triangulation problem is linear time equivalent to finding all vertex-edge horizontally visible pairs (or, equivalently, computing the horizontal visibility subdivision). The reduction of triangulation to computing the horizontal visibility information was independently obtained by Chazelle and Incerpi [2]. Since then, almost all researchers in the field have chosen this approach, namely, computing the horizontal visibility subdivision. Tarjan and Van Wyk [24] devised an $O(n \log \log n)$ time triangulation algorithm. This is a major theoretical breakthrough and shows that triangulation is easier than sorting. Tarjan and Van Wyk's algorithm uses the divide-and-conquer paradigm and computes the horizontal visibility information, using an adaptation of the Jordan sorting

algorithm [14] and homogeneous and heterogeneous finger search trees at two different levels of the algorithm. As a result, the computational overhead of their algorithm is very high. (Clarkson, Tarjan, and Van Wyk [4] have adapted the above algorithm and used random sampling [3] to devise a randomized triangulation algorithm with $O(n \log^* n)$ expected time.) Very recently, Kirkpatrick, Klawe, and Tarjan [16] propose a somewhat simpler algorithm than [24]. This shows that algorithmic simplicity is still an important issue regarding the polygon triangulation problem, but it is still a problem to be pursued vigorously.

In this paper we propose the new concept of *pseudo-triangulation*, a generalized version of triangulation, in which the member triangles need not all have the same orientation. We explore some combinatorial and topological properties of pseudo-triangulations. Our methodology in solving the triangulation problem may be viewed as a graph search method where the underlying graph is the *pseudo-triangulation-flip-graph* of P (or *flip-graph* of P , for short).

The remaining portion of this paper is organized as follows. Section 2 introduces the concept of pseudo-triangulations and shows some basic facts about it. Included in this section are the introduction of the characteristic function of a simple polygon and the proof of its additive property with respect to pseudo-triangulations (a fact which is used in subsequent sections), and a necessary and sufficient condition for a pseudo-triangulation to be a triangulation. Section 3 proves a connectivity theorem on triangulation-flip-graphs, one of whose corollaries is that in linear time we can decide whether a given triangulated polygon has a unique triangulation. The connectivity theorem of this section might have other applications such as in shortest paths and visibility problems. We also show that the maximum possible diameter of the triangulation-flip-graph is $\Theta(n^2)$. (However, the diameter of the flip-graph is known to be $\Theta(n)$ due to [23].) The main result of Section 4 is the Spin-Number Theorem. Section 5 introduces the notion of angular indices (integer weights on the vertices of the polygon) which provides yet another characterization of triangulations via pseudo-triangulations. This leads us to propose a generic method for the solution of the triangulation problem. Section 6 makes some concluding remarks and poses some open problems.

2. Pseudo-Triangulations

Consider two chords (v_i, v_j) and (v_k, v_l) of P . Without loss of generality assume $i < j$ and $k < l$. We say the two chords *interlace* if $i < k < j < l$ or $k < i < l < j$. Below, we will use some terminology from combinatorial topology. (See for example [1, 21].) Let T be a two dimensional simplicial-complex consisting of $n - 2$ triangles (or 2-simplices), $2n - 3$ segments (or 1-simplices), and n vertices (or 0-simplices). We say T is a *pseudo-triangulation* of P if it satisfies condition (i) below.

- (i) *the combinatorial property*: The vertices of T are the vertices of P . Each segment of T is either an edge or a chord of P . A segment in T is incident* to exactly one triangle in T if it is an edge of P , and it is incident to exactly two triangles in T if it is a chord of P . In the latter case we may call the segment a chord of T . If, in addition, it is a diagonal of P , we may call it a diagonal of T as well. Furthermore, T has no pair of chords that interlace.

It should be obvious that any triangulation of P is also a pseudo-triangulation of P . Let T be a pseudo-triangulation of P . We assign an orientation to each triangle T_j of T as follows. We consider ∂T_j to have the orientation on which its three vertices are seen in the same order as on ∂P . (Note, by the general position property the three vertices are not collinear.) Then T_j inherits the same orientation as ∂T_j . (See Figure 2.) As we shall see in Theorem 5 below, a pseudo-triangulation T of P is a triangulation of P if and only if it satisfies condition (ii) below.

* two simplices are *incident* if one is a face of the other.

(ii) *the topological property*: All triangles in T have the same orientation as P .




Figure 2. A pseudo-triangulation of the polygon (v_0, \dots, v_5) . Triangles (v_0, v_1, v_2) , (v_2, v_3, v_4) and (v_0, v_4, v_5) have positive orientation. Triangle (v_0, v_2, v_4) has negative orientation.

To help the intuition, there is a second way to define pseudo-triangulations. Let C be a *convex* n -vertex polygon, where ∂C contains the list of vertices u_0, u_1, \dots, u_{n-1} in positive orientation. Then the triangulations of C are in one-to-one correspondence with the pseudo-triangulations of P with the following correspondence: A triangulation T' of C corresponds to a pseudo-triangulation T of P if (u_i, u_j) is a diagonal of T' if and only if (v_i, v_j) is a chord of T . Let us call this correspondence the *natural mapping*. (See Figure 3.)




Figure 3. A pseudo-triangulation and its natural mapping. The sign in triangle (u_i, u_j, u_k) indicates the orientation of triangle (v_i, v_j, v_k) .

Now we define an elementary operation called *flip* which transforms one pseudo-triangulation of P to another. Let T be a pseudo-triangulation of P and (v_i, v_j) a chord of T incident to two triangles T_1 and T_2 of T . Let v_k and v_l be the other two vertices of T_1 and

T_2 . We say the chord (v_k, v_l) of P is the *dual* of (v_i, v_j) with respect to T . A *flip* is the operation of replacing a chord by its dual in a pseudo-triangulation. This operation is obviously reversible. If the flip operation is applied to convex polygons, it transforms one triangulation of the polygon to another. We also define the *flip-graph* of P to be the graph where its nodes correspond to pseudo-triangulations of P , and two nodes of the graph are adjacent if the corresponding pseudo-triangulations can be obtained one from the other by a single flip operation. The flip-graphs are isomorphic to the so called *rotation graphs* that have been studied in the literature [5, 18, 23]. It is known that these graphs are Hamiltonian and have diameter $O(n)$.

