The edge-flipping distance of triangulations

Institut für Informatik — Report 76*

Sabine Hanke *Thomas Ottmann*
Sven Schuierer

Institut für Informatik, Universität Freiburg
Am Flughafen 17, 79110 Freiburg, Germany
Email: schuiere@informatik.uni-freiburg.de

February 1996

*This paper will be held at the *European Workshop on Computational Geometry*, Münster, Germany, in March 1996.

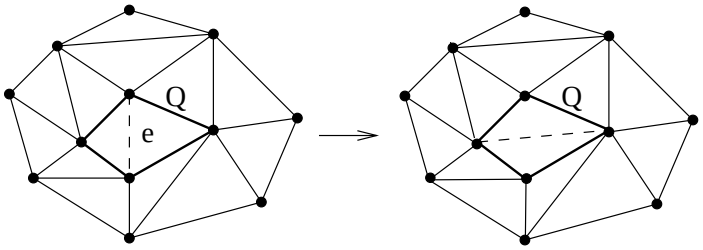

Figure 1: An edge-flipping operation replacing e

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 [Sleator et al. 88] 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 set of points in the plane by showing that not 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. Furthermore in the case of triangulations of convex polygons we present a combinatorical proof of a weaker lower bound of $\frac{3}{2}n - 5$ with the aid of two triangulations.

1 Introduction

Let S be a set of points in the plane, and T a triangulation of S. 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 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.

A natural idea to compare the structures of triangulation of the same set of points is to compute the *edge-flipping distance* which is defined as the least number of admissible edge-flipping operations to transform a triangulation into another. Of course, it must be shown that such a transformation is always possible.

Already in 1936 Wagner discussed a very similar problem 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 also 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-

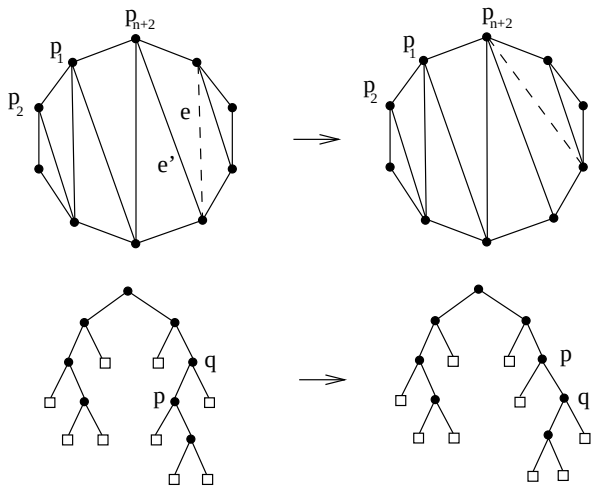

Figure 2: An edge-flipping operation and the corresponding rotation

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 exact 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 show 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 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'. In 1994 Cai and Hirsch [Cai, Hirsch 94] extend the results of Sleator, Tarjan, and Thurston to rotation distance problems of triangulations of planar surfaces. They give 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.

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 $O(n^2)$ 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 least $2n - 6$ 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 $O(n^2)$ 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. (Where we say that two different edges intersect, iff they intersect in their interiors.) And we 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:

Theorem: *Let S be a set of n points in the plane, T_1 and T_2 two triangulations of S, and let $\#(T_1, T_2)$ be the number of intersections of T_1 with T_2. With $flipdist(T_1, T_2)$ denote the edge-flipping distance between T_1 and T_2. 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 T_1 and T_2 respectively.

Proof: 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 T_1 and T_2 respectively. To prove the first inequality we present an algorithm that computes a sequence of edge-flipping operations with maximal length $\#(T_1, T_2)$ that transforms T_1 into T_2. It can be described as follows:

2.1 The algorithm

$\#(e, T)$	is a function returning the number of intersections between the edge e and the triangulation T
$diag(e, T)$	is a function returning the second diagonal of the quadrilateral around the edge e in the triangulation T

```
{
 init stack S;
 T'_1 := T_1; T'_2 := T_2;
 while ( T'_1 ≠ T'_2 ) do
 {
  while( there exists a diagonal e of a convex quadrilateral in T'_i with #(e, T'_j) = 1,
         where i, j ∈ {1,2}, i ≠ j ) do
  {
   Let e' be an edge in T'_j that intersects such an edge of T'_i;
   e'' := diag(e', T'_j);
   flip e' ⟶ e'' in T_j';
   if ( j = 1 ) then output( e' ⟶ e'' )
                else S.push(e'' ⟶ e' );
  }
  if ( T'_1 ≠ T'_2 ) then
  {
   Let e' be an edge of T'_i, i ∈ {1,2}, so that the number of intersections decreases
   at the most by flipping e';
   e'' := diag(e', T'_i);
   flip e' ⟶ e'' in T'_i;
   if ( i = 1 ) then output( e' ⟶ e'' )
                else S.push( e'' ⟶ e' );
  }
 }
```

