

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

Oswin Aichholzer¹, Wolfgang Mulzer², and Alexander Pilz¹

¹ Institute for Software Technology, Graz University of Technology, Austria

² Institute of Computer Science, Freie Universität Berlin, Germany

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 required to transform one triangulation into the other. For the special case of convex polygons, the problem of determining the shortest flip distance between two triangulations is equivalent to determining the rotation distance between two binary trees, a central problem which is still open after over 25 years of intensive study.

We show that computing the flip distance between two triangulations of a simple polygon is NP-complete. This complements a recent result that shows APX-hardness of determining the flip distance between two triangulations of a planar point set.

1 Introduction

Let P be a simple polygon in the plane, that is, a 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 consists of the vertices of P . The edges of T that are not on the outer face are called *diagonals*. Let d be a diagonal whose removal creates a convex quadrilateral f . Replacing d with the other diagonal of f yields another 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*, i.e., 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 [9] proved that any triangulation of a planar n -point set can be transformed into any other by $O(n^2)$ flips. Hence, for every planar n -point set the flip graph is connected with diameter $O(n^2)$. Later, he showed that in fact every triangulation can be transformed to the Delaunay triangulation by $O(n^2)$ flips that locally

* O.A. partially supported by the ESF EUROCORES programme EuroGIGA - Com-PoSe, Austrian Science Fund (FWF): I 648-N18. W.M. supported in part by DFG project MU/3501/1. A.P. is recipient of a DOC-fellowship of the Austrian Academy of Sciences at the Institute for Software Technology, Graz University of Technology.

fix the Delaunay property [10]. Hurtado, Noy, and Urrutia [7] gave an example where the flip distance is $\Omega(n^2)$, and they 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 $O(n+k^2)$. In particular, the flip graph of any planar polygon has diameter $O(n^2)$. Their result also generalizes the well-known fact that the flip distance between any two triangulations of a convex polygon is at most $2n - 10$, for $n > 12$, as shown by Sleator, Tarjan, and Thurston [15] in their work on the flip distance in convex polygons. The latter case is particularly interesting due to the correspondence between flips in triangulations of convex polygons and rotations in binary trees: The dual graph of such a triangulation is a binary tree, and a flip corresponds to a rotation in that tree; also, for every binary tree, a triangulation can be constructed.

We mention two further remarkable results on flip graphs for point sets. Hanke, Ottmann, and Schuierer [6] showed that the flip distance between two triangulations is bounded by the number of crossings in their overlay. Eppstein [5] gave a polynomial-time algorithm for calculating a lower bound on the flip distance. His bound is tight for point sets with no empty 5-gons; however, except for small instances, such point sets are not in general position (i.e., they must contain collinear triples) [1]. For a recent survey on flips see Bose and Hurtado [3].

Very recently, the problem of finding the flip distance between two triangulations of a point set was shown to be NP-hard by Lubiw and Pathak [11] and, independently, by Pilz [12], and the latter proof was later improved to show APX-hardness of the problem. Here, we show that the corresponding problem remains NP-hard even for simple polygons. This can be seen as a further step towards settling the complexity of deciding the flip distance between triangulations of convex polygons or, equivalently, the rotation distance between binary trees. This variant of the problem was probably first addressed by Culik and Wood [4] in 1982 (showing a flip distance of $2n - 6$).

The formal problem definition is as follows: given a simple polygon P , two triangulations T_1 and T_2 of P , and an integer l , decide whether T_1 can be transformed into T_2 by at most l flips. We call this decision problem POLYFLIP. To show NP-hardness, we give a polynomial-time reduction from RECTILINEAR STEINER ARBORESCENCE to POLYFLIP. RECTILINEAR STEINER ARBORESCENCE was shown to be NP-hard by Shi and Su [14]. In Section 2, we describe the problem in detail. We present the well-known *double chain* (used by Hurtado, Noy, and Urrutia [7] for giving their lower bound), a major building block in our reduction, in Section 3. Finally, in Section 4, we describe our reduction and prove that it is correct. An extended abstract of this work was presented at the 29th EuroCG, 2013; for omitted proofs, see [2].

