Journal of Universal Computer Science, vol. 2, no. 8 (1996), 570-579
submitted: 8/3/96, accepted: 11/7/96, appeared: 28/8/96 © Springer Pub. Co.

The Edge-flipping Distance of Triangulations

Sabine Hanke
(Institut für Informatik, Universität Freiburg, Germany
hanke@informatik.uni-freiburg.de)

Thomas Ottmann
(Institut für Informatik, Universität Freiburg, Germany
ottmann@informatik.uni-freiburg.de)

Sven Schuierer
(Institut für Informatik, Universität Freiburg, Germany
schuiere@informatik.uni-freiburg.de)

Abstract: An edge-flipping operation in a triangulation T of a set of points in the plane is a local restructuring that changes T into a triangulation that differs from T in exactly one edge. The edge-flipping distance between two triangulations of the same set of points is the minimum number of edge-flipping operations needed to convert one into the other. In the context of computing the rotation distance of binary trees Sleator, Tarjan, and Thurston show an upper bound of $2n - 10$ on the maximum edge-flipping distance between triangulations of convex polygons with n nodes, $n > 12$. Using volumetric arguments in hyperbolic 3-space they prove that the bound is tight. In this paper we establish an upper bound on the edge-flipping distance between triangulations of a general finite set of points in the plane by showing that no more edge-flipping operations than the number of intersections between the edges of two triangulations are needed to transform these triangulations into another, and we present an algorithm that computes such a sequence of edge-flipping operations.

Key Words: triangulation, edge-flipping operation, flip, edge-flipping distance, rotation distance

Category: F.2.2, G.2.2

1 Introduction

Triangulations of point sets play an important role in many applications. It is often desirable to compare two triangulations of the same point set. One way to measure the similarity between two triangulations is to compute the *edge-flipping distance* between them. If S is a set of points in the plane and T a triangulation of S, then an *edge-flipping operation* f in T replaces an inner edge e of T with the other diagonal of the quadrilateral Q which surrounds e if Q is convex (Fig. 1). So f transforms T into a triangulation of S that differs from T in exactly one edge. If another edge-flipping operation is used that is not the invers to f, then a triangulation of S is generated that differs from T in exactly two edges, and so on. In this way a triangulation can be changed gradually by a sequence of edge-flipping operations. In the literature this method is used to construct particular triangulations from any starting triangulation, where certain criteria (like the min-max angle criterion to construct the Delaunay triangulation [Lawson 77]) decide which edges are flipped.

The *edge-flipping distance* is now defined as the least number of admissible edge-flipping operations to transform one triangulation into another. Of course,

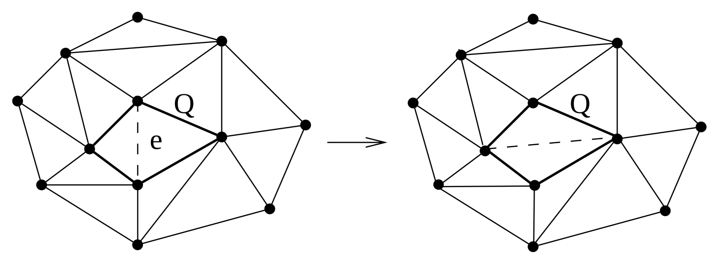

Figure 1: An edge-flipping operation replacing e.

it must be shown that such a transformation is always possible. A different interesting result on comparing two triangulations was recently presented by Aichholzer *et al.* who show that there exists a perfect matching between the edges of the triangulations that intersect [Aichholzer et al. 96].

Already in 1936 Wagner discussed a problem very similar to computing the edge-flipping distance in the context of arbitrary triangulated planar graphs [Wagner 36]: Wagner defines a diagonal transformation in any quadrilateral of a planar graph, and he shows that it is possible to transform any triangulated planar graphs with the same number of nodes into each other by a sequence of those diagonal transformations. In 1973 Dewdney extended Wagner's result to torus graphs [Dewdney 73].