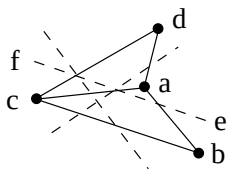

Figure 3: All edges that intersect ac also intersect bc and dc

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

2.2 Proof of correctness

First we 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 at least of one:
Without loss of generality let e be an edge of T_1' that has only one intersection with the edges of the triangulation T_2', i.e. $\#(e, T_2') = 1$. Let e' be the edge of T_2' that intersect e, and $e'' = diag(e', T_2')$ be the second diagonal of the quadrilateral around e in the triangulation T_2'. The equation $\#(e, T_2') = 1$ implies that e is an edge of T_1' not containing in T_2'. Let $e = ac$, that means a and c are the endpoints of the edge e, and let $a'bc'd$ be the quarilateral around $e' = bd$ in T_2'. Because the boundary edges $a'b$, bc', $c'd$, and $a'd$ of the quadrilateral do not intersect the edge e, and because the endpoints of e can not lie in the interior of the quadrilateral, it follows, that $a'bc'd$ must be convex and $a'bc'd = abcd$, in particular $e'' = e$.

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 intersect by at least two edges of the other triangulation. We show that in this case there is always an edge-flipping operation that decreases the number of intersections between the triangulations T_1' and T_2':
Without loss of generality we discuss this problem only for the triangulation T_1'.
Let $\#(e, e', T)$ denote the number of edges in T that intersect the edge e as well as edge e', and $\#_p(T, e)$:= 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. $\#(ac, T_2') = max\{\#(e, T_2') \mid e \text{ is an edge of } T_1'\}$.*

Proof: Let $abcd$ be a quadrilateral in T_1' with diagonal ac so that ac has a maximum number of intersections with T_2'. Assume that $abcd$ is not convex and the angle at point a inside the quadrilateral is larger than π. Then all edges that intersect the edge ac also intersect the edges bc and dc [Fig. 3].
If there is an edge ef in T_2' that intersects the edge dc with either $e = b$ or that intersects the edge ab as well, then T_2' can not contain an edge that intersects ad as well as bc, so $\#(ad, bc, T_2') = \#_d(T_2, bc') = 0$ and $\#(ac, T_2') = \#(bc, dc, T_2') + \#(ab, dc, T_2') + \#_b(T_2', dc)$. Hence, all edges that intersect ac also intersect dc.

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 dc. This edge does not intersect ac and therefore $\#(ac, T_2') \leq \#(dc, T_2') + 1$ which contradics the maximality of $\#(ac, T_2')$. For symmetry reasons there exists no edge ef in T_2' that intersects the edge dc with either $e = b$ or that intersects the edge ab as well. Therefore, it follows

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

In particular bc and dc can not be boundary edges of the triangulations. If we next take the

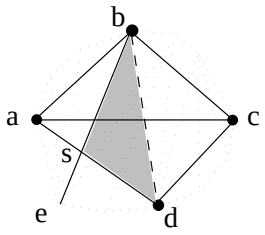

Figure 4: Definition of *bds* in [Lemma 2]

quadrilateral in T_1' with diagonal bc then this quadrilateral is either convex or the next neighbouring quadrilateral has a diagonal with a maximal number of intersections, and so on. Finally if such a quadrilateral contains a boundary edge of the triangulation, it 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 be that intersects ad 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 be an edge of T_2' that intersects ad. This implies that all edges of T_2' that intersects bd cut the edge ad as well, because otherwise at least one point of the set S would lie inside the triangle bds where s is the intersection point of be with ad [Fig. 4], and so in the interior of the quadrilateral $abcd$. Since be does not intersect bd, and because of the assumption that ac has the maximum number of intersections with T_2', it follows that

$$\#(bd, T_2') < \#(ad, T_2') \leq \#(ac, T_2')$$

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

Lemma 3: *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 no edge be that intersects ad or cd and no edge dg that intersects ab or bc, then either the edge-flipping operation $ac \longrightarrow bd$ decrements the number of intersections between the triangulations T_1' and T_2', or T_1' contains another quadrilateral with maximal diagonal which complies with the conditions of [Lemma 2].*

