

Flip Distance Between Triangulations of a Simple Polygon is NP-Complete

Oswin Aichholzer*

Wolfgang Mulzer†

Alexander Pilz‡

Abstract

Let T be a triangulation of a simple polygon. A *flip* in T is the operation of removing one diagonal of T and adding a different one such that the resulting graph is again a triangulation. The flip distance between two triangulations is the smallest number of flips that is necessary to transform one triangulation into the other. We show that computing the flip distance between two triangulations of a simple polygon is NP-hard.

1 Introduction

Let P be a simple polygon in the plane, that is, the closed region bounded by a piece-wise linear, simple cycle. A *triangulation* T of P is a geometric (straight-line) maximal outerplanar graph whose outer face is the complement of P and whose vertex set are the vertices of P . Let d be a diagonal of P whose removal creates a convex quadrilateral f . By replacing d with the other diagonal of f , we again get a triangulation of P . This operation is called a *flip*. The *flip graph* of P is the abstract graph whose vertices are the triangulations of P and in which two triangulations are adjacent if and only if they differ by a single flip. We study the *flip distance*, that is, the minimum number of flips required to transform a given source triangulation into a target triangulation.

Edge flips became popular in the context of Delaunay triangulations. Lawson [6] proved that the flip graph is connected with diameter $O(n^2)$. Hurtado, Noy, and Urrutia [5] gave an example where the flip distance is $\Omega(n^2)$, and showed that the same bounds hold for triangulations of simple polygons. They also proved that if the polygon has k reflex vertices, then the flip graph has diameter $O(n + k^2)$. This generalizes the well-known fact that the flip distance between

any two triangulations of a convex polygon is at most $2n - 10$, for $n > 12$ [11].

Hanke, Ottmann, and Schuierer [4] showed that the flip distance between two triangulations of a point set is at most the number of crossings in the overlay of the source and the target triangulation. Eppstein [3] gave a polynomial-time algorithm for computing a lower bound on the flip distance. This bound is tight for point sets that do not contain empty 5-gons. For a survey on flip operations see Bose and Hurtado [2]. Recently, the problem of finding the flip distance between two triangulations of a point set was shown to be NP-complete by Lubiw and Pathak [7] and, independently, to be APX-hard by Pilz [8]. Here, we obtain the following result.

Theorem 1 (main result) *Let P be a simple polygon. It is NP-complete to decide whether the flip distance between two triangulations of P is at most k .*

In this extended abstract we omit all proofs; they can be found in the preprint [1]. Our reduction uses a variant of the RECTILINEAR STEINER ARBORESCENCE PROBLEM. Let S be a set of N points in the plane, called *sinks*, whose coordinates are nonnegative integers. A rectilinear Steiner tree T is called a *rectilinear Steiner arborescence* (RSA) for S if (i) T is rooted at the origin; (ii) each leaf of T lies at a sink in S ; and (iii) for each $s = (x, y) \in S$, the length of the path in T from the origin to s equals $x + y$, i.e., all edges in T point north or east, as seen from the origin [9]. In the RSA problem, we are given a set of sinks S and an integer k and ask whether S has an RSA of length at most k . The problem is strongly NP-complete; in particular, it remains NP-complete if S is contained in an $n \times n$ grid, with n polynomially bounded in $N = |S|$ [10]. For our reduction we need a restricted version of the RSA problem, called the YRSA problem. In an instance (S, k) of the YRSA problem, we require that no two sinks in S have the same y -coordinate. One can show that this variant of the problem remains NP-complete.

2 Double Chains

We use definitions (and illustrations) along the lines of [8]. A *double chain* D consists of two chains, an *upper chain* and a *lower chain*. There are n points

*Institute for Software Technology, Graz University of Technology, Austria. Partially supported by the ESF EUROCORES programme EuroGIGA - ComPoSe, Austrian Science Fund (FWF): I 648-N18. oaich@ist.tugraz.at.

