

Flip Distance Between Two Triangulations of a Point-Set is NP-complete

Anna Lubiw*

Vinayak Pathak*

Abstract

Given two triangulations of a convex polygon, computing the minimum number of flips required to transform one to the other is a long-standing open problem. It is not known whether the problem is in P or NP-complete. We prove that two natural generalizations of the problem are NP-complete, namely computing the minimum number of flips between two triangulations of (1) a polygon with holes; (2) a set of points in the plane.

1 Introduction

Given a triangulation in the plane, a *flip* operates on two triangles that share an edge and form a convex quadrilateral. The flip replaces the diagonal of the convex quadrilateral by the other diagonal to form two new triangles. A sequence of flips can transform any triangulation to any other triangulation—this is true for triangulations of a convex polygon, and more generally for triangulations of a point set, and for triangulations of a polygon with holes.

In this paper we investigate the complexity of computing the *flip distance*, which is the minimum number of flips to transform one triangulation to another. This is particularly interesting for convex polygons, where the flip distance is the rotation distance between two binary trees (see below).

The main result of our paper is that it is NP-complete to compute the flip distance between two triangulations of a polygon with holes, or of a set of points in the plane.

After submitting this paper, we learned that Pilz [20] independently proved the same result. The differences between our proofs are discussed later on.

1.1 Flip distance and rotation distance

Binary search trees are a widely used data structure, and *rotations* are the basic operations used to balance them. Despite the importance of rotations, the complexity of computing the minimum number of rotations to convert one labelled binary search tree to another, called the “rotation-distance”, has been open since at least 1982 [6]. It is not known if the problem is NP-complete.

There is a bijection between binary trees with $n - 1$ labeled leaves and triangulations of an n -vertex convex polygon. Moreover, a rotation in the tree corresponds to a flip in the polygon. Thus, computing the rotation distance between two trees is exactly equivalent to computing the flip distance between two triangulations of a convex polygon. See [23].

1.2 Generalizations and related work

Flips have been studied in the geometric setting for triangulations of point sets and of polygons. In this context, a convex polygon is equivalent to a point set in convex position. The former generalizes to simple polygons, and the latter to planar point sets. Both of these are contained in the most general case of a polygon with holes (a “polygonal region”), so long as we consider a point as a one vertex polygonal hole. There is a survey on flips by Bose and Hurtado [4]. It also covers flips in the combinatorial setting of maximal planar graphs, which we will not discuss. Flips are often studied in terms of the *flip graph* which has a vertex for every triangulation and an edge when two triangulations differ by one flip, see e.g., [10].

The foundational result is that the flip graph is connected. This was proved first by Lawson [14] for the case of point sets. He then re-proved the result [13] by arguing that any triangulation can be flipped to the Delaunay triangulation, which then acts as a “canonical” triangulation from which any other triangulation can be reached. The constrained Delaunay triangulation can be used in the same way to argue that any polygonal region has a connected flip graph [2]. For more direct proofs see [9, 12, 18].

Regarding the number of flips needed to transform one triangulation to another, flipping via the [constrained] Delaunay triangulation takes $O(n^2)$ flips—in fact, a more exact bound is the number of visibility edges, see [2]. Hurtado, Noy and Urrutia [12] proved that $\Omega(n^2)$ flips may be required even for triangulations of a polygon. For the case of a convex polygon, Sleator et al. [23] proved that for large values of n , the flip distance between two triangulations of an n -gon is at most $2n - 10$, and that $2n - 10$ flips are sometimes necessary.

The problem of computing the exact flip distance between two given triangulations is especially interesting for convex polygons, as mentioned above. Lucas [16] gave a polynomial time algorithm for special cases. The

*Cheriton School of Computer Science, University of Waterloo, Waterloo, Canada {alubiw,vpathak}@uwaterloo.ca

best approximation factor is trivially 2, and can be improved in some special cases [15]. Recently it was proved that the problem is fixed parameter tractable in the flip distance [5]. Attempts have also been made to compute good upper and lower bounds on the flip distance efficiently. See, for example, [1, 19, 17, 7].