Proof: If the edge-flipping operation $ac \longrightarrow bd$ decrements 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 decrement the number of intersections: Because the triangulation T_2' does not contain edges be and dg like above, there exist a triangle bef adjacent to b, where ef intersects ab as well as bc, and a triangle dgh adjacent to d, where gh intersects ad as well as cd [see Fig. 5]. Therefore ef and gh do not intersect the edge ac. Because of the assumption $\#(ac, T_2') \geq max\{\#(ac, T_2'), \#(bc, T_2', \#(cd, T_2'), \#(ad, T_2')\}$ T_2' also contains triangles apq and cxy, where pq intersects ab as well as ad 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 intersects the edge ac must intersect bc as well (analogous to 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 ad and bc [see Fig. 5], then follows $\#(ab, cd, T_2') = 0$, and

$$\#(ac, T_2') \quad = \quad \#(ab, ad, T_2') + \#(bc, dc, T_2') + \#(ad, bc, T_2') \tag{1}$$

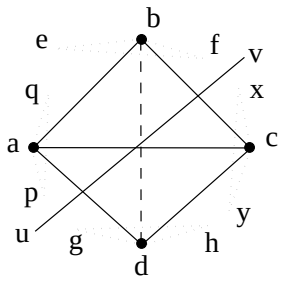

Figure 5: if T_2' doesn't contain be and dg like above and $ac \to bd$ doesn't decrement the intersections

$$\begin{aligned}
\#(bc, T_2') &= \#(ab, bc, T_2') + \#(bc, dc, T_2') + \#(ad, bc, T_2') && (2)\\
\#(ad, T_2') &= \#(ab, ad, T_2') + \#(ad, dc, T_2') + \#(ad, bc, T_2') && (3)\\
\#(bd, T_2') &= \#(ab, bc, T_2') + \#(ad, dc, T_2') + \#(ad, bc, T_2') && (4)
\end{aligned}$$

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

$$\begin{aligned}
(1) \geq (2): &\quad \#(ab, ad, T_2') \geq \#(ab, bc, T_2') && (5)\\
(1) \geq (3): &\quad \#(bc, dc, T_2') \geq \#(ad, dc, T_2') && (6)
\end{aligned}$$

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

$$\#(ab, ad, T_2') + \#(bc, dc, T_2') \leq \#(ab, bc, T_2') + \#(ad, dc, T_2'), \qquad (7)$$

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

$$\#(ab, ad, T_2') + \#(bc, dc, T_2') = \#(ab, bc, T_2') + \#(ad, dc, T_2'). \qquad (8)$$

Therefore

$$\#(ab, bc, T_2') \overset{(5)}{\leq} \#(ab, ad, T_2') \overset{(8)}{=} \#(ab, bc, T_2') + \#(ad, dc, T_2') - \#(bc, dc, T_2') \qquad (9)$$

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

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

Now take 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] occur so that an edge-flipping operation replacing this diagonal decreases the number of intersections between T_1' and T_2'. □

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 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 putted 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. □

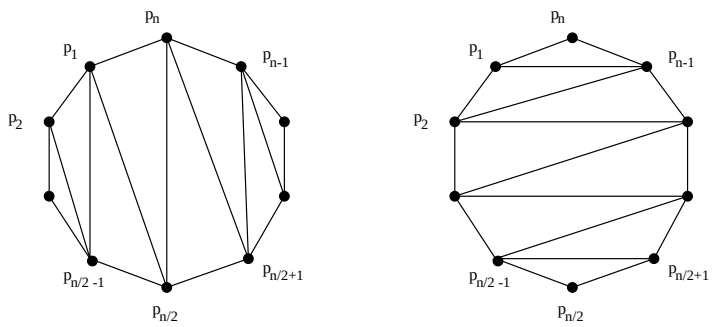

Figure 6: The triangulations T_1^n and T_2^n

3 A lower bound example

In the context of computing the rotation distance of binary trees Sleator, Tarjan, and Thurston [Sleator et al. 88] 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 this bound is tight, but they do not provide an example of a pair of triangulations of a convex polygon with edge-flipping distance $2n-10$. In the following we present a combinatorical proof of a weaker lower bound of $\frac{3}{2}n-5$, solving an open problem of Knuth [Knuth 93], who has suspecected that the following two triangulations have quite a large edge-flipping distance.