†Institute of Computer Science, Freie Universität Berlin, Germany. Supported in part by DFG project MU/3501/1. mulzer@inf.fu-berlin.de.

‡Recipient of a DOC-fellowship of the Austrian Academy of Sciences at the Institute for Software Technology, Graz University of Technology. Part of this work has been done while this author was visiting the Dept. de Matemáticas, Universidad de Alcalá, Spain. apilz@ist.tugraz.at.




Figure 1: The upper and lower extreme triangulations of P_D with a flip distance of $(n - 1)^2$, as shown in [5].




Figure 2: The extra point p in the kernel of D allows flipping one extreme triangulation of P_D^p to the other in $4n - 4$ flips.

on each chain, $\langle u_1, \dots, u_n \rangle$ on the upper chain and $\langle l_1, \dots, l_n \rangle$ on the lower chain, both numbered from left to right. The upper chain is reflex w.r.t. any point of the lower chain, and vice versa. Let P_D be the polygon defined by $\langle l_1, \dots, l_n, u_n, \dots, u_1 \rangle$. We call the triangulation T_u of P_D where u_1 has maximum degree the *upper extreme triangulation*; observe that this triangulation is unique. The triangulation T_l of P_D where l_1 has maximum degree is called the *lower extreme triangulation*. The flip distance between T_u and T_l is $(n - 1)^2$ [5], see Figure 1.

Definition 1 Let D be a double chain. Let W_1 be the double wedge defined by the supporting lines of u_1u_2 and l_1l_2 whose interior does not contain a point of D . W_n is defined analogously by the supporting lines of u_nu_{n-1} and l_nl_{n-1} . Let $W = W_1 \cup W_n$ be called the wedge of D . A point is outside of D if it is not contained in $W \cup P_D$. The kernel of D is the intersection of the closed half-planes below u_1u_2 and $u_{n-1}u_n$, as well as above l_1l_2 and $l_{n-1}l_n$.

We refer to a polygon as in Figure 2, where p is in the kernel of D , by P_D^p . As mentioned in [12], the flip distance between the two extreme triangulations from Figure 1 is much smaller in P_D^p than in P_D . Figure 2 shows that $4n - 4$ flips suffice. It turns out that this is optimal, even for more general polygons:

Lemma 2 Let P be a polygon that completely contains P_D and has $\langle l_1, \dots, l_n \rangle$ and $\langle u_n, \dots, u_1 \rangle$ as part of its boundary. Further, let T_1 and T_2 be two triangulations that contain the upper extreme triangulation and the lower extreme triangulation of P_D as a subtriangulation, respectively. Then T_1 and T_2 have flip distance at least $4n - 4$.

The proof by Lubiw and Pathak [7] for constant-size double chains directly generalizes to the above result. The following is a special case of a result from [8].




Figure 3: A double chain extended by a vertex z . The vertex z is incident to u_7 and l_8 , represented by the blue point b in the grid. The brown chain path represents the chain triangles. If we flip edges to z , b will move along that path. A flip between chain triangles (dotted edge replaced by the dashed one) changes a bend in that path (from the dotted one).

Lemma 3 Let P be a polygon that completely contains P_D and has $\langle l_1, \dots, l_n \rangle$ and $\langle u_n, \dots, u_1 \rangle$ as part of its boundary, and let T_1 and T_2 be two triangulations that contain the upper extreme triangulation and the lower extreme triangulation of P_D as subtriangulation, respectively. Suppose there is no vertex in the interior of the wedge of P_D . Then the flip distance between T_1 and T_2 is at least $(n - 1)^2$.