2 The Rectilinear Steiner Arborescence Problem

Let S be a set of N points in the plane whose coordinates are nonnegative integers. The points in S are called *sinks*. A *rectilinear tree* T is a connected acyclic collection of horizontal and vertical line segments that intersect only

at their endpoints. The *length* of T is the total length of all segments in T (cf. [8, p. 205]). The tree T is a *rectilinear Steiner tree* for S if each sink in S appears as an endpoint of a segment in T . We call T 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 [13]. In the *RSA problem*, we are given a set of sinks S and an integer k . The question is whether there is an RSA for S of length at most k . Shi and Su showed that the RSA 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 , the number of points [14].³

We recall an important structural property of the RSA. Let A be an RSA for a set S of sinks. Let e be a vertical segment in A that does not contain a sink. Suppose there is a horizontal segment f incident to the upper endpoint a of e . Since A is an arborescence, a is the left endpoint of f . Suppose further that a is not the lower endpoint of another vertical edge. Take a copy e' of e and translate it to the right until e' hits a sink or another segment endpoint (this will certainly happen at the right endpoint of f); see Fig. 1. The segments e and e' define a rectangle R . The upper and left side of R are completely covered by e and (a part of) f . Since a has only two incident segments, every sink-root path in A that goes through e or f contains these two sides of R , entering the boundary of R at the upper right corner d and leaving it at the lower left corner b . We reroute every such path at d to continue clockwise along the boundary of R until it meets A again (this certainly happens at b), and we delete e and the part of f on R . In the resulting tree we subsequently remove all unnecessary segments (this happens if there are no more root-sink paths through b) to obtain another RSA A' for S . Observe that A' is not longer than A . This operation is called *sliding e to the right*. If similar conditions apply to a horizontal edge, we can *slide it upwards*. The *Hanan grid* for a point set P is the set of all vertical and horizontal lines through the points in P . In essence, the following theorem can be proved constructively by repeated segment slides in a shortest RSA.

Theorem 2.1 ([13]). *Let S be a set of sinks. There is a minimum-length RSA A for S such that all segments of A are on the Hanan grid for $S \cup \{(0, 0)\}$. \square*




Fig. 1. The slide operation. The dots depict sinks; the rectangle R is drawn gray. The dotted segments are deleted, since they no longer lead to a sink.

³ Note that a polynomial-time algorithm was claimed [16] that later has been shown to be incorrect [13].

We use a restricted version of the RSA problem, called YRSA. An instance (S, k) of the YRSA problem differs from an instance for the RSA problem in that we require that no two sinks in S have the same y -coordinate. The NP-hardness of YRSA follows by a simple perturbation argument; see the full version for all omitted proofs.

Theorem 2.2. *YRSA is strongly NP-complete.*

3 Double Chains

Our definitions (and illustrations) follow [12]. A *double chain* D consists of two chains, an *upper chain* and a *lower chain*. There are n vertices 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. Any point on one chain sees every point on the other chain, and any quadrilateral formed by three vertices of one chain and one vertex of the other chain is non-convex. Let P_D be the polygon defined by $\langle l_1, \dots, l_n, u_n, \dots, u_1 \rangle$; see Fig. 2 (left). 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 two extreme triangulations are used to show that the diameter of the flip graph is quadratic; see Fig. 2 (right).

Theorem 3.1 ([7]). *The flip distance between T_u and T_l is $(n - 1)^2$.* \square

Through a slight modification of D , we can reduce the flip distance between the upper and the lower extreme triangulation to linear. This will enable us in our reduction to impose a certain structure on short flip sequences. To describe this modification, we first define the *flip-kernel* of a double chain.




Fig. 2. Left: The polygon and the hourglass (gray) of a double chain. The diamond-shaped flip-kernel can be extended arbitrarily by flattening the chains. Right: The upper extreme triangulation T_u and the lower extreme triangulation T_l .

Let W_1 be the wedge defined by the lines through u_1u_2 and l_1l_2 whose interior contains no point from D but intersects the segment u_1l_1 . Define W_n analogously by the lines through u_nu_{n-1} and l_nl_{n-1} . We call $W := W_1 \cup W_n$ the *hourglass of D* . The unbounded set $W \cup P_D$ is defined by four rays and the two chains. The *flip-kernel* of D is the intersection of the closed half-planes below the lines through u_1u_2 and $u_{n-1}u_n$, as well as above the lines through l_1l_2 and $l_{n-1}l_n$.⁴

⁴ The flip-kernel of D might not be completely inside the polygon P_D . This is in contrast to the “visibility kernel” of a polygon.