Let Q be a simple polygon with vertex set U in the plane R^2 . We define the *characteristic function* $\chi_Q : R^2 - U \rightarrow \{ 0, \pm 1/2, \pm 1 \}$ of Q as follows. For a point $p \in R^2 - U$ we define the magnitude of $\chi_Q(p)$ to be 0 if p is in the exterior of Q , $1/2$ if p is on $\partial Q - U$, and 1 if p is in the interior of Q . The sign of $\chi_Q(p)$ is defined to be positive or negative if the orientation of Q is, respectively, positive or negative. We leave $\chi_Q(p)$ undefined if p is a vertex of Q .

Theorem 1. *Let T be a pseudo-triangulation of P which contains the triangles T_j , for $1 \leq j \leq n - 2$. Let V denote the set of vertices of P . Then for every point $q \in R^2 - V$ the following identity holds:*

$$\chi_P(q) = \sum_{j=1}^{n-2} \chi_{T_j}(q)$$

Proof: We prove the above identity by showing that the flip operation leaves the value of the right hand side of the equation invariant. Suppose (without loss of generality) that the triangles T_1 and T_2 of T are replaced by two new triangles T_1' and T_2' by a flip. A case analysis easily shows that $\chi_{T_1}(q) + \chi_{T_2}(q) = \chi_{T_1'}(q) + \chi_{T_2'}(q)$. (See Figure 4.) The proof is complete by the two facts that the flip-graph of P is connected and that the equation obviously holds if T is a triangulation of P . \square

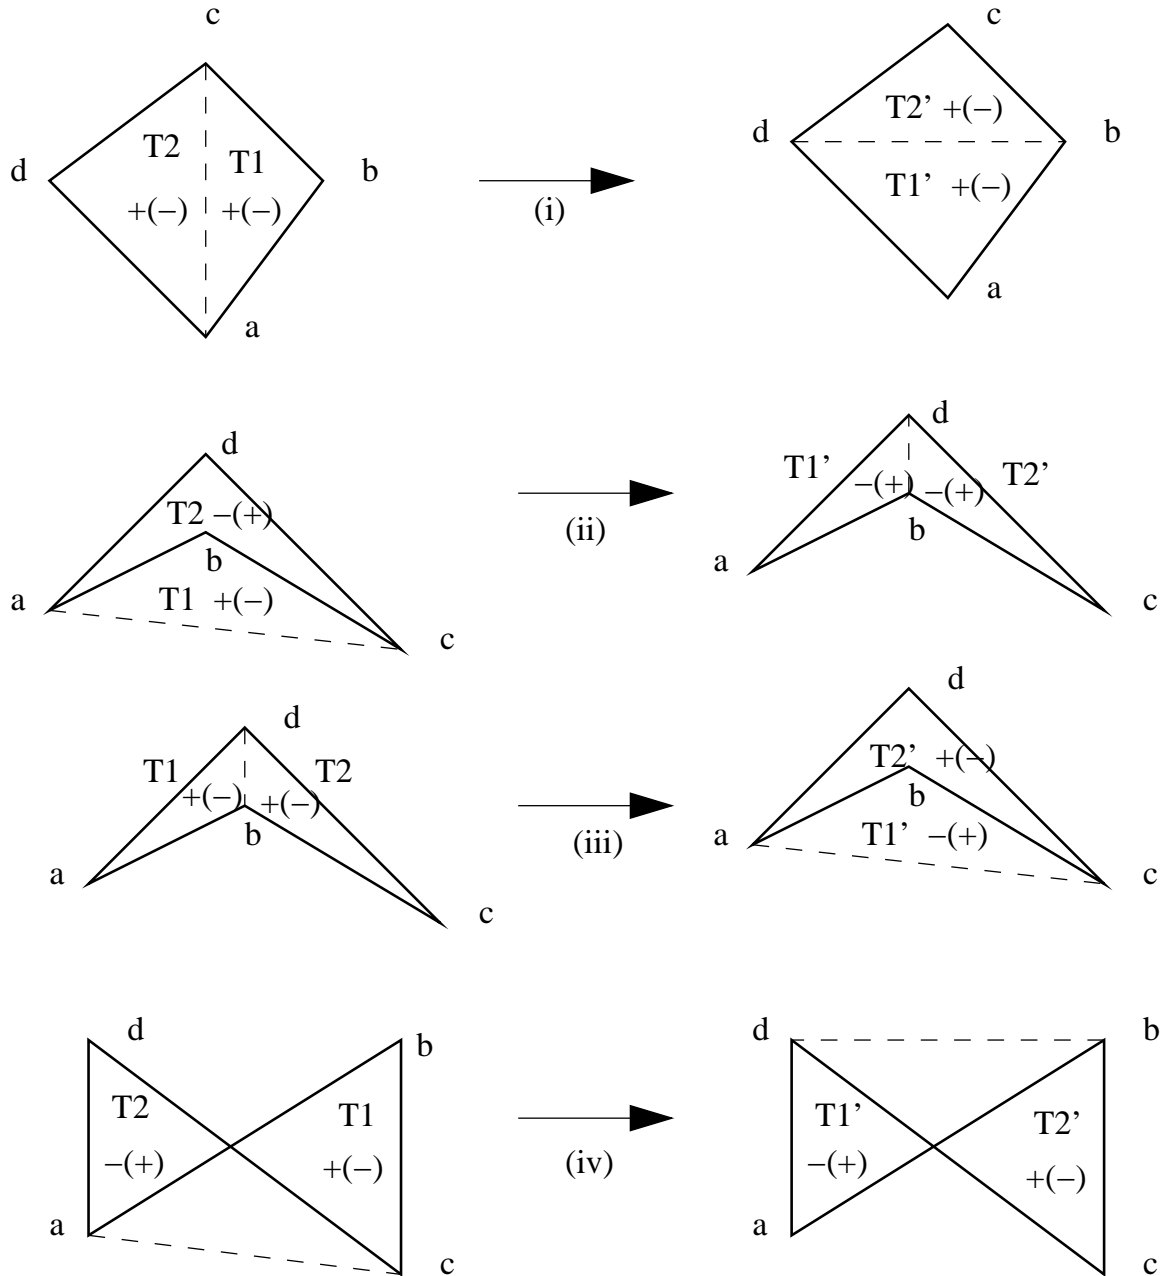


Figure 4. The flip operation. Case (i): no overlap; cases (ii) and (iii): total overlap; case (iv): partial overlap.