Take a polygon P_D^z and consider a triangulation T of P_D^z . A *chain edge* is an edge of T between the upper and the lower chain of D . A *chain triangle* is a triangle that contains two chain edges. We use the chain edges to define the *chain path*, an abstract path on the $n \times n$ grid. Let e_1, e_2, \dots, e_m be the chain edges, sorted from left to right according to their intersections with a line ℓ that separates the upper from the lower chain. For $i = 1, \dots, m$, write $e_i = (u_v, l_w)$ and set $c_i = (v, w)$. Note that, in particular, $c_1 = (1, 1)$, which we use as the *root* of our setting. Since T is a triangulation, any two consecutive edges e_i, e_{i+1} share one endpoint, while the other endpoints are adjacent on the corresponding chain. Thus, c_{i+1} dominates c_i and $\|c_{i+1} - c_i\|_1 = 1$. The chain path is defined as the path $c_1c_2 \dots c_m$. See Figure 3 for an example.

The chain path is an x - and y -monotone path in the $n \times n$ grid. We call its upper right endpoint b . By observing the changes of the chain path by flips of different types, the following lemma can be obtained.

Lemma 4 Let T be a triangulation of P_D^z . Then T uniquely determines an x - and y -monotone path (i.e., the chain path) in the $n \times n$ grid starting at the root $(1, 1)$. Conversely, any chain path uniquely determines a triangulation of T . The possible flips of T correspond to the following operations on the chain path: (i) extend the right endpoint north or east; (ii) shorten the path at the right endpoint; (iii) change an east-north bend to a north-east bend, or vice versa.

3 Installing Sinks

We show how to reduce YRSA to our flip distance problem. Let S be a set of N sinks with root at $(1, 1)$ on an $(n-1) \times (n-1)$ grid (recall that n is polynomial in N). We describe how to construct a polygon P_D^* for S . Our construction has two integral parameters β and d . With foresight, we set $\beta = 2N$ and $d = nN$.

Let P_D^z be the polygon from Section 2, but with βn vertices on each chain. As we saw in Section 2, we can interpret a triangulation of P_D^z as a chain path in the $\beta n \times \beta n$ grid. We imagine that the sinks of S are in this grid, with their coordinates multiplied by β . For each sink $s = (x, y)$, we place a (rotated) small double chain D_s of size d such that $l_{\beta y}$ and $l_{\beta y+1}$ correspond to the last point on the lower and upper chain of D_s , respectively. In addition, $u_{\beta x}$ is the only point in the kernel of D_s and $u_{\beta x}$ is also the only point in the interior of the wedge of D_s . We call the resulting polygon P_D^* . If β is large enough, the small double chains D_s do not interfere with each other, and P_D^* is simple. Since the y -coordinates in S are pairwise distinct, we create at most one double chain at each edge of the lower chain of P_D^z . Observe that we have some flexibility for the precise placement of the points of each D_s . Thus we can choose their placement in a way that their coordinates are polynomial in n .

Next, we describe the source and target triangulation for P_D^* . The source triangulation T_1 contains all edges of P_D^z . The interior of P_D^z is triangulated such that all edges are incident to z , i.e., b is at the root. The small double chains are all triangulated with the upper extreme triangulation. The target triangulation T_2 is defined similarly, but now all the small double chains are triangulated with the lower extreme triangulation (note that the choice of the upper and lower chain is arbitrary for the small double chains).

Hence, each corresponding pair of small double chains in T_1 and T_2 has flip distance $(d-1)^2$ due to Lemma 3, unless the appropriate vertex on the upper chain of P_D^* is used. Intuitively, if d is large enough, a shortest flip sequence will have to “traverse” each sink, inducing an arborescence for S . Vice versa, every arborescence for S gives a short flip sequence between T_1 and T_2 .