Fig. 3. The extra point p in the flip-kernel of D allows flipping one extreme triangulation of P_D^p to the other in $4n - 4$ flips.

Definition 3.2. Let D be a double chain whose flip-kernel contains a point p to the right of the directed line $l_n u_n$. The polygon P_D^p is given by the sequence $\langle l_1, \dots, l_n, p, u_n, \dots, u_1 \rangle$. The upper and the lower extreme triangulation of P_D^p contain the edge $u_n l_n$ and otherwise are defined in the same way as for P_D .

The flip distance between the two extreme triangulations for P_D^p is much smaller than for P_D [17]. Fig. 3 shows how to transform them into each other with $4n - 4$ flips. The next lemma shows that this is optimal, even for more general polygons. The lemma is a slight generalization of a lemma by Lubiw and Pathak [11] on double chains of constant size.

Lemma 3.3. Let P be a polygon that 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 sub-triangulation, respectively. Then T_1 and T_2 have flip distance at least $4n - 4$.

The following result can be seen as a special case of [12, Proposition 1].

Lemma 3.4. Let P be a polygon that contains P_D and has $\langle u_n, \dots, u_1, l_1, \dots, l_n \rangle$ as part of its boundary. Let T_1 and T_2 be two triangulations that contain the upper and the lower extreme triangulation of P_D as a sub-triangulation, respectively. Consider any flip sequence σ from T_1 to T_2 and suppose there is no triangulation in σ containing a triangle with one vertex at the upper chain, the other vertex at the lower chain, and the third vertex at a point in the interior of the hourglass of P_D . Then $|\sigma| \geq (n - 1)^2$.

4 The Reduction

We reduce YRSA to POLYFLIP. Let S be a set of N sinks on an $n \times n$ grid with root at $(1, 1)$ (recall that n is polynomial in N). We construct a polygon P_D^* and two triangulations T_1, T_2 in P_D^* such that a shortest flip sequence from T_1 to T_2 corresponds to a shortest RSA for S . To this end, we will describe how to interpret any triangulation of P_D^* as a *chain path*, a path in the integer grid that starts at the origin and uses only edges that go north or east. It will turn out that flips in P_D^* essentially correspond to moving the endpoint of the chain path along the grid. We choose P_D^*, T_1 , and T_2 in such a way that a shortest flip sequence between T_1 and T_2 moves the endpoint of the chain path according to an Eulerian traversal of a shortest RSA for S . To force the chain path to

visit all sites, we use the observations from Section 3: the polygon P_D^* contains a double chain for each sink, so that only for certain triangulations of P_D^* it is possible to flip the double chain quickly. These triangulations will be exactly the triangulations that correspond to the chain path visiting the appropriate site.

4.1 The Construction

Our construction has two integral parameters, β and d . With foresight, we set $\beta = 2N$ and $d = nN$. We imagine that the sinks of S lie on a $\beta n \times \beta n$ grid, with their coordinates multiplied by β .

We take a double chain D with βn vertices on each chain such that the flip-kernel of D extends to the right of $l_{\beta n} u_{\beta n}$. We add a point z to that part of the flip-kernel, and we let P_D^+ be the polygon defined by $\langle l_1, \dots, l_{\beta n}, z, u_{\beta n}, \dots, u_1 \rangle$. Next, we add double chains to P_D^+ in order to encode the sinks. For each sink $s = (x, y)$, we remove the edge $l_{\beta y} l_{\beta y+1}$, and we replace it by a (rotated) double chain D_s with d vertices on each chain, such that $l_{\beta y}$ and $l_{\beta y+1}$ correspond to the last point on the lower and the upper chain of D_s , respectively. We orient D_s in such a way that $u_{\beta x}$ is the only point inside the hourglass of D_s and so that $u_{\beta x}$ lies in the flip-kernel of D_s ; see Fig. 4. We refer to the added double chains as *sink gadgets*, and we call the resulting polygon P_D^* . For β large enough, the sink gadgets do not overlap, and P_D^* is a simple polygon. Since the y -coordinates in S are pairwise distinct, there is at most one sink gadget per edge of the lower chain of P_D^+ . The precise placement of the sink gadgets is flexible, so we can make all coordinates polynomial in n ; see the full version for details.