For a simple polygon Q , let $area(Q)$ denote the signed area of Q (with the Euclidean metric), whose magnitude is the area of Q and whose sign is the orientation of Q . There is an analog of Theorem 1 about signed areas; the related identity is:

$$area(P) = \sum_{j=1}^{n-2} area(T_j)$$

A proof similar to that of Theorem 1 can be used to prove the area identity. (For an alternative proof of a special case see [17].) We will use the characteristic function again later.

Lemma 2. *In any pseudo-triangulation T of P no pair of diagonals cross each other.*

Proof: Otherwise under the natural mapping we would get a triangulation of the convex polygon C with a pair of crossing diagonals, a contradiction. \square

Lemma 3 (*The Enclosure Lemma*). *Suppose R and Q are two simple polygons such that Q does not intersect the exterior of R and has at least three vertices in common with R and its remaining vertices are in the interior of R . Then the common vertices of R and Q appear in the same order around ∂R and ∂Q if R and Q have the same orientation, and appear in reverse order if R and Q have opposite orientation.*

Proof: We use mathematical induction on the number of noncommon vertices of R and Q . Suppose the vertices common to both R and Q are w_0, w_1, \dots, w_{k-1} in the order around ∂Q . Let $w_k := w_0$. First assume there is a vertex of either R or Q which is not a vertex of the other. This implies that there exists an index i , $0 \leq i < k$, so that the portion of ∂Q from w_i to w_{i+1} is not a subset of ∂R . Let the oriented polygonal chain Π be the portion of ∂Q from w_i to w_{i+1} . Π partitions the interior of R into two simple polygons R_i and L_i , so that R_i is to the right of Π and L_i is to the left. We assume R_i and L_i have the same orientation as P . By the Jordan Curve Theorem, the interior of Q must be either entirely in R_i , or entirely in L_i . If Q is in R_i , then by mathematical induction and the fact that w_i and w_{i+1} appear in the same order around ∂R_i and ∂Q , Q must have the same orientation as R_i and their common vertices appear in the same order around ∂R_i and ∂Q . Since in this case every vertex common to both R and Q is also common to R_i and Q , the inductive step follows. If Q is in L_i , then again by induction and the fact that w_i and w_{i+1} appear in opposite order around ∂L_i and ∂Q , Q must have the opposite orientation to L_i and their common vertices appear in opposite order around ∂L_i and ∂Q . Since in this case every vertex common to both R and Q is also common to L_i and Q , the inductive step follows. The base of the induction is the case when R and Q have the same set of vertices. In this case a similar argument holds in which either R_i or L_i is identical to R , for all $0 \leq i < k$. In other words, in this case R and Q must be the same polygon with either the same or opposite orientation. \square

Corollary 4. *Suppose t is one of the triangles in a pseudo-triangulation of P . If t has negative orientation, then at least one of its sides is neither a diagonal nor an edge of P .*

Proof: If all three sides of t are either edges or diagonals of P , then The Enclosure Lemma applies and therefore t has the same orientation as P , which is positive, a contradiction. \square

A much stronger version of Corollary 4 will be proved in Section 4.

Theorem 5. *Let T be a pseudo-triangulation of P . Then T is a triangulation of P if and only if all triangles of T have positive orientation.*

Proof: If T has a negatively oriented triangle, Corollary 4 implies that T is not a triangulation of P . If all triangles of T have positive orientation, then Theorem 1 implies that the triangles of T cannot intersect the exterior of P ; that no two triangles of T can have common interior points; and that each interior point of P is either in the interior of one triangle of T or on the common chord of two such triangles. In other words, the triangles of T partition the interior of P and hence T is a triangulation of P . \square

An implication of Theorem 5 is that in any pseudo-triangulation T of P at least one triangle is positively oriented. This is because if all the triangles of T are negatively oriented, then T must be a triangulation of the polygon identical to P but in opposite orientation. But this is absurd. It was Theorem 5 and the fact that the diameter of the flip-graph of P is $O(n)$ that was the first motivating factor in our study of pseudo-triangulations. The underlying implication is that for each pseudo-triangulation of P as a starting one, there is a sequence of $O(n)$ flip operations, that converts it to a triangulation of P . Furthermore,

after each step we can determine in $O(1)$ time whether the current pseudo-triangulation is indeed a triangulation of P by simply keeping the count of how many negatively oriented triangles it contains.

3. The Triangulation-Flip-Graph

From the literature, we already know a few facts about the structure of the flip-graph [5, 18, 23]. As we mentioned earlier, we know that the flip-graph is Hamiltonian and has diameter $O(n)$. It is desirable to know more about this graph. As an obvious additional property, we can mention that the flip-graph is triangle free, i.e. has no 3-clique. Let us call the subgraph of the flip-graph induced by the nodes which correspond to triangulations of P the *triangulation-flip-graph* of P . What can we say about the connectivity and the diameter of the triangulation-flip-graph? This graph may not be Hamiltonian. Indeed, it may not even be biconnected and may not contain a Hamiltonian path either. For such an example see Figure 5. However, we have the following facts.

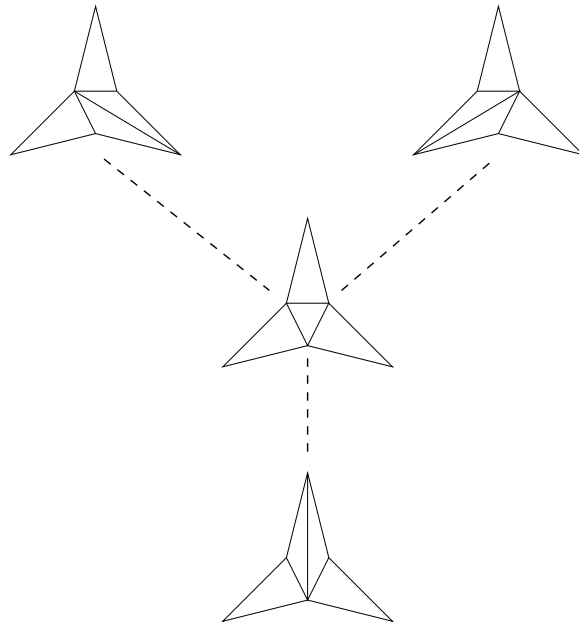


Figure 5. The subgraph of the flip-graph induced by the nodes that correspond to the triangulations of P may not be biconnected, and may not have a Hamiltonian path.