In 1987 Pallo established a $O(n^2)$ algorithm for computing efficient lower and upper bounds of the rotation distance between binary trees [Pallo 87]. Because there is a 1-1-relationship between edge-flipping operations in triangulations of convex polygons with $n+2$ vertices and rotations in binary trees of size n [Sleator et al. 88], Pallo's results are also interesting for the study of edge-flipping distances. Every binary tree with n internal nodes can be represented as a triangulation of a convex polygon $P = \{p_1, ..., p_{n+2}\}$ in the following way: The edge $p_{n+2}p_1$ represents the root, and every other boundary edge p_ip_{i+1} $(i = 1, ..., n+1)$ represents a leaf of the tree. The triangulation contains a triangle with edges e_1, e_2, and e_3 if and only if the node represented as e_1 (without loss of generality) is the father of the nodes represented as e_2 and e_3. So a rotation in the tree corresponds to exactly one edge-flipping operation in the triangulation. Let p be a node in the tree with father q, and let e and e' respectively be the corresponding edges in the triangulation. If we rotate at p, we get the triangulation representing the generated tree by flipping e. Then the new edge represents q, and e' represents p (Fig. 2). Therefore, rotation distances of binary trees and edge-flipping distances of triangulations of convex polygons are equivalent.

In the context of computing the rotation distance of binary trees Sleator, Tarjan, and Thurston showed in 1988 that a transformation of triangulations of convex polygons into each other by using edge-flipping operations is always possible and they prove a tight bound of $2n-10$ on the admissible edge-flipping operations, where $n > 12$ is the number of points of the polygon [Sleator et al. 88]. Furthermore they showed that if it is possible to flip one edge in a triangulation T_1 creating T_1' so that an edge of a triangulation T_2 is generated then there exists

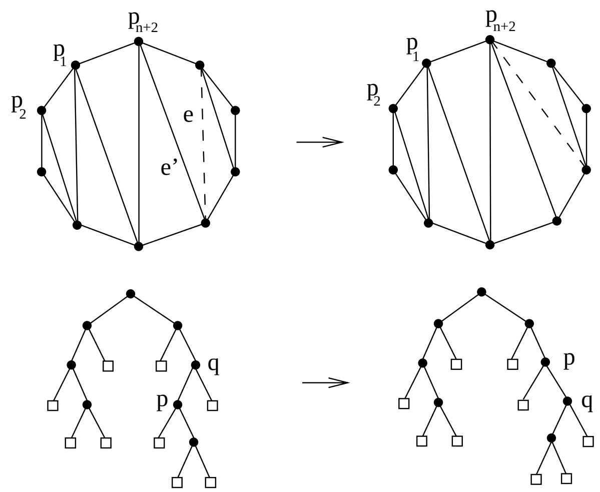

Figure 2: An edge-flipping operation and the corresponding rotation

a sequence of edge-flipping operations of minimal length that transforms T_1 into T_2 in which the first edge-flipping operation creates T_1'.

In 1994 Cai and Hirsch [Cai, Hirsch 94] extended the results of Sleator, Tarjan, and Thurston to rotation distance problems of triangulations of planar surfaces. They give an upper and lower bound for this problem, and analogous to Sleator *et al.* they show that in the case of triangulations of the annulus the flip distance decreases by one through a flip operation which creates a common edge.

In a recent paper Hurtado, Noy, and Urrutia [Hurtado et al. 96] study the problem of flipping edges in triangulations of polygons and point sets. They prove that if a polygon has k reflex vertices, then any two triangulations of this polygon can be transformed into another by flipping at most $O(n + k^2)$ edges. They give examples of polygons with triangulations T and T' such that to transform T into T' requires $O(n^2)$ edge-flipping operations, and they extend these results to triangulations of point sets. Furthermore they show that any triangulation of n points in the plane contains at least $(n - 4)/2$ edges that can be flipped.