Let n be even, $P = \{p_1, ..., p_n\}$ a convex polygon, and let T_1^n and T_2^n be the triangulations of P [see Fig. 6] with the following interior edges:
$T_1^n : ..., p_2p_{\frac{n}{2}-1}, p_{\frac{n}{2}-1}p_1, p_1p_{\frac{n}{2}}, p_{\frac{n}{2}}p_n, p_np_{\frac{n}{2}+1}, p_{\frac{n}{2}+1}p_{n-1}, \cdots$
$T_2^n : p_1p_{n-1}, p_{n-1}p_2, p_2p_{n-2}, p_{n-2}p_3, ..., p_{\frac{n}{2}-1}p_{\frac{n}{2}+1}$

Claim: *T_1^n and T_2^n have edge-flipping distance $\frac{3}{2}n-5$.*

Proof: We will reduce the problem to a polygon of $n-2$ points to determine the edge-flipping distance between T_1^n and T_2^n recursively.
Let $F = (f_1, f_2, ..., f_m)$ be a sequence of edge-flipping operations that transforms T_1^n into T_2^n. Assume that F generates no edge of T_2^n twice. We will show that for all such sequences of edge-flipping operations there exists a sequence of edge-flipping operations $G = (g_1, ..., g_M)$, so that:

(1) G transforms T_1^n into T_2^n

(2) $g_1 = p_{\frac{n}{2}}p_n \longrightarrow p_1p_{\frac{n}{2}+1}$
$g_2 = p_1p_{\frac{n}{2}} \longrightarrow p_{\frac{n}{2}-1}p_{\frac{n}{2}+1} \in T_2^n$
$g_3 = p_{\frac{n}{2}+1}p_n \longrightarrow p_1p_{n-1} \in T_2^n$

(3) $M \leq m$, i.e. G does not contain more edge-flipping operations than F

In particular there exists an optimal sequence of edge-flipping operations G with properties (1) and (2), *optimal* means to be a sequence of minimal length. Let T' be the triangulation that is obtained from T_1^n by using g_1, g_2, and g_3 [Fig. 7]. Then $flipdist(T_1^n, T_2^n) = flipdist(T_1^n, T') + flipdist(T', T_2^n)$, and it follows from [Sleator et al. 88] that every optimal sequence of edge-flipping operations that transforms T' into T_2^n does not move the edges $p_{\frac{n}{2}-1}p_{\frac{n}{2}+1}$ and p_1p_n. So if we renumber the points $p_1, ..., p_{\frac{n}{2}-1}, p_{\frac{n}{2}+1}, ..., p_{n-1}$ in a suitable way, we can take T' restricted to $P - \{p_{\frac{n}{2}}, p_n\}$ as T_1^{n-2} and T_2^n restricted to $P - \{p_{\frac{n}{2}}, p_n\}$ as T_2^{n-2}. For $n = 4$ $\quad flipdist(T_1^4, T_2^4) =$

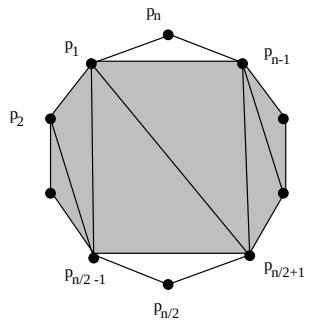

Figure 7: The triangulation T'

1. Therefore we can calculate the edge-flipping distance $flipdist(T_1^n, T_2^n)$ between T_1^n and T_2^n recursively:

$$\begin{aligned} flipdist(T_1^n, T_2^n) &= 3 + flipdist(T_1^{n-2}, T_2^{n-2}) \\ &= 3 + ... + 3 + flipdist(T_1^4, T_2^4) \\ &= \frac{n-4}{2} \cdot 3 + 1 \\ &= \frac{3}{2}n - 5 \end{aligned}$$