Theorem 6. *The triangulation-flip-graph of P is connected.*

Proof: Assume $n \geq 4$, otherwise there is nothing to prove. An *ear* of a triangulation of P is any of its triangles that has at least two edges of P as its sides. It is well known that when $n \geq 4$ any triangulation of P has at least two ears. (See for example [13, 19].)

Suppose T and T' are two triangulations of P . We will show that there is a sequence of flip operations that converts T to T' in such a way that all intermediate pseudo-triangulations obtained by this sequence are also triangulations of P . Let t be an ear of T' . If t is also an ear of T , then we may discard t and apply the proof to the remaining smaller polygon. So let us assume t is not an ear of T . Suppose the two edges of P that are the sides of t are (v_{i-1}, v_i) and (v_i, v_{i+1}) . Suppose the other ends of the diagonals in T incident to v_i are w_1, w_2, \dots, w_{k-1} in order around ∂P . Let $w_0 := v_{i+1}$ and $w_k := v_{i-1}$. We claim that there is an index j , $1 \leq j < k$, such that the quadrangle $v_i, w_{j-1}, w_j, w_{j+1}$ is convex. Since t is an ear of T , none of the vertices w_1, w_2, \dots, w_{k-1} can be in t , otherwise they

would prevent the visibility between v_{i-1} and v_{i+1} . Therefore, w_1, w_2, \dots, w_{k-1} are on the opposite side of the line through v_{i-1} and v_{i+1} with respect to v_i . Among the vertices w_1, w_2, \dots, w_{k-1} let w_j be orthogonally farthest from the line through v_{i-1} and v_{i+1} . Then, obviously, the quadrangle $v_i, w_{j-1}, w_j, w_{j+1}$ is convex, as desired. Therefore, (w_{j-1}, w_{j+1}) is a diagonal of P . (See Figure 6.) If we flip (v_i, w_j) to replace it with (w_{j-1}, w_{j+1}) , we will obtain a new triangulation of P in which the number of diagonals incident to v_i has decreased by one. If we continue this process, through a sequence of flip operations which produces only triangulations, eventually t becomes an ear of T . Now, we can discard t and apply the same argument on the remaining portion of the polygon. \square




Figure 6. The figure used in the proof of Theorem 6.

Corollary 7. *Given a triangulation T of P , we can determine whether T is the unique triangulation of P in $O(n)$ time.*

Proof: Theorem 6 implies that T is the unique triangulation of P if and only if it does not have any diagonal whose two incident triangles together form a convex quadrangle. The latter condition can easily be checked in $O(1)$ time for each diagonal of T . \square

For an example of a simple polygon with a unique triangulation see Figure 7.




Figure 7. A simple polygon with a unique triangulation.

Theorem 8. *The maximum diameter of the triangulation-flip-graph of any simple polygon with n vertices is $\Theta(n^2)$.*

Proof: Let P be a simple n -vertex polygon and $D(P)$ denote the diameter of its triangulation-flip-graph. From the proof of Theorem 6 it follows that $D(P)$ is $O(n^2)$ since we need less than n flips for each "ear removal". The fact that $D(P)$ could be as large as $\Omega(n^2)$ is apparent from the simple example in Figure 8. The number of diagonal flips needed to convert the triangulation in Figure 8 (a) to the triangulation shown in Figure 8 (b), in such a way that all intermediate pseudo-triangulations are actually triangulations, is $(\lfloor n/2 \rfloor - 1)(\lceil n/2 \rceil - 1)$. To see this, it suffices to notice that after $n - 3$ diagonal flips the instance is converted to one with 2 less vertices. (An ear removed from each side.) The reason why this is the minimum number of flips necessary follows from the fact that at each step there are at most two flips possible; one on each side of the "middle diagonal". \square




Figure 8. Two triangulations of an hourglass polygon.

4. Spin Numbers

Let O be a point outside the convex hull of P . Let Γ be an open oriented simple polygonal chain from point O to vertex a of P such that it does not intersect P . Furthermore, we make the simplifying assumption that Γ has only finitely many intersection points with any chord of P , with all such intersections being crossing intersections (i.e., no tangential intersections). We call Γ a *probe* to vertex a . Let t be the triangle with vertices a, b and c of P seen in that order around ∂P . We define the *spin* at vertex a of

triangle t with respect to Γ , denoted $\psi_a(t; \Gamma)$, to be the algebraic number of times Γ intersects (b, c) , i.e. the side of t opposite vertex a , where an intersection is counted +1 if at that point Γ enters t , and counted -1 if at that point Γ exits t . (See Figure 9.) The spins at vertices b and c are defined similarly.




Figure 9. The spin at vertex a of triangle $t = (a, b, c)$ with respect to the probe Γ ; $\psi_a(t; \Gamma) = +2$.

Lemma 9. Let Γ and Γ' be two probes to vertex a of triangle t . Then $\psi_a(t; \Gamma) = \psi_a(t; \Gamma')$.

Proof: There exists a third probe Γ'' to a that has no points in common with Γ or Γ' . Let Q be the simple polygon whose boundary is Γ followed by a , followed by the reverse of Γ'' , followed by O . Vertices b and c are in the exterior of Q . As we move from b to c along the chord (b, c) count the algebraic number of times we move in and out of Q , counting +1 as we move into Q and count -1 as we move out of Q . by the Jordan Curve Theorem, the magnitude of that count is $|\psi_a(t; \Gamma) - \psi_a(t; \Gamma'')| = 0$. With a similar argument we have $\psi_a(t; \Gamma') - \psi_a(t; \Gamma'') = 0$. The lemma follows. \square

Lemma 9 says the spin numbers at vertices of a triangle are a property of the triangle itself (and of P of course) and do not depend on which particular probes we choose. Because of this we may use the shorter notations $\psi_a(t)$, $\psi_b(t)$ and $\psi_c(t)$. (See Figure 10.)




Figure 10. The spin numbers at vertices of triangle $t = (a, b, c)$; $\psi_a(t) = +1$, $\psi_b(t) = +1$ and $\psi_c(t) = -1$.

Remark. We could have defined the spin numbers in terms of *winding numbers*. For instance, $\psi_a(t)$ is the winding number of point a with respect to the closed curve that is formed by going from b to c along ∂P followed by the segment from c back to b .