Fig. 4. The sink gadget for a site (x, y) is obtained by replacing the edge $l_{\beta y} l_{\beta y+1}$ by a double chain with d vertices on each chain. The double chain is oriented such that $u_{\beta x}$ is the only point inside its hourglass and its flip-kernel. In our example, $\beta = 1$.

Next, we describe the source and target triangulation for P_D^* . In the source triangulation T_1 , the interior of P_D^+ is triangulated such that all edges are incident to z . The sink gadgets are all triangulated with the upper extreme triangulation. The target triangulation T_2 is similar, but now the sink gadgets are triangulated with the lower extreme triangulation.

To get from T_1 to T_2 , we must go from one extreme triangulation to the other for each sink gadget D_s . By Lemma 3.4, this requires $(d - 1)^2$ flips, unless the

flip sequence creates a triangle that allows us to use the vertex in the flip-kernel of D_s . In this case, we say that the flip sequence *visits* the sink s . For d large enough, a shortest flip sequence must visit each sink, and we will show that this induces an RSA for S of similar length. Conversely, we will show how to derive a flip sequence from an RSA. The precise statement is given in the following theorem.

Theorem 4.1. *Let $k \geq 1$. The flip distance between T_1 and T_2 w.r.t. P_D^* is at most $2\beta k + (4d - 2)N$ if and only if S has an RSA of length at most k .*

We will prove Theorem 4.1 in the following sections. But first, let us show how to use it for our NP-completeness result.

Theorem 4.2. *POLYFLIP is NP-complete.*

Proof. As mentioned in the introduction, the flip distance in polygons is polynomially bounded, so POLYFLIP is in NP. We reduce from YRSA. Let (S, k) be an instance of YRSA such that S lies on a grid of polynomial size. We construct P_D^* and T_1, T_2 as described above. This takes polynomial time (see the full version for details). Set $l = 2\beta k + (4d - 2)N$. By Theorem 4.1, there exists an RSA for S of length at most k if and only if there exists a flip sequence between T_1 and T_2 of length at most l . \square

4.2 Chain Paths

Now we introduce the *chain path*, our main tool to establish a correspondence between flip sequences and RSAs. Let T be a triangulation of P_D^+ (i.e., the polygon P_D^* without the sink gadgets, cf. Section 4.1). A *chain edge* is an edge of T between the upper and the lower chain of P_D^+ . A *chain triangle* is a triangle of T that contains two chain edges. Let e_1, \dots, e_m be the chain edges, sorted from left to right according to their intersection with a line that separates the upper from the lower chain. For $i = 1, \dots, m$, write $e_i = (u_v, l_w)$ and set $c_i = (v, w)$. In particular, $c_1 = (1, 1)$. 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 $\|c_{i+1} - c_i\|_1 = 1$. It follows that $c_1 c_2 \dots c_m$ is an x - and y -monotone path in the $\beta n \times \beta n$ -grid, beginning at the root. It is called the *chain path* for T . Each vertex of the chain path corresponds to a chain edge, and each edge of the chain path corresponds to a chain triangle. Conversely, every chain path induces a triangulation T of P_D^+ ; see Fig. 5. In the following, we let b denote the upper right endpoint of the chain path. The next lemma describes how the chain path is affected by flips; see Fig. 5.

Lemma 4.3. *Any triangulation T of P_D^+ uniquely determines a chain path, and vice versa. A flip in T corresponds to one of the following operations on the chain path: (i) move the endpoint b north or east; (ii) shorten the path at b ; (iii) change an east-north bend to a north-east bend, or vice versa.* \square




Fig. 5. A triangulation of P_D^+ and its chain path. Flipping edges to and from z moves the endpoint b along the grid. A flip between chain triangles changes a bend.

4.3 From an RSA to a Short Flip Sequence

Using the notion of a chain path, we now prove the “if” direction of Theorem 4.1.