The more general problem of computing the flip distance between two triangulations of a point set is stated as an open problem in the survey by Bose and Hurtado [4], and the book by Devadoss and O'Rourke [8, Unsolved Problem 12]. Hanke et al. [11] proved that the flip-distance is upper bounded by the total number of intersections between the overlap of the initial and final triangulations. Eppstein [10] provided an algorithm to compute a lower bound on the flip-distance efficiently. He also showed that the lower bound is equal to the flip-distance for certain special kinds of point-sets.

2 Triangulations of polygonal regions

Theorem 1 *The following problem is NP-complete: Given two triangulations of a polygon with holes and a number k , is the flip distance between the two triangulations at most k ?*

2.1 Proof idea

Note that the problem lies in NP. We prove hardness by giving a polynomial time reduction from vertex cover on 3-connected cubic planar graphs [3, 24], which is known to be NP-complete [3, 24].

The idea is to take a planar straight-line drawing of the graph and create a polygonal region by replacing each edge by a “channel” and each vertex by a “vertex gadget”. We then construct two triangulations of the polygonal region that differ on the channels, and show that a short flip sequence corresponds to a small vertex cover in the original graph.

We begin by describing channels and their triangulations, because this gives the intuition for the proof. A *channel* is a polygon that consists of two 7-vertex reflex chains joined by two *end* edges, as shown in Figures 1(a) and 1(b). Note that every vertex on the upper reflex chain sees every vertex on the lower reflex chain and vice versa. We identify two triangulations of a channel: a *left-inclined triangulation* as shown in Figure 1(a); and a *right-inclined triangulation* as shown in Figure 1(b).

A channel is the special case $n = 7$ of the polygons H_n of Hurtado et al. [12]. They prove in Theorem 3.8 that the flip distance between the right-inclined and left-inclined triangulations of H_n is $(n - 1)^2$. We include a different proof in order to generalize:

Property 1 *Transforming a left-inclined triangulation of a channel to a right-inclined triangulation takes at least 36 flips.*

Proof. In any triangulation of a channel, each edge of the upper reflex chain is in a triangle whose apex lies on the bottom reflex chain. This apex must move from lower right (B_7) to lower left (B_1), in order to transform the left-inclined triangulation to the right-inclined triangulation. Similarly, each edge of the lower reflex chain is in a triangle whose apex lies on the upper reflex chain, and must move from upper left to upper right. However, one flip can only involve one edge of the upper chain and one edge of the lower chain (no other 4 vertices form a convex quadrilateral), and thus can only move one upper and one lower apex, and only by one vertex along the chain. Twelve triangles times six apex moves per triangle divided by two apex moves per flip gives a lower bound of 36 flips. \square

We now show that the number of flips goes down if a channel has a *cap*, an extra vertex that is visible to all the channel vertices, as shown in Figure 1(c).

Property 2 *The flip distance from a left-inclined to a right-inclined triangulation of a capped channel is 24.*

Proof. The “canonical” triangulation shown in Figure 1(d) is 12 flips away from both the left-inclined and the right-inclined triangulations of a capped channel: To flip the left-inclined triangulation to the canonical triangulation, flip edges A_1B_1, \dots, A_1B_7 followed by edges A_2B_7, \dots, A_6B_7 in that order. Similarly for the right-inclined triangulation.

For the lower bound, we follow the same idea as above. In any triangulation, each edge of the upper [lower] reflex chain is in a triangle whose apex is either the cap or a vertex of the lower [upper] chain. There are only two kinds of flips: (1) a flip involving the cap vertex, an edge of one chain, and a vertex of the other chain; and (2) a flip involving one edge of each chain. A flip of type (1) moves the apex of only one triangle, and moves the apex to or from the cap. If a triangle is altered by a flip of type (1) then at least two such flips are required, one to move the apex to the cap and one to move the apex from the cap. If a triangle is only altered by flips of type (2), then, as above, it costs 3 flips to get the apex to its destination. Thus the 12 triangles require at least 24 flips. \square

We now elaborate on the idea of our reduction. We create a polygonal region by replacing each edge in the planar drawing by a channel, and each vertex by a vertex gadget. We make two triangulations of the polygonal region. In triangulation T_1 all edge channels are left-inclined and in T_2 all edge channels are right-inclined. The triangulations are otherwise identical. We design vertex gadgets so that making a few flips in a vertex gadget creates a cap for a channel connected to it. Since transforming a channel from left-inclined to right-inclined is less costly if it is capped, the minimum flip




Figure 1: Channels

sequence that transforms all the channels is obtained by choosing the smallest set of vertices that covers all the edges and using them to cap all the channels. Thus, intuitively, a minimum flip sequence corresponds to a minimum vertex cover.