Suppose x is a vertex of P but not of t . Let $\alpha_x(t)$ denote the algebraic number of times a probe to vertex x intersects ∂t ; again, counting an intersection $+1$ if the probe enters t at that point, and -1 otherwise. By the Jordan Curve Theorem we know $\alpha_x(t)$ is 0 if x is outside t , and is $+1$ if x is inside t . We have the following fact about spin numbers:

Theorem 10 (*The Spin-Number Theorem*). *Let t be a triangle with three vertices a, b and c of P in that order around ∂P . Then the quantity $\psi_a(t) + \psi_b(t) + \psi_c(t)$ is 0 if t is positively oriented, and is $+1$ if t is negatively oriented.*

Proof: Suppose $a = v_i$, $b = v_j$ and $c = v_k$. Consider the triple $(j-i, k-j, i-k)$ where index arithmetic is done modulo n (i.e., add n to the result of the subtraction if it is negative). We consider all such possible triples in lexicographic order. So, without loss of generality, assume $j-i \leq k-j \leq i-k$.

Case (i): Suppose $1 = j-i = k-j \leq i-k$. That is, (a, b) and (b, c) are edges of P . Since probes by definition do not intersect P , we must have $\psi_a(t) = \psi_c(t) = 0$. Triangle t is negatively oriented if the internal angle at vertex b of P is reflex, and t is positively oriented otherwise. Now consider a probe Γ to vertex b . As we move along Γ sufficiently close to b (where distance is measured along Γ), we will be inside triangle t if t is negatively oriented, and we will be outside t otherwise. Since the origin of Γ is outside t and the only way Γ can intersect the boundary of t is by intersecting the chord (a, c) , we conclude $\psi_b(t)$ must be $+1$ if t is negatively oriented, and must be 0 if t is positively oriented. (See Figure 11 (a) and (b).)

Case (ii): Now suppose $1 = j-i < k-j \leq i-k$. In this case (a, b) is an edge of P . Therefore, $\psi_c(t) = 0$. Let Γ be a probe to a . Let γ be a point on Γ so that the distance

from γ to a along Γ is sufficiently small. Let Γ' be the probe to b which consists of the portion of Γ from O to γ followed by point γ , followed by the open segment from γ to b . If γ is inside t , then so is the entire segment from γ to b . If γ is outside t , then the segment from γ to b is either entirely outside t or it intersects the boundary of t only once with the intersection being on the chord (a, c) . If the segment (γ, b) , in the sufficiently small neighborhood of b , is inside t , then by the Jordan Curve Theorem the algebraic number of times Γ' intersects the boundary of t is $+1$. Since no such intersection can occur on (a, b) , we conclude $\psi_a(t; \Gamma) + \psi_b(t; \Gamma') = +1$. With a similar reasoning we conclude that $\psi_a(t; \Gamma) + \psi_b(t; \Gamma') = 0$ if the segment (γ, b) , in the sufficiently small neighborhood of b , is outside t . Furthermore, Since the interior of the triangle (a, b, γ) has no point in common with P , we conclude that the orientation of t is positive if and only if the segment (γ, b) , in the sufficiently small neighborhood of b , is outside t . (See Figure 11 (c) and (d).)

Case (iii): Now suppose $1 < j - i \leq k - j \leq i - k$. Let $d = v_{i+1}$. Consider the triangles $t_1 = (a, d, b)$, $t_2 = (a, d, c)$, and $t_3 = (d, b, c)$. The theorem holds for t_1 and t_2 due to case (ii) above. Also, since triangle t_3 is lexicographically smaller than t , the theorem is assumed proven for t_3 already. Triangle t might be positively or negatively oriented. For each of these two cases there are seven sets of possibilities (fourteen cases in all) for the location of point d as shown in Figure 11 (e) and (f). We will prove the theorem for case 1 as shown in Figure 11 (e). The proof of the other cases, being similar, are left to the reader. In case 1 triangles t and t_2 are positively oriented, and triangles t_1 and t_3 are negatively oriented. For this case we have the following equations:

$$0 = \alpha_a(t_3) = +\psi_a(t_1) + \psi_a(t_2) - \psi_a(t) \tag{1}$$

$$0 = \alpha_b(t_2) = -\psi_b(t_1) + \psi_b(t_3) + \psi_b(t) \tag{2}$$

$$0 = \alpha_c(t_1) = -\psi_c(t_2) + \psi_c(t_3) + \psi_c(t) \tag{3}$$

$$0 = \alpha_d(t) = +\psi_d(t_1) + \psi_d(t_2) - \psi_d(t_3) \tag{4}$$

Let us use the abbreviation $\Psi(t') := \psi_u(t') + \psi_v(t') + \psi_w(t')$ for any triangle $t' = (u, v, w)$. From equations (1)–(4) we see that $-\alpha_a(t_3) + \alpha_b(t_2) + \alpha_c(t_1) - \alpha_d(t) = 0 = \Psi(t) - \Psi(t_1) - \Psi(t_2) + \Psi(t_3)$. Since $\Psi(t_1) = +1$, $\Psi(t_2) = 0$, and $\Psi(t_3) = +1$, we conclude $\Psi(t) = 0$. \square