Let us now examine our initial problem, whether it is always possible to transform two triangulations T_1 and T_2 of the same set of points in the plane by a sequence of edge-flipping operations. Every triangulation (in particular T_1 and T_2) can be transformed into a Delaunay triangulation with $\mathrm{O}(n^2)$ edge-flipping operations [Bern, Eppstein 92], where n is the number of points. The resulting Delaunay triangulations may be different, if more than three points lie on a circle, but then these points form a convex polygon, which triangulations can be transformed into each other with at most $2n - 6$ edge-flipping operations [Sleator et al. 88]. Since edge-flipping is reversible, it is possible to construct T_2 from T_1 with at most $\mathrm{O}(n^2)$ edge-flipping operations.

In Section 2 we improve this rough estimate of the edge-flipping distance by showing that the number of intersections between the edges of two triangulations is an upper bound on the edge-flipping distance between these triangulations. (We say that two different edges intersect, iff they intersect in their interiors.) In Section 3 we also present an algorithm that computes a sequence of edge-flipping operations that is no longer than the number of intersections.

2 An upper bound on the edge-flipping-distance

In the following we show that the number of intersections between the edges of two triangulations is an upper bound on the edge-flipping distance between these triangulations.

Let T_1 and T_2 be two triangulations of the same set of n points in the plane. We denote by $\#(T_1, T_2)$ the number of intersections of T_1 with T_2 and by $flipdist(T_1, T_2)$ the edge-flipping distance between T_1 and T_2.

Theorem 1: If T_1 and T_2 are two triangulations of the same set of n points in the plane, then

$$flipdist(T_1, T_2) \leq \#(T_1, T_2) < (3n - 2n_b - 3)^2,$$

where n_b is the number of boundary points of both T_1 and T_2.

The basic idea to prove the theorem is to show that for any two triangulations T_1 and T_2 of the same set of points there always exists an edge-flipping operation in T_1 that decreases the number of intersections between these two triangulations. Then, using such kinds of flips we can easily transform T_1 into T_2 with at most $\#(T_1, T_2)$ edge-flipping operations since two triangulations are the same iff the number of intersections is zero.

In order to find such an edge-flipping operation we consider the edges of T_1 that have a maximal number of intersections with the edges in T_2. In the following $\#(e, T)$ denotes the number of intersections between the edge e and the triangulation T, $\#(e_1, e_2, T)$ the number of edges in T that intersect the edge e_1 as well as edge e_2, and $\#_p(e, T)$ the number of edges in T adjacent to p that intersect the edge e.

Lemma 1: *T_1 contains a convex quadrilateral abcd with diagonal ac so that ac has the maximum number of intersections with T_2, i.e. $\#(ac, T_2) = max\{\#(e, T_2) \mid e$ is an edge of $T_1\}$.*

Proof: Let $Q = abcd$ be a quadrilateral in T_1 with diagonal ac so that ac has a maximum number of intersections with T_2. Assume that Q is not convex and that the angle at point a inside the quadrilateral is larger than π. We claim that all edges that intersect the edge ac also intersect both edges bc and cd.

The proof is by contradiction. Clearly, all edges that intersect the edge ac also intersect at least one of the edges bc and cd since the angle at a is larger than π (Fig. 3), and we obtain

$$\begin{aligned}\#(ac, T_2) = {} & \#(ab, cd, T_2) + \#_b(cd, T_2) + \#(bc, cd, T_2) + \\ & \#_d(bc, T_2) + \#(da, bc, T_2)\end{aligned}$$

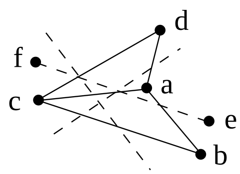

Figure 3: All edges that intersect ac also intersect bc or cd or both

Assume w.l.o.g. there is one edge ef in T_2 that intersects ac and cd but not bc. This implies that either ef intersects ab or one of the end points of ef equals b, and also that ef separates da from bc. Hence, there is no edge in T_2 that intersects both da and bc (or d and bc), i.e. $\#_d(bc, T_2) = \#(da, bc, T_2) = 0$. Therefore,

$$\#(ac, T_2) = \#(ab, cd, T_2) + \#_b(cd, T_2) + \#(bc, cd, T_2),$$

i.e. all the edges that intersect ac also intersect cd.