Costruction of G: First we show that the given sequence $F = (f_1, ..., f_m)$ of edge-flipping operations induces a sequence $F' = (f_1', f_2', ..., f_{m+3}')$ of edge-flipping and empty operations, and then we obtain the sequence G by removing the empty operations from F'. In order to do so we set $f_1' := p_{\frac{n}{2}}p_n \longrightarrow p_1p_{\frac{n}{2}+1}$, $f_2' := p_1p_{\frac{n}{2}} \longrightarrow p_{\frac{n}{2}-1}p_{\frac{n}{2}+1} \in T_2^n$, and $f_3' := p_{\frac{n}{2}+1}p_n \longrightarrow p_1p_{n-1} \in T_2^n$. For simplification we insert three empty operations in F at the beginning and renumber the flip-operations, i.e. $F = (f_1, f_2, ..., f_{m+3})$ with $f_1 = f_2 = f_3 = \emptyset$.
Now we explain how to obtain $f_4', ...f_{m+3}'$ from $f_4, ...f_{m+3}$: We define a map h from the set of points $\{p_1, ..., p_n\}$ of P to the set of points $S := \{p_1, ..., p_{\frac{n}{2}-1}, p_{\frac{n}{2}+1}, ..., p_{n-1}\}$ by $h(p_{\frac{n}{2}}) := p_{\frac{n}{2}+1}, h(p_n) := p_1$, and h restricted to S is the identity on S. Then we extend this map to edges of triangulations of P by $h(p_ip_j) := h(p_i)h(p_j)$ for all $i, j \in \{1, .., n\}$. In this way any edge excluding p_1p_n and $p_{\frac{n}{2}}p_{\frac{n}{2}+1}$ is mapped to an edge between points of S. More precisely, if we define a map H on triangulations T of P by $H(T)$ contains $h(e)$ for all edges e of T, then $H(T)$ is a triangulation of $P - \{p_{\frac{n}{2}}, p_n\}$, because h keeps all triangles that are not adjacent to $p_{\frac{n}{2}}$ or p_n fixed, and each triangle adjacent to p_n (or $p_{\frac{n}{2}}$) is mapped to a corresponding triangle or edge respectively adjacent to p_1 (or $p_{\frac{n}{2}+1}$). In the following for all triangulations T of P we identify $H(T)$ with the triangulation of P that contains $H(T)$ and the triangles $p_1p_{n-1}p_n$ and $p_{\frac{n}{2}-1}p_{\frac{n}{2}}p_{\frac{n}{2}+1}$. So by definition $H(T_1^n) = T'$ [Fig. 7], and $H(T_2^n) = T_2^n$. Let $(T^{(0)}, T^{(1)}, ..., T^{(m)})$ be the sequence of triangulations induced by $f_4, ...f_{m+3}$ with $T^{(3)} = T_1^n$, $T^{(m+3)} = T_2^n$, and $T^{(i-1)} \longrightarrow T^{(i)}$ by using f_i. Then by H we obtain a sequence $(H(T^{(3)}), H(T^{(3)}), ..., H(T^{(m+3)}))$ of triangulations with starting triangulation $T' = H(T^{(3)})$ and final triangulation $T_2^n = H(T^{(m+3)})$. Now we show, that for all $i = 4, ..., m+3$ there is an edge-flipping or empty operation f_i' induced by f_i so that $H(T^{(i-1)}) \longrightarrow H(T^{(i)})$ by using f_i'. Let ac be the edge in $T^{(i-1)}$ that is flipped by f_i into the edge bd. If $h(a)h(b)h(c)h(d)$ is a non-degenerated quadrilateral in $H(T^{(i-1)})$, then clearly $f_i' := h(ac) \longrightarrow h(bd)$ transforms $H(T^{(i-1)})$ into $H(T^{(i)})$. Else $h(a)h(b)h(c)h(d)$ forms a triangle or an edge in $T^{(i-1)}$, since p_1p_n and $p_{\frac{n}{2}}p_{\frac{n}{2}+1}$ are the only two edges that are mapped to a point, and therefore $h(a) \neq h(c)$ and $h(b) \neq h(d)$. Hence $h(ac), h(bd) \in \{h(ab), h(bc), h(cd), h(da)\}$, and so $H(T^{(i-1)}) = H(T^{(i)})$, and we define $f_i' := \emptyset$. In this way the given sequence $(f_1, ..., f_{m+3})$ of edge-flipping operations induces a sequence $F' = (f_1', ..., f_{m+3}')$ of edge-flipping and empty operations with properties (1) and (2).
Because $H(T^{(3)})$ has the edges $p_{\frac{n}{2}-1}p_{\frac{n}{2}+1}$ and p_1p_{n-1} with T_2^n in common in contrast to $T^{(3)}$,