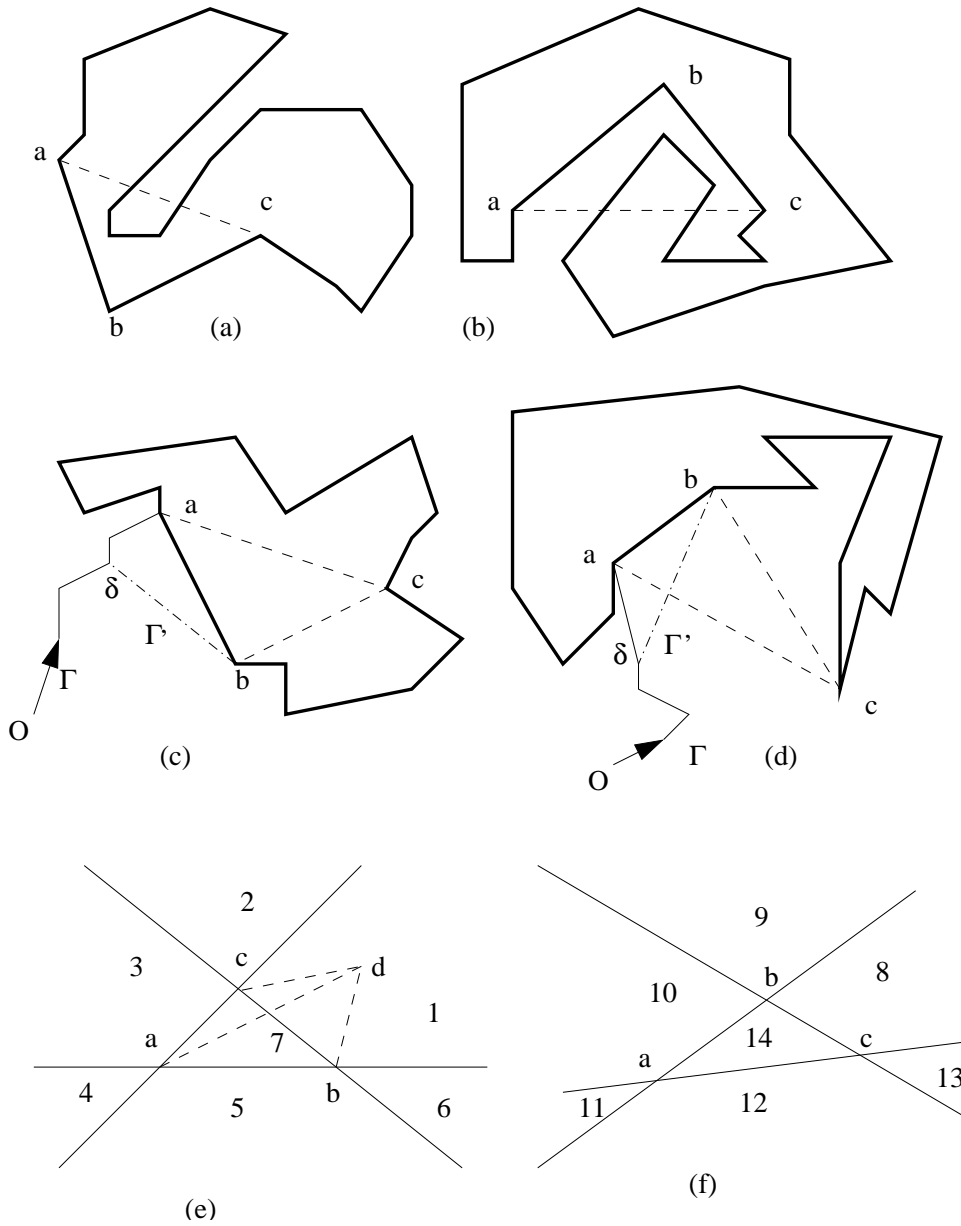


Figure 11. The cases used in the proof of Theorem 10.

As a first application, we see that Corollary 4 is now an obvious consequence of Theorem 10. This is because if triangle t , with vertices a, b and c , is negatively oriented, then by the theorem at least one of $\psi_a(t)$, $\psi_b(t)$ or $\psi_c(t)$ is positive. Say $\psi_a(t)$ is positive. Then (b, c) is neither an edge nor a diagonal of P , since it must intersect any probe to vertex a .

Our initial attempt was to assign some kind of integer weights to the vertices of each triangle in a pseudo-triangulation of P with the aim of using these weights as a guide for the selection of a suitable chord to flip. It was this attempt that led us to the definition of spin numbers and the discovery of the Spin-Number Theorem. However, the effective use of spin numbers is still open. The difficulty comes in updating the spin numbers after a flip operation is performed. Suppose the chord (a, c) is flipped and replaced by its dual (b, d) . Before the flip operation takes place the two triangles incident with the chord (a, c) are $t_1 = (a, b, c)$ and $t_2 = (a, c, d)$. After the flip operation the two new

triangles incident with the new chord (b, d) are $t_3 = (a, b, d)$ and $t_4 = (b, c, d)$. Ideally, we would want the six new spin numbers of the vertices of the triangles t_3 and t_4 to be completely determined by the six old spin numbers of the triangles t_1 and t_2 and (the positions of) the vertices a, b, c, d . If this were so, then the updating could be done in $O(1)$ time; however, this is not the case. (See Figure 12.) In the next section we will define an alternative notion of integer weights at vertices which seem to be more promising.




Figure 12. The triangles are $t_1 = (a, b, c)$ and $t_2 = (a, c, d)$ before the flip, and $t_3 = (a, b, d)$ and $t_4 = (b, c, d)$ after the flip. In both cases (i) and (ii) prior to the flip operation we have $\psi_a(t_1) = 1, \psi_b(t_1) = 0, \psi_c(t_1) = 0, \psi_a(t_2) = 1, \psi_c(t_2) = 0, \psi_d(t_2) = -1$. However, after the flip operation in case (i) we have: $\psi_a(t_3) = 0, \psi_b(t_3) = 1, \psi_d(t_3) = 0, \psi_b(t_4) = 1, \psi_c(t_4) = 0, \psi_d(t_4) = -1$; but in case (ii) we have: $\psi_a(t_3) = 0, \psi_b(t_3) = 0, \psi_d(t_3) = 1, \psi_b(t_4) = 0, \psi_c(t_4) = 0, \psi_d(t_4) = 0$.

5. Angular (Deficit) Indices

In this section we define a local quantity that shows some promise towards successful application in an ultimate triangulation algorithm. Let $angle(a)$ denote the internal angle of P at vertex a . Suppose t is an oriented triangle with vertices a, b and c . We define the *signed-angle at vertex a of triangle t* , denoted $\omega_a(t)$, so that its magnitude is the internal angle of t at vertex a and its sign is the orientation of t . $\omega_b(t)$ and $\omega_c(t)$ are defined similarly. Therefore, $\omega_a(t) + \omega_b(t) + \omega_c(t)$ equals $+\pi$ radians if t is positively oriented and is $-\pi$ radians if t is negatively oriented. Suppose T is a pseudo-triangulation of P consisting of triangles T_j , for $1 \leq j \leq n-2$. Assume m of these triangles are negatively oriented and the remaining $n-m-2$ are positively oriented. The total

signed-angle sum of these triangles is therefore $\pi (n - 2) - 2m\pi$. This last quantity can be interpreted as follows. The term $\pi (n - 2)$ is the sum of the internal angles of P , and the deficit term $2m\pi$ is due to the fact that T contains m negatively oriented triangles. Suppose (without loss of generality) that the triangles of T incident to vertex a are T_1, T_2, \dots, T_l . Let us define the *signed-angle of vertex a with respect to the pseudo-triangulation T* , denoted $\omega_a (T)$, to be

$$\omega_a (T) := \sum_{i=1}^l \omega_a (T_i) .$$

Because of connectivity we must have

$$\omega_a (T) = \text{angle}(a) + 2 k_a \pi$$

for some *integer* k_a (negative, zero or positive). We call k_a the *angular index* of vertex a with respect to the pseudo-triangulation T . We say, with respect to T , vertex a is *balanced* if $k_a = 0$, has *angular surplus* if $k_a > 0$, and has *angular deficit* if $k_a < 0$. (See Figure 13.)