Lemma 4.4. *Let $k \geq 1$ and A an RSA for S of length k . Then the flip distance between T_1 and T_2 w.r.t. P_D^* is at most $2\beta k + (4d - 2)N$.*

Proof. The triangulations T_1 and T_2 both contain a triangulation of P_D^+ whose chain path has its endpoint b at the root. We use Lemma 4.3 to generate flips inside P_D^+ so that b traverses A in a depth-first manner. This needs $2\beta k$ flips.

Each time b reaches a sink s , we move b north. This creates a chain triangle that allows the edges in the sink gadget D_s to be flipped to the auxiliary vertex in the flip-kernel of D_s . The triangulation of D_s can then be changed with $4d - 4$ flips; see Lemma 3.3. Next, we move b back south and continue the traversal. Moving b at s needs two additional flips, so we take $4d - 2$ flips per sink, for a total of $2\beta k + (4d - 2)N$ flips. \square

4.4 From a Short Flip Sequence to an RSA

Finally, we consider the “only if” direction in Theorem 4.1. Let σ_1 be a flip sequence on P_D^+ . We say that σ_1 *visits* a sink $s = (x, y)$ if σ_1 has at least one triangulation T that contains the chain triangle $u_{\beta x} l_{\beta y} l_{\beta y + 1}$. We call σ_1 a *flip traversal* for S if (i) σ_1 begins and ends in the triangulation whose corresponding chain path has its endpoint b at the root and (ii) σ_1 visits every sink in S . The following lemma shows that every short flip sequence in P_D^* can be mapped to a flip traversal.

Lemma 4.5. *Let σ be a flip sequence from T_1 to T_2 w.r.t. P_D^* with $|\sigma| < (d-1)^2$. Then there is a flip traversal σ_1 for S with $|\sigma_1| \leq |\sigma| - (4d - 4)N$.*

Proof. We show how to obtain a flip traversal σ_1 for S from σ . Let T^* be a triangulation of P_D^* . A triangle of T^* is an *inner triangle* if all its sides are diagonals. It is an *ear* if two of its sides are polygon edges. By construction, every inner triangle of T^* must have (i) one vertex incident to z (the rightmost vertex of P_D^+), or (ii) two vertices incident to a sink gadget (or both). In the latter case,

there can be only one such triangle per sink gadget. The weak (graph theoretic) dual of T^* is a tree in which ears correspond to leaves and inner triangles have degree 3.

Let D_s be a sink gadget placed between the vertices l_s and l'_s . Let u_s be the vertex in the flip-kernel of D_s . We define a triangle Δ_s for D_s . Consider the bottommost edge e of D_s , and let Δ be the triangle of T^* that is incident to e . By construction, Δ is either an ear of T^* or is the triangle defined by e and u_s . In the latter case, we set $\Delta_s = \Delta$. In the former case, we claim that T^* has an inner triangle Δ' with two vertices on D_s : follow the path from Δ in the weak dual of T^* ; while the path does not encounter an inner triangle, the next triangle must have an edge of D_s as a side. There is only a limited number of such edges, so eventually we must meet an inner triangle Δ' . We then set $\Delta_s = \Delta'$; see Fig. 6. Note that Δ_s might be $l_s l'_s u_s$.




Fig. 6. Triangulations of D_s in P_D^* with $\Delta_s = \Delta$ (left), and with Δ being an ear (red) and Δ_s an inner triangle (right). The fat tree indicates the dual.

For each sink s , let the polygon $P_{D_s}^{u_s}$ consist of the D_s extended by the vertex u_s (cf. Definition 3.2). Let T^* be a triangulation of P_D^* . We show how to map T^* to a triangulation T^+ of P_D^+ and to triangulations T_s of $P_{D_s}^{u_s}$, for each s .

We first describe T^+ . It contains every triangle of T^* with all three vertices in P_D^+ . For each triangle Δ in T^* with two vertices on P_D^+ and one vertex on the left chain of a sink gadget D_s , we replace the vertex on D_s by l_s . Similarly, if the third vertex of Δ is on the right chain of D_s , we replace it by l'_s . For every sink s , the triangle Δ_s has one vertex at a point u_i of the upper chain. In T^+ , we replace Δ_s by the triangle $l_s l'_s u_i$. No two triangles overlap, and they cover all of P_D^+ . Thus, T^+ is indeed a triangulation of P_D^+ .