One complication is that we cannot construct a vertex gadget for a *sharp* vertex—a vertex of degree 3 where one of the three incident angles in the planar drawing is $\geq \pi$. Therefore, we first show how to eliminate sharp vertices. Let G be our given 3-connected cubic planar graph. Using a result of Rote and Bárány [21], we can find, in polynomial time, a *strictly convex* drawing of G on a polynomial-sized grid. *Strictly convex* means that each face is a strictly convex polygon. Thus the only sharp vertices of this drawing are the vertices of the outer face. We replace each sharp vertex v by a 3-vertex chain v_1, v_2, v_3 as shown in Figure 2. We claim that G has a vertex cover of size $\leq k$ if and only if the modified graph has a vertex cover of size $\leq k + t$, where t is the number of vertices on the outer face of G . This is because any minimum vertex cover of the modified graph can be adjusted to use either $\{v_1, v_3\}$ (corresponding to v being in the vertex cover of G), or $\{v_2\}$ (corresponding to v not being in the vertex cover of G).




Figure 2: Eliminating sharp vertices

We remark that Pilz’s independent NP-hardness reduction [20] is from general (non-planar) vertex cover. His construction begins with the same channel gadgets, but then uses channels that overlap geometrically while flipping independently.

2.2 Details of the reduction

For the remainder of the proof we will assume that we have a graph G with vertices of degree 2 and 3, and a straight-line planar drawing, Γ , of the graph on a polynomial sized grid with no sharp vertices.

We define the *narrow* and *wide* mouths of a channel as shown in Figure 1(e). Any point inside the narrow mouth but outside the channel can be a potential cap for the channel. We show below that a vertex outside the wide mouth does not reduce the flip distance.

We now describe the triangulated vertex gadgets. See Figures 3(a) and 3(b). Each of the 2 or 3 channels attached to the vertex gadget will have one potential cap. We place a convex quadrilateral $CDEF$ with diagonal

CE , called the *lock*, that separates each channel from its potential cap. Thus the lock CE must be flipped, or “unlocked”, in order to cap any channel.

For the degree-2 gadget (see Figure 3(a)), place point C in the smaller angular sector (of angle $< \pi$) between the two channels, so that C is outside the wide mouths of both channels. Place points D , E , and F in the other angular sector, with D inside channel 1’s narrow mouth and outside channel 2’s wide mouth, E outside the wide mouth of both channels, and F inside channel 2’s narrow mouth and outside channel 1’s wide mouth. Triangulate as shown. Thus D is a potential cap for channel 1 and F is a potential cap for channel 2.

For the degree-3 gadget (see Figure 3(b)), note that because the vertex is not sharp, the mouth of each channel exits between the other two channels. We place vertices in the angular sectors as shown in the figure. Place D inside the intersection of the narrow mouths of channels 1 and 2, and outside the wide mouth of channel 3. Place F inside channel 3’s narrow mouth and outside channel 1 and 2’s wide mouths. Place C and E outside the wide mouths of all the channels. Triangulate as shown. Thus D is a potential cap for both channel 1 and 2 and F is a potential cap for channel 3.

Observe that every channel is blocked from its unique potential cap by exactly 3 edges. (For example, in Figure 3(b), channel 1 is separated from its potential cap D by edges FA , FE , and CE .) Observe furthermore that for each vertex gadget, the sets of blocking edges of the channels have one edge in common, namely the locking edge CE , and are otherwise disjoint. These properties are crucial for correctness.

We will say that a vertex gadget is *locked* if the diagonal CE exists and *unlocked* otherwise. We first show what is possible with unlocked vertex gadgets.