Because ef is an edge of T_2, and because T_2 contains the points a, b, c, and d as well, there exists an edge adjacent to a in T_2 that intersects the edge cd. This edge does not intersect ac and, therefore, $\#(ac, T_2) \leq \#(cd, T_2) - 1$ which contradicts the maximality of $\#(ac, T_2)$.

Hence, all edges which intersect ac also intersect the edges bc and cd and

$$\#(bc, T_2) = \#(cd, T_2) = \#(ac, T_2) = max\{\#(e, T_2) \mid e \text{ is an edge of } T_1\}$$

In particular, bc and cd cannot be boundary edges of the triangulation. Now consider the quadrilateral Q' in T_1 with diagonal bc. Q' is either convex or there is again a neighbouring quadrilateral Q'' which has a diagonal D'' with the maximal number of intersections and, in addition, D'' is intersected by the same edges as ac and bc. Continuing this process we finally reach the convex hull of the point set. As we observed before a quadrilateral that contains a boundary edge of the triangulation and a diagonal with the maximal number of intersections is convex. □

Lemma 2: *Let abcd be a convex quadrilateral in T_1 with diagonal ac so that ac has the maximum number of intersections with T_2. If T_2 contains an edge eb that intersects da or cd (or an edge dg that intersects ab or bc respectively), then the edge-flipping operation $ac \longrightarrow bd$ decreases the number of intersections of T_1 with T_2.*

Proof: Without loss of generality let eb be an edge of T_2 that intersects da. This implies that all edges of T_2 that intersect bd intersect the edge da as well, because otherwise at least one point of the set S lies inside the triangle bds, where s is the intersection point of eb with da (Fig. 4), and thus lies inside the quadrilateral $abcd$. Since eb does not intersect bd, and because of the assumption that ac has the maximum number of intersections with T_2, it follows that

$$\#(bd, T_2) < \#(da, T_2) \leq \#(ac, T_2)$$

So the edge-flipping operation $ac \longrightarrow bd$ decreases the number of intersections of T_1 with T_2. □

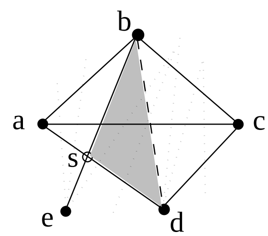

Figure 4: Definition of bds in Lemma 2

Lemma 3: *Let $Q = abcd$ be a convex quadrilateral in T_1 with diagonal ac so that ac has the maximum number of intersections with T_2. If there is no edge eb in T_2 that intersects da or cd and there is no edge dg in T_2 that intersects ab or bc, then either*

1. *the edge-flipping operation $ac \longrightarrow bd$ reduces the number of intersections between the triangulations T_1 and T_2, or*
2. *there is a different quadrilateral Q' in T_1 such that the diagonal of Q' has the maximum number of intersections with T_2 and Q' fulfills the conditions of Lemma 2.*

Proof: If the edge-flipping operation $ac \longrightarrow bd$ decreases the number of intersections between the triangulations T_1 and T_2, then we are done. So in the following we assume that $ac \longrightarrow bd$ does not decrease the number of intersections.

Because the triangulation T_2 does not contain edges eb and dg as assumed in Lemma 2 above, there exists a triangle ebf adjacent to b, where ef intersects ab as well as bc, and a triangle dgh adjacent to d, where gh intersects da as well as cd (see Fig. 5). Therefore, ef and gh do not intersect the edge ac. Because of the assumption that $\#(ac,T_2) \geq max\{\#(ab,T_2), \#(bc,T_2), \#(cd,T_2), \#(da,T_2)\}$ T_2 also contains triangles apq and cxy, where pq intersects ab as well as da and analogous xy intersects bc and cd, because otherwise there is a contradiction: Assume that none or only one of these triangles exist, then without loss of generality T_2 contains an edge aq that intersects bc. This implies that all edges of T_2 that intersect the edge ac intersect bc as well (analogous to the proof of Lemma 2). Because aq does not intersect ac, it follows $\#(ac,T_2) \leq \#(bc,T_2) - 1$, which contradicts the maximality of $\#(ac,T_2)$.