Figure 13. Angular index of vertex v_1 in this pseudo-triangulation is -1 .

Lemma 11. *Let T be a pseudo-triangulation of P . Then T is a triangulation of P if and only if no vertex of P has angular deficit with respect to T .*

Proof: From above definitions we see that the sum of the angular indices of vertices of P is $-m$ where m is the number of negatively oriented triangles of T . If $m > 0$, then obviously some vertex has angular deficit. If $m = 0$, then by Theorem 5 T is a triangulation of P ; hence, all vertices of P are balanced. \square

Lemma 12. *Let T be a pseudo-triangulation of P . Let k_v be the angular index of vertex v with respect to T . Then there is some natural number α such that $2\alpha - k_v$ triangles of T contain vertex v in their interior; α of these triangles are negatively oriented and $\alpha - k_v$ of these triangles are positively oriented.*

Proof: Suppose the triangles of T that are incident to v are T_i , for $1 \leq i \leq l$, and the remaining triangles of T are T_i , for $l < i \leq n-2$. Let N be a sufficiently small neighborhood of vertex v . (We assume N is small enough that it does not intersect with any chord of P that is not incident to v .) Let p be an arbitrary point in $N \cap \text{exterior} (P)$. Since vertex v has angular index k_v , by connectivity we must have $\sum_{i=1}^l \chi_{T_i} (p) = k_v$. From Theorem 1 we have $\sum_{i=1}^{n-2} \chi_{T_i} (p) = 0$. Therefore, $\sum_{i=l+1}^{n-2} \chi_{T_i} (p) = -k_v$. Let α denote

the number of negatively oriented triangles of T that are not incident to v and contain p in their interior. Since p is not on the boundary of any triangle of T , there must be $\alpha - k_v$ positively oriented triangles of T that are not incident to v and contain p in their interior. Since v is sufficiently close to p , the lemma follows. \square

Corollary 13. *If vertex v has angular deficit with respect to T , then some positively oriented triangle T_j of T contains vertex v in its interior.*

Proof: Obvious consequence of Lemma 12. \square

In the pseudo-triangulation T let us say vertex v is *saturated* if all $n - 3$ chords of T are incident to v .

Corollary 14. *A saturated vertex v in a pseudo-triangulation T must be balanced.*

Proof: Since all triangles of T are incident to v , the proof follows from Lemma 12. \square

Notice that the triangle T_j of Corollary 13, although positively oriented, cannot have all its sides as edges or diagonals of P , since it includes a vertex of P in its interior. Motivated by the above two corollaries, let us define a *chain of triangles* in a pseudo-triangulation T of P to be a sequence T'_0, T'_1, \dots, T'_l of distinct triangles of T , so that T'_{i-1} and T'_i , are incident to a common chord c_i , for $1 \leq i \leq l$. We say the length of this chain is l . A chain, in a pseudo-triangulation, is determined by its initial and final triangles. Let v be the vertex incident to T'_0 but not to T'_1 . We might also refer to the above mentioned chain as the chain from vertex v to triangle T'_l . We say the sequence c_1, c_2, \dots, c_l is the *chain of separating chords* from vertex v to triangle T'_l . Suppose vertex v has an angular deficit and hence by Corollary 13 it is in the interior of a positively oriented triangle T_j of T . Let c_1, c_2, \dots, c_l be the chain of separating chords from vertex v to triangle T_j . Intuitively, we would want to flip c_1 . There are two advantages for doing so. One is that the flip will shorten the length of the chain and eventually "destroy" it. The other is that if we keep performing such flips with respect to vertex v , it will eventually have to become balanced by Corollary 14. The question is, how do we determine which triangle of T is incident to vertex v whose side opposite v is c_1 ?

The effect of a flip operation to the angular indices is strictly local. (This is the desired property that spin numbers, as explained in the previous section, do not possess.) A flip operation can change the angular index of only one of the four vertices involved in the operation, and that will happen if and only if that vertex is in the convex hull of the other three. (See cases (ii) and (iii) of Figure 4.) Furthermore, this change is only additive by the amount $+1$ or -1 .

Consider a pseudo-triangulation T in which one vertex, say v_0 is saturated. Then, by Corollary 14, v_0 is balanced. Since v_1 and v_{n-1} are incident to one triangle each, they are either balanced or have angular deficit index -1 depending on whether the internal angle of P at those vertices is, respectively, convex or reflex. Each of the remaining vertices v_2, \dots, v_{n-2} is incident to two triangles of T . Therefore the signed-angle of those vertices with respect to T must be in the range from -2π to $+2\pi$. This implies that the angular indices of these vertices must be either 0 or -1 . In summary, the angular indices of all vertices with respect to T is 0 or -1 and these can be determined in total $O(n)$ time.

In a pseudo-triangulation T let us call a chord (u, w) a *candidate* chord if there is a vertex v of P such that v has angular deficit and (v, u, w) is one of the triangles in T . Below, we propose our heuristic for the solution of the triangulation problem. (Note that in its generic form, this heuristic is nondeterministic due to line 5.)

Heuristic *triangulate* (P);

0. let T be the pseudo-triangulation in which v_0 is saturated;
1. compute the angular indices of vertices with respect to T ;
2. {comment: all angular indices at this point are 0 or -1 }
3. let C be the list of candidate chords of T ;
4. **while** C is not empty **do**
5. select and remove a candidate chord (u , w) from C ;
6. flip (u , w);
7. update T , the angular indices, and C
8. **end while**;
9. output T .