Property 3 *If we unlock a vertex gadget then, for each channel attached to it, there is a sequence of 28 flips that transforms the channel triangulation and returns the vertex gadget to its (unlocked) state.*

Proof. We first claim that there is a 2-flip sequence that caps the channel. We enumerate the possibilities (refer to Figure 3). Note that we handle channels one at a time, not simultaneously. For the degree-2 gadget: flip CF followed by CA for channel 1; flip CD followed by CB' for channel 2. For the degree-3 gadget: flip FE followed by FA for channel 1; flip CF followed by CA' for channel 2; flip ED followed by EA'' for channel 3. Once the channel is capped, we can transform the left-inclined triangulation to the right-inclined triangulation in 24 flips by Property 2. Then we undo the 2 flips that capped the channel. The total number of flips is 28. \square

Next we give lower bounds on the number of flips. First, note that the proof of Property 1 carries over to:

Property 4 *Transforming a left-inclined triangulation of a channel to a right-inclined triangulation takes at least 36 flips even in the presence of other vertices, so long as the other vertices lie outside the wide mouths at either end of the channel.*

We now consider what happens when we unlock some vertex gadgets. Let T'_1 be the triangulation obtained from T_1 by unlocking some vertex gadgets. Let T'_2 be the triangulation obtained from T_2 by unlocking the same vertex gadgets. Let C be the set of channels that have a locked vertex gadget at both ends. Then:

Property 5 *If the vertex gadgets at the ends of the channels of C remain locked, then the number of flips required to transform T'_1 to T'_2 is at least $28|E-C|+36|C|$.*

Proof. Consider a channel of C , with a locked vertex gadget at both ends. The cap vertices of the channel are not useable. By construction, the other vertices are outside the wide mouths of the channel. Therefore, by Property 4, we need 36 flips to transform it.

Consider the channels with an unlocked vertex gadget at one end. We only save flips by capping the channel. To do this, we must flip the two edges that block the channel from its cap. Because the edges that block one channel are disjoint from the edges that block any other channel, we must do two flips per channel, and we must re-flip those edges to return to the original state. Finally, by Property 2 it takes at least 24 flips to transform a capped channel. (Note that the proof of Property 2 carries over even if the channel is capped at both ends.) The total number of flips is 28 per channel. \square

2.3 Putting it all together

Lemma 2 *G has a vertex cover of size $\leq k$ if and only if the flip distance between the two triangulations T_1 and T_2 is $\leq 2k + 28|E|$.*

Proof. Suppose that G has a vertex cover of size k . Unlock the corresponding k vertex gadgets. Each edge channel has an unlocked gadget at one end, so by Property 3 we can transform between the two triangulations of the channel in 28 flips. When all channels have been transformed, we relock the k vertex gadgets. The total number of flips is $2k + 28|E|$.

For the other direction, suppose that there is a flip sequence between T_1 and T_2 of length $\leq 2k + 28|E|$. Let L be the set of vertices whose gadgets are unlocked in the flip sequence. Let C be the set of edges not covered by vertex set L . By adding one vertex to cover each edge of C , we observe that G has a vertex cover of size $|L| + |C|$. Thus it suffices to show that $|L| + |C| \leq k$. By Properties 4 and 5 the number of flips is at least $2|L| + 36|C| + 28(|E - C|) \geq 2|L| + 28|E| + 8|C|$.




Figure 3: Gadgets for vertices

By assumption, the number of flips was $\leq 2k + 28|E|$. Therefore $2|L| + 8|C| \leq 2k$, which implies that $|L| + |C| \leq k$, as required. \square

The last ingredient of the NP-completeness proof is to show that the reduction takes polynomial time. We need the following claim.

Claim 1 *The size of the coordinates used in the construction is bounded by a polynomial in n .*

Proof. We begin with a straight line drawing of a graph on a polynomial size grid. Expand the grid, and allocate a square region around each vertex for the vertex gadget

(Figure 4). Expand each edge to two parallel line segments. These line segments will become the channel, but for now, the reflex vertices of the channel are all collinear, which means that the channel's wide mouth is equal to its narrow mouth. The points C, D, E, F of the vertex gadget go in *feasible* regions defined by the wide and narrow mouths (e.g. in the 3-channel gadget, point D lies in the narrow mouth of channels 1 and 2, but outside the wide mouth of channel 3). We make the channels narrow enough so that all the feasible regions intersect the region allocated to the gadget.