Without loss of generality let T_2 contain an edge uv that intersects the edges da and bc (see Fig. 5), then follows $\#(ab, cd, T_2) = 0$, and

$$\#(ac,T_2) = \#(ab,da,T_2) + \#(bc,cd,T_2) + \#(da,bc,T_2) \quad (1)$$

$$\#(bc,T_2) = \#(ab,bc,T_2) + \#(bc,cd,T_2) + \#(da,bc,T_2) \quad (2)$$

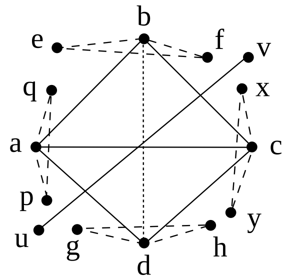

Figure 5: If T_2 does not contain edges eb and dg as in Lemma 2 and $ac \to bd$ does not reduce the number of intersections, then there exist eight points $e, \ldots, y$ such that the edges between Q and $e, \ldots, y$ intersect as displayed.

$$\#(da, T_2) = \#(ab, da, T_2) + \#(da, cd, T_2) + \#(da, bc, T_2) \quad (3)$$
$$\#(bd, T_2) = \#(ab, bc, T_2) + \#(da, cd, T_2) + \#(da, bc, T_2) \quad (4)$$

Because of the maximality of $\#(ac, T_2)$, it follows that

$$(1) \geq (2): \quad \#(ab, da, T_2) \geq \#(ab, bc, T_2) \quad (5)$$
$$(1) \geq (3): \quad \#(bc, cd, T_2) \geq \#(da, cd, T_2) \quad (6)$$

Because the edge-flipping operation $ac \longrightarrow bd$ does not decrease the number of intersections between T_1 and T_2,

$$\#(ab, da, T_2) + \#(bc, cd, T_2) \leq \#(ab, bc, T_2) + \#(da, cd, T_2), \quad (7)$$

and so by (7) and (5) + (6):

$$\#(ab, da, T_2) + \#(bc, cd, T_2) = \#(ab, bc, T_2) + \#(da, cd, T_2). \quad (8)$$

Therefore, by (8) and (5)

$$\begin{aligned} \#(ab, bc, T_2) + \#(da, cd, T_2) - \#(bc, cd, T_2) &= \#(ab, da, T_2) \\ &\geq \#(ab, bc, T_2) \end{aligned} \quad (9)$$

By addition of $\#(bc, cd, T_2)$ in (9), it follows that $\#(bc, cd, T_2) \leq \#(da, cd, T_2)$ and using (6) we obtain $\#(bc, cd, T_2) = \#(da, cd, T_2)$. Analogous we obtain $\#(ab, da, T_2) = \#(ab, bc, T_2)$. So

$$\#(ac, T_2) = \#(bc, T_2) = \#(da, T_2) = \#(bd, T_2).$$

Now consider the triangle bct of the triangulation T_1 that neighbours on $abcd$, then the quadrilateral $abtc$ has a diagonal with a maximal number of intersections, too. If the edge-flipping operation that replaces bc decreases the number of intersections, then we are done, else next take the quadrilateral of T_1 with diagonal bt, and so on. In the end, such a quadrilateral with maximal diagonal contains a boundary edge, and the conditions of Lemma 2 are fulfilled so that an edge-flipping operation replacing this diagonal decreases the number of intersections between T_1 and T_2. □

Proof of Theorem 1: It is clear that $\#(T_1, T_2) < (3n - 2n_b - 3)^2$, because $3n - 2n_b - 3$ is the number of inner edges of both T_1 and T_2.