6. Remarks and Open Problems

The only place where the heuristic is not deterministically specified is at line 5. What should be the criterion in selecting a candidate chord? There certainly is a sequence of $O(n)$ flips that will solve the problem. (This is the diameter of the flip-graph.) However, even the termination of the heuristic is not quite clear and depends on further refinement of line 5. The resolution of this point needs further research work.

A second question is whether spin numbers, which we studied in Section 4, can have effective algorithmic application.

Another point which deserves further work is exploring additional applications of Theorem 6. One is that, it might be possible to generate the *visibility graph* of P efficiently and with a low computational overhead. For existing work on this problem see [8, 11, 22]. It might be possible to start from one triangulation of P and judiciously go through a sequence of flip operations and collect diagonals of P as they are generated for the first time by the flip operations. The efficiency of the method would depend on the amortized number of flip operations needed to produce one more "new" diagonal. However, the reader should keep in mind the rather negative result stated in Theorem 8.

A similar idea might be possible for computing the single source shortest path trees inside P . For the related work in the literature see [9, 10, 15]. There is a unique shortest path tree for each vertex of P considered as the source. For a vertex v , the edges of the shortest path tree with source v is a subset of at least one triangulation T of P . Let us call a triangulation T with this property a *shortest path triangulation* of P with respect to source v . Once such a triangulation is known, the corresponding shortest path tree can be extracted from it in linear time. Theorem 6 implies the existence of a sequence of flip operations that generates a corresponding sequence of triangulations of P which contains, as a subsequence, a sequence of shortest path triangulations, one for each vertex as the source. If such a sequence of flip operations is relatively short and can be found effectively, it will result in an efficient algorithm to generate all the shortest path trees of the polygon.

References

1. Alexandroff, P.S., *Combinatorial Topology, Vol. I* (1956).
2. Chazelle, B. and J. Incerpi, "Triangulation and shape complexity," *ACM Trans. on Graphics*, pp. 135-152 (1984).
3. Clarkson, K.L., "Applications of random sampling in computational geometry, II" in *Proc. 4th ACM Symp. on Computational Geometry*, pp. 1-11 (1988).
4. Clarkson, K.L., R.E. Tarjan, and C.J. Van Wyk, "A fast Las Vegas algorithm for triangulating a simple polygon" in *Proc. 4th Annual Symp. on Computational Geometry*, pp. 18-22 (1988).

5. Culik, K. and D. Wood, "A note on some tree similarity measures," *Inform. Process. Lett.*, 15, 1, pp. 39-42 (1982).
6. Fournier, A. and D.Y. Montuno, "Triangulating simple polygons and equivalent problems," *ACM Trans. on Graphics*, pp. 153-174 (1984).
7. Garey, M.R., D.S. Johnson, F.D. Preparata, and R.E. Tarjan, "Triangulating a simple polygon," *Inform. Process. Lett.*, pp. 175-180 (1978).
8. Ghosh, S.K. and D.M. Mount, "An output sensitive algorithm for computing visibility graphs" in *Proc. 28th Annual Symp. on Foundations of Computer Science*, pp. 11-19 (1987).
9. Guibas, L.J. and J. Hershberger, "Optimal Shortest path queries in a simple polygon" in *Proc. 3rd Annual Symp. on Computational Geometry*, pp. 50-63 (1987).
10. Guibas, L.J., J. Hershberger, D. Leven, M. Sharir, and R.E. Tarjan, "Linear time algorithms for visibility and shortest path problems inside triangulated simple polygons," *Algorithmica*, 2, 2, pp. 209-234 (1987).
11. Hershberger, J., "Finding the visibility graph of a simple polygon in time proportional to its size" in *Proc. 3rd Annual Symp. on Computational Geometry*, pp. 11-20, Waterloo (1987).
12. Hertel, S. and K. Mehlhorn, "Fast triangulation of a simple polygon" in *Proc. Conf. Found. of Comput. Theory*, pp. 207-218, Springer-Verlag, Berlin, New York (1983).
13. Highnam, P.T., "The ears of a polygon," *Inform. Process. Lett.*, 15, 5, pp. 196-198 (1982).
14. Hoffman, K., K. Mehlhorn, P. Rosenstiehl, and R.E. Tarjan, "Sorting Jordan sequences in linear time using level-linked search trees," *Inform. and Control*, 68, pp. 170-184 (1986).
15. Kapoor, S. and S.N. Maheshwari, "Efficient algorithms for Euclidean shortest path and visibility problems with polygonal obstacles" in *Proc. 4th Annual Symp. on Computational Geometry*, pp. 172-182 (1988).
16. Kirkpatrick, D.G., M.M. Klawe, and R.E. Tarjan, "Polygon triangulation in $O(n \log \log n)$ time with simple data-structures" in *Proc. 6th Annual Symp. on Computational Geometry*, pp. 34-43 (1990).
17. Lopshits, A.M., *Computation of Areas of Oriented Figures* (1963).
18. Lucas, J.M., "The rotation graph of binary trees is Hamiltonian," *J. Algorithms*, 8, 4, pp. 503-535 (1987).
19. Meisters, G.H., "Polygons have ears," *Amer. Math. Monthly*, 82, pp. 648-651 (1975).
20. Mirzaian, A., "Triangulating simple polygons: pseudo triangulations," Technical Report CS-88-12, Dept. of Computer Science, York University (1988).
21. Moise, E.E., *Geometric Topology in Dimensions 2 and 3* (1977).
22. Overmars, M.H. and E. Welzl, "New methods for computing visibility graphs" in *Proc. 4th Annual Symp. on Computational Geometry*, pp. 164-171 (1988).
23. Sleator, D.D., R.E. Tarjan, and W.P. Thurston, "Rotation distance, triangulations, and hyperbolic geometry" in *Proc. 18th Annual ACM Symp. on Theory of Comput.*, pp. 122-135 (1986).
24. Tarjan, R.E. and C.J. Van Wyk, "An $O(n \log \log n)$ -time algorithm for triangulating simple polygons," *SIAM J. Comp.*, 17, 1, pp. 143-178 (1988).
25. Woo, T.C. and S.Y. Shin, "A linear time algorithm for triangulating a point-visible polygon," *ACM Trans. on Graphics*, 4, pp. 60-69 (1985).