To do this, note that the edges incident to the vertex corresponding to the gadget in the straight line drawing and their extensions (the dotted lines in Figure 4) intersect the square at points whose coordinates have polynomial size. Let S be the set of intersection points and corners of the square. For the edge corresponding to channel 1, consider the point p_1 where it intersects the square and find the point p other than itself in S that lies on the same edge of the square and is closest to it. Setting A to be the point on the boundary of the square a distance $pp_1/3$ away from p_1 towards p and B the symmetric point on the opposite side determines the channel and its width. Do the same thing at the other end of the edge corresponding to channel 1 and obtain another width. Finally, pick the narrower of the two options for channel 1. Since A and B lie on the edge of the square and their distance to p_1 is polynomial, we need polynomial number of bits to express the coordinates of A and B as well. Repeat the above for A', B', A'' and B'' . Since all the possible intersection points between the upper and lower chains of the channels occur inside the square, all the feasible regions have non-empty intersections with the interior of the square.

Now we pick points C, D, E, F inside the appropriate regions. Because the boundaries of the feasible regions are determined by pairs of points on the expanded grid, the new points can be chosen to have polynomial size (because solutions to linear programs have polynomial size as shown in Theorem 10.1 of [22]).

Finally we place the reflex points of each channel. The feasible region wherein each set of reflex points can be placed is bounded by lines through pairs of points already placed. Thus, we can choose reflex points of polynomial size. \square

3 Triangulations of point-sets

We prove the NP-hardness of computing the flip distance between two triangulations of a point set by reducing from computing the flip distance between two triangulations of a polygonal region. Given two triangulations T_1 and T_2 of a polygonal region R with n vertices, we triangulate all the holes and pockets of R the same way in both triangulations. Next, we repeat each




Figure 4: Constraints for vertex gadgets.

edge on the boundary of the holes and pockets n^2 times (as shown in Figure 5). This gives two triangulations T'_1 and T'_2 of a point set. Flip distance between the two triangulations will be the same as the flip distance between the original T_1 and T_2 . We use the fact that the flip distance to the constrained Delaunay triangulation is at most $\binom{n}{2}$ [2]. Thus the flip distance between T_1 and T_2 is less than n^2 , but dismantling a single set of repeated edges will itself require more flips.

We now describe the details about repeating the edges. Consider a triangulated hole or a pocket, say, $v_1v_2v_3v_4$. For each vertex on the boundary, draw the angle bisector of the angle formed by the two edges incident on that vertex. Thus at v_2 , we draw the angle bisector of $\angle v_1v_2v_3$. Next, we choose n^2 equally spaced points on a polynomially small portion of the bisector inside the hole. All of these points will have polynomial sized co-ordinates. Next, we repeat edges as shown in Figure 5.

Note that flipping a few edges to one of the new vertices will not help because the new vertices behave like the vertex whose angle bisector they were drawn on. This intuition is captured in the following lemma, which then implies the NP-completeness:

Lemma 3 *Flip distance between T_1 and T_2 is equal to the flip distance between T'_1 and T'_2 .*

Proof. For each vertex v_i on the boundary of a hole, there is a set of n^2 points associated with it. Call the union of the set together with v_i itself the *cluster corresponding to v_i* . Define two triangulations to lie in the same *equivalence class* if they are the same when we collapse each cluster into one point.




Figure 5: Repeating edges on the boundary of pockets and holes.

Flip distance between T'_1 and T'_2 is at most the flip distance between T_1 and T_2 because imitating a flip sequence that transforms T_1 to T_2 gives a flip sequence that transforms T'_1 to T'_2 .

But flip distance between T_1 and T_2 is at most the flip distance between T'_1 and T'_2 as well. Consider the smallest flip sequence that transforms T'_1 to T'_2 . It has a length of at most n^2 . At each step, consider the triangulation obtained by collapsing each cluster into one point. If this triangulation does not change, ignore the step. Thus flipping an edge that uses only “old” vertices to one that uses some “new” ones will be ignored. If the change happens only inside a hole or a pocket, ignore that step as well. Since the number of flips is at most n^2 , the hole is not broken during the flip sequence. Thus the resulting sequence transforms T_1 to T_2 . \square

Theorem 4 *The following problem is NP-complete: Given two triangulations of a point set in the plane, and a number k , is the flip distance between the triangulations at most k ?*

4 Conclusion

We have shown that it is NP-complete to compute the flip distance for triangulations of a polygonal region, or a point set. The problem remains open for a convex polygon, or a simple polygon, and also for more combinatorial objects such as labelled and unlabelled maximal planar graphs.