By Lemma 1, 2, and 3 imply that for all triangulations T_1 and T_2 which are not equal there is an edge-flipping operation in T_1 that decreases the number of intersections between these triangulations. Therefore, if we use such kinds of flips, we can easily transform T_1 into T_2 with at most $\#(T_1, T_2)$ edge-flipping operations, since two triangulations are the same iff the number of intersections is zero. □

3 The Algorithm

In the case of triangulations of convex polygons Sleator, Tarjan, and Thurston [Sleator et al. 88] show that if it is possible to flip one edge in a triangulation T_1 creating T_1' so that an edge of a triangulation T_2 is generated, then there exists a sequence of edge-flipping operations of minimal length that transforms T_1 into T_2 in which the first edge-flipping operation creates T_1'. Therefore, if we want to transform two given triangulations of a convex polygon into each other by using only a minimum number of edge-flipping operations, we start to flip edges such that an edge of the other triangulation is created by each operation. But what to do, if at least two edge-flipping operations are needed to generate a common edge, is still an open question.

The simple algorithm implied by the proof of Theorem 1 gives a heuristic what to do in this case. The strategy to create common edges, whenever possible, seems to be a good heuristic in the case of triangulations of point sets as well. So in the following we present a combined algorithm to transform two given triangulations of the same point set into each other.

The algorithm we present makes use of the following lemma.

Lemma 4: *Let T and T' be two triangulations of the same point set. If e' is an edge of T' which is intersected by exactly one edge e in T, then e' is an edge that can be created by one flip, i.e. by flipping $e \longrightarrow e'$.*

Proof: Assume that T and T' are two triangulations of the same point set and e' is an edge of T' that has only one intersection with the triangulation T. Let e be the edge of T that intersects e'. Then e' is not contained in T. Let $e' = a'c'$, that means a' and c' are the endpoints of the edge e', and let $abcd$ be the quarilateral around $e = bd$ in T. Because the boundary edges ab, bc, cd, and da of the quadrilateral do not intersect the edge e', and because the endpoints of e' cannot lie in the interior of the quadrilateral, it follows that $ac = a'c'$. Since $e = bd$ intersects $e' = a'c' = ac$, $abcd$ is convex, and thus e' is the second diagonal of the quadrilateral around e in the triangulation T. □

Using pseudo code we can now describe the algorithm as follows:

```
flip_sequence(T1, T2)
{
  init stack S;
```

$T_1' := T_1$; $T_2' := T_2$;

while $(T_1' \neq T_2')$ **do**
{
while (there exists an edge e' in T_i' with $\#(e', T_j') = 1$, $i, j \in \{1, 2\}$) **do**
{ Let e be the edge in T_j' that intersects e';
flip $e \longrightarrow e'$ in T_j';
if $(j = 1)$ **then** output($e \longrightarrow e'$);
else S.push($e' \longrightarrow e$);
}

if $(T_1' \neq T_2')$ **then**
{ Let e be an edge of T_i', $i \in \{1, 2\}$, such that flipping e decreases $\#(T_1, T_2)$ by the greatest amount;
Let e' be the second diagonal of the quadrilateral around the edge e in T_i';
flip $e \longrightarrow e'$ in T_i';
if $(i = 1)$ **then** output($e \longrightarrow e'$);
else S.push($e' \longrightarrow e$);
}
}

while (S not empty) **do** output(S.pop);
}

In order to prove the correctness of the algorithm we have to show that the outer while-loop terminates. Observe that:

1. By each execution of the inner while-loop a common edge of the triangulations T_1' and T_2' is generated, and so the number of intersections is decreased by at least one.

2. The if-statement is executed if and only if $T_1' \neq T_2'$ and all edges that T_1' and T_2' do not have in common are intersected by at least two edges of the other triangulation. By Lemma 1, 2, and 3 we have shown that for all triangulations T_1' and T_2' which are not equal there is an edge-flipping operation that decreases the number of intersections between these triangulations.