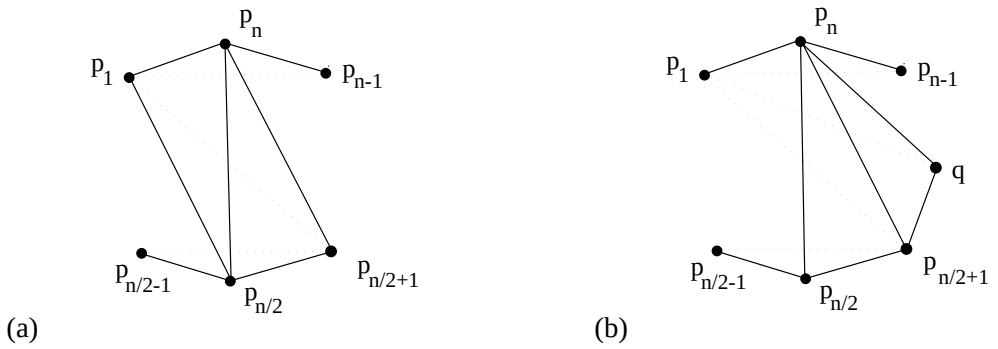

Figure 8: Occurence of the third empty operation

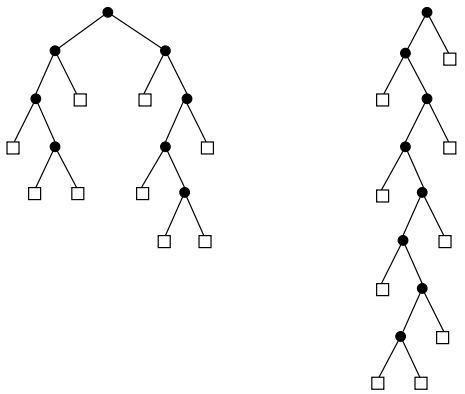

Figure 9: Two binary trees with rotation distance $\frac{3}{2}n - 2$

it is clear that F' must contain two empty operations. But F' contains at least a further empty operation: If $p_{\frac{n}{2}}p_n$ is flipped before $p_1p_{\frac{n}{2}}$ and $p_{\frac{n}{2}+1}p_n$, then the edge-flipping operation that removes $p_{\frac{n}{2}}p_n$ induces the empty operation, since h maps the quadrilateral $p_1p_{\frac{n}{2}}p_{\frac{n}{2}+1}p_n$ to the edge $p_1p_{\frac{n}{2}+1}$ [Fig. 8 a]. Else without loss of generality $p_{\frac{n}{2}+1}p_n$ is flipped before $p_{\frac{n}{2}}p_n$. Observe the edge-flipping operation f that transforms $p_{\frac{n}{2}+1}p_n$ into an edge $p_{\frac{n}{2}}q$ [Fig. 8 b]. Since $h(p_{\frac{n}{2}+1}p_n) = h(p_{\frac{n}{2}}p_n) = p_1p_{\frac{n}{2}+1}$, the operation f' induced by f is the empty operation. Therefore F' contains at least three empty operations. So if we get the sequence G by removing the empty operations from F', it does not contain more edge-flipping operations than F. Because of that G has the wanted properties (1) - (3). □

Lemma: *The example above gives an example of two binary trees of size n that have rotation distance* $\frac{3}{2}n - 2$. [Fig. 9]

4 Conclusion

We have introduced the *edge-flipping distance* between two triangulations T_1 and T_2 of the same set of points in the plane, and we have given 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, if it is possible, always to flip an edge such, that an edge of T_2 is generated. In the case of triangulations of convex polygons this strategy is very

useful. Sleator, Tarjan, and Thurston [Sleator et al. 88] have shown that through such a kind of flip the edge-flipping distance between the resulting triangulation and the final triangulation decreases by one. 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 occurence 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 [Sleator et al. 88] 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 triangulation into each other the number of intersections between the generated triangulation during the transformation and the final triangulation does not need to decrease after each step. There are a lot of examples of pairs of triangulations and optimal sequences of edge-flipping operations where the number of intersections increases by some edge-flipping operations during the transformation.

An open question is what the structure of a given sequence of edge-flipping operations must have to be of minimal length.

References

[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.

[Knuth 93] Knuth, D.: "Computer Musings: The Associative Law, or The Anatomy of Rotations in Binary Trees"; The Distinguished Lecture Series VII, University Video Communications, Stanford (1993).

[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.