Lemma 5 *Let A be an arborescence for S of length k . Then the flip distance on P_D^* between T_1 and T_2 is at most $2\beta k + (4d-2)N$.*

Next we consider the opposite direction of the correspondence. In the proof of the following lemma, we will describe a mapping from each triangulation T of P_D^* to a triangulation T_z of P_D^z . For each sink $s \in S$, the corresponding chain triangle t_s in T_z is defined as the chain triangle in P_D^z that allows the double chain D_s to be flipped quickly. We say that a flip sequence σ_1 on P_D^z visits a sink $s \in S$, if σ_1 has at least one




Figure 4: A part of a triangulation of P_D^* and the two corresponding triangulations T_z and T_s .

triangulation T that contains the corresponding chain triangle t_s . We call σ_1 a *flip traversal* for S if (i) the sequence σ_1 begins and ends in the same triangulation T_z such that T_z corresponds to b lying on the root; (ii) the sequence σ_1 visits every sink in S .

Lemma 6 *Let σ be a flip sequence on P_D^* from T_1 to T_2 with $|\sigma| < (d-1)^2$. Then there exists a flip sequence σ_1 on P_D^z such that σ_1 is a flip traversal for S with $|\sigma_1| \leq |\sigma| - (4d-4)N$.*

Sketch of Proof. Let T^* be a triangulation of P_D^* . Let D_s be a small double chain placed between the vertices l_s and l'_s with u_s being the vertex in the kernel of D_s . We define Δ_s as the triangle that is either the inner triangle (i.e., all three sides are diagonals) incident to two vertices of D_s or the triangle that is incident to both convex vertices of D_s but is not an ear. Note that in the first case the third vertex might be u_s and that in the latter case the third vertex has to be u_s . Due to the structure of P_D^* there always exists exactly one such triangle Δ_s per sink. Let the polygon $P_D^{u_s}$ consist of the double chain D_s extended by the vertex u_s , and let T_s denote a triangulation of it. We define a mapping of any triangulation T^* of P_D^* to a triangulation T_z of P_D^z and to triangulations T_s for all sinks s . The triangulation T_z contains every triangle that has all three vertices in P_D^z . For each triangle ∇ that has two vertices on P_D^z and one on the left chain of D_s , we replace the apex on D_s by l_s . The analogous is done if the apex of a triangle ∇ is on the right chain of D_s ; we replace that apex by l'_s . For every sink s , the triangle Δ_s is known to have an apex at a point u_i of the upper chain. In T_z , we replace Δ_s by the triangle $l_s l'_s u_i$. Since these are exactly the triangles needed for a triangulation of P_D^z and no two triangles overlap, T_z is indeed a triangulation of P_D^z . Similarly, all triangles in T^* with all three vertices on $P_D^{u_s}$ are also in T_s , and the triangles having two points on D_s and whose apex is not in $P_D^{u_s}$ get their apex at u_s in T_s (note that this includes Δ_s). See Figure 4.

Using a case analysis, one can show that each flip changes at most one of the triangulations that the original triangulation is mapped to. \square

4 Arborescences and Traces

Next, we define traces (domains drawn on the grid) and relate them to flip traversals. A *trace* is drawn on the $\beta n \times \beta n$ grid. It consists of *edges* and *boxes*: an edge is a line segment of length 1 whose endpoints have positive integer coordinates; a box is a square of side length 1 whose corners have positive integer coordinates. Similar to arborescences, we require that a trace R (i) is (topologically) connected; (ii) contains the root $(1, 1)$; and (iii) from every grid point contained in R there exists an x - and y -monotone path to the root that lies completely in R . We say R is a *covering trace* for S (or, R *covers* S) if every sink in S is part of R .

Let σ_1 be a flip traversal as in Lemma 6. By Lemma 4, we can interpret the sequence σ_1 as the evolution of a chain path. This gives a covering trace R for S in the following way. For every flip in σ_1 that extends the chain path, we add the corresponding edge to R . For every chain flip in σ_1 , we add the corresponding box to R . Afterwards, we remove from R all edges that coincide with a side of a box in R . Clearly, R is (topologically) connected. Since σ_1 is a flip traversal for S , every sink is covered by R (i.e., incident to a box or edge in R). Note that every grid point p in R is connected to the root by an x - and y -monotone path on R , since at some point p belonged to a chain path in σ_1 . Hence, R is indeed a trace, the unique *trace of* σ_1 .