Now we describe how to obtain T_s , for a sink $s \in S$. Each triangle of T^* with all vertices on $P_{D_s}^{u_s}$ is also in T_s . Each triangle with two vertices on D_s and one vertex not in $P_{D_s}^{u_s}$ is replaced in T_s by a triangle whose third vertex is moved to u_s in T_s (note that this includes Δ_s); see Fig. 7. Again, all triangles cover $P_{D_s}^{u_s}$ and no two triangles overlap.




Fig. 7. Obtaining T^+ and T_s from T^* .

Eventually, we show that a flip in T^* corresponds to at most one flip either in T^+ or in precisely one T_s for some sink s . We do this by considering all the possibilities for two triangles that share a common flippable edge. Note that by construction no two triangles mapped to triangulations of different polygons $P_{D_s}^{u_s}$ and $P_{D_t}^{u_t}$ can share an edge (with $t \neq s$ being another sink).

Case 1. We flip an edge between two triangles that are either both mapped to T^+ or to T_s and are different from Δ_s . This flip clearly happens in at most one triangulation.

Case 2. We flip an edge between a triangle Δ_1 that is mapped to T_s and a triangle Δ_2 that is mapped to T^+ , such that both Δ_1 and Δ_2 are different from Δ_s . This results in a triangle Δ'_1 that is incident to the same edge of $P_{D_s}^{u_s}$ as Δ_1 (for each such triangle, the point not incident to that edge is called the *apex*), and a triangle Δ'_2 having the same vertices of P_D^+ as Δ_2 . Since the apex of Δ_1 is a vertex of the upper chain or z (otherwise, it would not share an edge with Δ_2), it is mapped to u_s , as is the apex of Δ'_1 . Also, the apex of Δ'_2 is on the same chain of D_s as the one of Δ_2 . Hence, the flip affects neither T^+ nor T_s .

Case 3. We flip the edge between a triangle Δ_2 mapped to T^+ and Δ_s . By construction, this can only happen if Δ_s is an inner triangle. The flip affects only T^+ , because the new inner triangle Δ'_s is mapped to the same triangle in T_s as Δ_s , since both apexes are moved to u_s .

Case 4. We flip the edge between a triangle Δ of T_s and Δ_s . Similar to Case 3, this affects only T_s , because the new triangle Δ'_s is mapped to the same triangle in T^+ as Δ_s , since the two corners are always mapped to l_s and l'_s .

Thus, σ induces a flip sequence σ_1 in P_D^+ and flip sequences σ_s in each $P_{D_s}^{u_s}$ so that $|\sigma_1| + \sum_{s \in S} |\sigma_s| \leq |\sigma|$. Furthermore, each flip sequence σ_s transforms $P_{D_s}^{u_s}$ from one extreme triangulation to the other. By the choice of d and Lemma 3.4, the triangulations T_s have to be transformed so that Δ_s has a vertex at u_s at some point, and $|\sigma_s| \geq 4d - 4$. Thus, σ_1 is a flip traversal, and $|\sigma_1| \leq |\sigma| - N(4d - 4)$, as claimed. \square

In order to obtain a static RSA from a changing flip traversal, we use the notion of a *trace*. A *trace* is a domain 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 4.5. By Lemma 4.3, each triangulation in σ_1 corresponds to 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 flip in σ_1 that changes a bend, 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 . For example, the cost of a tree is twice the number of its edges, and the cost of an $a \times b$ rectangle is $ab + 2(a + b)$. An edge can be interpreted as a degenerated box, having two boundary sides and no interior. The following proposition is proved in the full version.

Proposition 4.6. *Let σ_1 be a flip traversal and R the trace of σ_1 . Then $\text{cost}(R) \leq |\sigma_1|$.*

Now we relate the length of an RSA for S to the cost of a covering trace for S , and thus to the length of a flip traversal. Since each sink (s_x, s_y) is connected in R to the root by a path of length $s_x + s_y$, traces can be regarded as generalized RSAs. In particular, we make the following observation.