Acknowledgements. We thank Therese Biedl for helpful suggestions.

References

- [1] J.-L. Baril and J.-M. Pallo. Efficient lower and upper bounds of the diagonal-flip distance between triangulations. *Information Processing Letters*, 100:131–136, 2006.

- [2] M. Bern and D. Eppstein. Mesh generation and optimal triangulation. In *Computing in Euclidean Geometry, 2nd edition*, pages 47–123. World Scientific, Lecture Notes Series on Computing, 1995.
- [3] T. Biedl, G. Kant, and M. Kauffman. On triangulating planar graphs under the four-connectivity constraint. *Algorithmica*, 19(4):427–446, 1997.
- [4] P. Bose and F. Hurtado. Flips in planar graphs. *Computational Geometry: Theory and Applications*, 42:60–80, 2009.
- [5] S. Cleary and K. St. John. Rotation distance is fixed-parameter tractable. *Inf. Process. Lett.*, 109(16):918–922, July 2009.
- [6] K. Culik II and D. Wood. A note on some tree similarity measures. *Information Processing Letters*, 15(1):39–42, 1982.
- [7] P. Dehornoy. On the rotation distance between binary trees. *Advances in Mathematics*, 223(4):1316–1355, 2010.
- [8] S. L. Devadoss and J. O’Rourke. *Discrete and Computational Geometry*. Princeton University Press, 2011.
- [9] N. Dyn, I. Goren, and S. Rippa. Transforming triangulations in polygonal domains. *Computer Aided Geometric Design*, 10:531–536, 1993.
- [10] D. Eppstein. Happy endings for flip graphs. *J. Computational Geometry*, 1, 2010.
- [11] S. Hanke, T. Ottmann, and S. Schuierer. The edge-flipping distance of triangulations. *Journal of Universal Computer Science*, 2(8):570–579, 1996.
- [12] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete and Computational Geometry*, 22:333–346, 1999.
- [13] C. Lawson. Software for C^1 surface interpolation. In J. Rice, editor, *Mathematical Software III*, pages 161–194. Academic Press, New York, 1977.
- [14] C. L. Lawson. Transforming triangulations. *Discrete Mathematics*, 3(4):365–372, 1972.
- [15] M. Li and L. Zhang. Better approximation of diagonal-flip transformation and rotation transformation. In *Proceedings of the 4th Annual International Conference on Computing and Combinatorics, COCOON ’98*, pages 85–94. Springer-Verlag, 1998.
- [16] J. M. Lucas. Untangling binary trees via rotations. *The Computer Journal*, 47(2), 2004.
- [17] F. Luccio, A. M. Enriquez, and L. Pagli. Lower bounds on the rotation distance of binary trees. *Information Processing Letters*, 110(21):934–938, 2010.
- [18] E. Osherovich and A. Bruckstein. All triangulations are reachable via sequences of edge-flips: an elementary proof. *Computer Aided Geometric Design*, 25:157–161, 2008.
- [19] J. Pallo. An efficient upper bound of the rotation distance of binary trees. *Information Processing Letters*, 73:87–92, 2000.
- [20] A. Pilz. Flip distance between triangulations of a planar point set is NP-complete. <http://arxiv.org/abs/1206.3179>.
- [21] G. Rote and I. Bárány. Strictly convex drawings of planar graphs. *Documenta Mathematica*, 11:369–391, 2006.
- [22] A. Shrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, 1986.
- [23] D. D. Sleator, R. E. Tarjan, and W. P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.*, 1:647–681, 1988.
- [24] T. Watanabe, T. Ae, and A. Nakamura. On the node cover problem of planar graphs. In *Proceedings of the International Symposium on Circuits and Systems*, pages 78–81, 1979.