Next, we define the *cost* of a trace R , $\text{cost}(R)$, so that if R is the trace of a flip traversal σ_1 , then $\text{cost}(R)$ gives a lower bound on $|\sigma_1|$. An edge has cost 2. Let B be a box in R . A *boundary side* of B is a side that is not part of another box. The cost of B is 1 plus the number of boundary sides of B . Then, $\text{cost}(R)$ is the total cost over all boxes and edges in R .

Proposition 7 *Let σ_1 be a flip traversal and R a trace for σ_1 . Then $\text{cost}(R) \leq |\sigma_1|$.*

Observation 1 *Any shortest path tree A_{σ_1} in R for the root w.r.t. S is an arborescence.*

If σ_1 contains no chain flips, the corresponding trace R has no boxes, but it may not be acyclic. However, due to Observation 1 it contains an arborescence A_{σ_1} , in particular with $2|A_{\sigma_1}| \leq \text{cost}(R)$.

Lemma 8 *Let σ_1 be a flip traversal of S . Then there exists a covering trace R for S in the $\beta n \times \beta n$ grid such that R does not contain a box and such that $\text{cost}(R) \leq |\sigma_1|$.*

Corollary 9 *Let σ be a flip sequence on P_D^* from T_1 to T_2 with $|\sigma| \leq 2\beta k + (4d - 2)N$. Then there exists a rectilinear Steiner arborescence for S of length at most k .*

Sketch of Proof. Since there is always an arborescence on S of length less than $2nN$, we may assume that $k < 2nN$. We can use Lemma 6, and then apply Lemma 8 to the resulting sequence to obtain an arborescence A of length at most $\beta k + N$. It is well-known that there exists a minimal arborescence A' for S whose length is a multiple of β . Thus, since $\beta > N$, we get that A' has length at most βk , so the corresponding arborescence for S on the original grid has length at most k . \square

Together with Lemma 5, this implies Theorem 1.

References

- [1] O. Aichholzer, W. Mulzer, and A. Pilz. Flip Distance Between Triangulations of a Simple Polygon is NP-Complete. *ArXiv e-prints*, Sept. 2012. arXiv:1209.0579.
- [2] P. Bose and F. Hurtado. Flips in planar graphs. *Comput. Geom.*, 42(1):60–80, 2009.
- [3] D. Eppstein. Happy endings for flip graphs. *JoCG*, 1(1):3–28, 2010.
- [4] S. Hanke, T. Ottmann, and S. Schuierer. The edge-flipping distance of triangulations. *J.UCS*, 2(8):570–579, 1996.
- [5] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete Comput. Geom.*, 22:333–346, 1999.
- [6] C. L. Lawson. Transforming triangulations. *Discrete Math.*, 3(4):365–372, 1972.
- [7] A. Lubiw and V. Pathak. Flip distance between two triangulations of a point-set is NP-complete. In *Proc. 24th CCCG*, pages 127–132, Charlottetown, Canada, August 2012.
- [8] A. Pilz. Flip distance between triangulations of a planar point set is APX-hard. Submitted, preprint available at arXiv:1206.3179, 2012.
- [9] S. K. Rao, P. Sadayappan, F. K. Hwang, and P. W. Shor. The rectilinear Steiner arborescence problem. *Algorithmica*, 7:277–288, 1992.
- [10] W. Shi and C. Su. The rectilinear Steiner arborescence problem is NP-complete. In *Proc. 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 780–787, 2000.
- [11] D. Sleator, R. Tarjan, and W. Thurston. Rotation distance, triangulations and hyperbolic geometry. *J. Amer. Math. Soc.*, 1:647–682, 1988.
- [12] J. Urrutia. Algunos problemas abiertos. In N. Coll and J. Sellares, editors, *Proc. IX ECG*, pages 13–24. Univ. De Girona, July 2001.