Observation 4.7 *Let R be a covering trace for S that contains no boxes, and let A_{σ_1} be a shortest path tree in R from the root to all sinks in S . Then A_{σ_1} is an RSA for S . \square*

If σ_1 contains no flips that change bends, the corresponding trace R has no boxes. Then, R contains an RSA A_{σ_1} with $2|A_{\sigma_1}| \leq \text{cost}(R)$, by Observation 4.7. The next lemma shows that, due to the size of β , there is always a shortest covering trace for S that does not contain any boxes. See the full version for the proof.

Lemma 4.8. *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 $\text{cost}(R) \leq |\sigma_1|$.*

Now we can finally complete the proof of Theorem 4.1 by giving the second direction of the correspondence.

Lemma 4.9. *Let $k \geq 1$ and 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 an RSA for S of length at most k .*

Proof. Trivially, there always exists an RSA on S of length less than $2nN$, so we may assume that $k < 2nN$. Hence (recall that $\beta = 2N$ and $d = nN$),

$$2\beta k + 4dN - 2N < 2 \times 2N \times 2nN + 4nN^2 - 2N < 12nN^2 < (d - 1)^2,$$

for $n \geq 14$ and positive N . Thus, since σ meets the requirements of Lemma 4.5, we can obtain a flip traversal σ_1 for S with $|\sigma_1| \leq 2\beta k + 2N$. By Lemma 4.8 and Observation 4.7, we can conclude that there is an RSA A for S that has length at most $\beta k + N$. By Theorem 2.1, there is an RSA A' for S that is not longer than A and that lies on the Hanan grid for S . The length of A' must be 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 $n \times n$ grid has length at most k . \square

References

- [1] Z. Abel, B. Ballinger, P. Bose, S. Collette, V. Dujmović, F. Hurtado, S. Kominers, S. Langerman, A. Pór, and D. Wood. Every large point set contains many collinear points or an empty pentagon. *Graphs Combin.*, 27:47–60, 2011.
- [2] O. Aichholzer, W. Mulzer, and A. Pilz. Flip Distance Between Triangulations of a Simple Polygon is NP-Complete. *ArXiv e-prints*, 2012, 1209.0579. Extended abstract in Proc. 29th EuroCG, Braunschweig 2013.
- [3] P. Bose and F. Hurtado. Flips in planar graphs. *CGTA*, 42(1):60–80, 2009.
- [4] K. Culik II and D. Wood. A note on some tree similarity measures. *Inf. Process. Lett.*, 15(1):39–42, 1982.
- [5] D. Eppstein. Happy endings for flip graphs. *JoCG*, 1(1):3–28, 2010.
- [6] S. Hanke, T. Ottmann, and S. Schuierer. The edge-flipping distance of triangulations. *J.UCS*, 2(8):570–579, 1996.
- [7] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete Comput. Geom.*, 22:333–346, 1999.
- [8] F. Hwang, D. Richards, and P. Winter. *The Steiner Tree Problem*. Annals of Discrete Mathematics. North-Holland, 1992.
- [9] C. L. Lawson. Transforming triangulations. *Discrete Math.*, 3(4):365–372, 1972.
- [10] C. L. Lawson. Software for C^1 surface interpolation. In J. R. Rice, editor, *Mathematical Software III*, pages 161–194. Academic Press, NY, 1977.
- [11] A. Lubiw and V. Pathak. Flip distance between two triangulations of a point-set is NP-complete. In *Proc. 24th CCCG*, pages 127–132, 2012.
- [12] A. Pilz. Flip distance between triangulations of a planar point set is APX-hard. *ArXiv e-prints*, 2012, 1206.3179.
- [13] S. K. Rao, P. Sadayappan, F. K. Hwang, and P. W. Shor. The rectilinear Steiner arborescence problem. *Algorithmica*, 7:277–288, 1992.
- [14] W. Shi and C. Su. The rectilinear Steiner arborescence problem is NP-complete. In *Proc. 11th SODA*, pages 780–787, 2000.
- [15] D. Sleator, R. Tarjan, and W. Thurston. Rotation distance, triangulations and hyperbolic geometry. *J. Amer. Math. Soc.*, 1:647–682, 1988.
- [16] V. Trubin. Subclass of the Steiner problems on a plane with rectilinear metric. *Cybernetics*, 21:320–324, 1985.
- [17] J. Urrutia. Algunos problemas abiertos. In *Proc. IX Encuentros de Geometría Computacional*, pages 13–24, 2001.