Therefore, the outer while-loop of the presented algorithm terminates after at least $\#(T_1', T_2')$ steps, because $T_1' = T_2'$ iff $\#(T_1', T_2') = 0$. Then T_1' and T_2' have been transformed to the same triangulation T, where the sequence of edge-flipping operations that transforms T_1' into T has been written on the output during the transformation. Now the sequence that transforms T_2' into T is put out in reverse order by the last while-loop of the algorithm. So the given algorithm computes a sequence of edge-flipping operations with length of at least $\#(T_1', T_2')$ that transforms T_1' into T_2'.

4 Conclusion

We introduce the *edge-flipping distance* between two triangulations T_1 and T_2 of the same set of points in the plane, and we give an algorithm that computes a sequence of edge-flipping operations to transforms T_1 into T_2 which is no longer than the number of intersections between T_1 and T_2. This algorithm uses the heuristic of always flipping so as to create an edge of T_2, whenever possible. So in the case of triangulations of convex polygons the algorithm computes the beginning and end of an optimal sequence of edge-flipping operations until the first occurrence of the if-statement. And, therefore, because of the 1-1-relationship between edge-flipping operations in triangulations of convex polygons with $n+2$ vertices and rotations in binary trees of size n this algorithm is also interesting for the computing of an upper bound on the rotation distance between two given binary trees.

Note that in a sequence of edge-flipping operations with minimal length that transform two triangulations into each other, the number of intersections between the generated triangulation during the transformation and the final triangulation does not necessarily decrease after each step. There are some examples of pairs of triangulations and optimal sequences of edge-flipping operations where the number of intersections even increases by some edge-flipping operations during the transformation.

Of course, one of the main open questions is the complexity status of computing the edge-flipping distance between two triangulations. It is not known whether the problem is in NP even if we restrict ourselves to convex polygons. In particular, the structure of a minimal length sequence of edge-flipping operations is unknown.

References

[Aichholzer et al. 96] Aichholzer, O., Aurenhammer, F., Cheng, S.-W., Rote, G., Taschwer, M., and Xu, Y.-F.: "Triangulations Intersect Nicely"; Distcr. & Comput. Geom. (1996), to appear. Also Proc. 11th ACM Symp. Computational Geometry (1995), 220-229.

[Bern, Eppstein 92] Bern, M., Eppstein, D.: "Mesh generation and optimal triangulation"; Computing in Euclidean Geometry, Lecture Notes Series on Computing, 1, World Scientific, Singapore (1992), 23-90.

[Cai, Hirsch 94] Cai, J.-Y., Hirsch, M.: "Rotation Distance, Triangulations of Planar Surfaces and Hyperbolic Geometry"; Proc. ISAAC'94 (5th International Symposium on Symbolic and Algebraic Computation), Beijing (1994), 172-180.

[Dewdney 73] Dewdney, A.K.: "Wagner's theorem for torus graphs"; Discrete Mathematics, 4 (1973), 139-149.

[Hurtado et al. 96] Hurtado, F., Noy, M., and Urrutia, J.: "Flipping Edges in Triangulations"; Proc. 12th Annual ACM Symposium on Computational Geometry (1996), 214-223.

[Pallo 87] Pallo, J.: "On the rotation distance in the lattice of binary trees"; Information Processing Letters, 25 (1987), 369-373.

[Lawson 77] Lawson, C.L.: "Software for C^1 surface interpolation"; Mathematical Software III, Academic Press, New York (1977), 161-194.

[Sleator et al. 88] Sleator, D.D., Tarjan, R.E., Thurston, W.R.: "Rotation distance, triangulations and hyperbolic geometry"; Journal of the American Mathematical Society, 1, 3 (1988), 647-681.

[Wagner 36] Wagner, K.: "Bemerkungen zum Vierfarbenproblem"; Jahresbericht der Deutschen Mathematiker-Vereinigung, 46 (1936), 26-32.