A combinatorial method to find sharp lower bounds on flip distances

Lionel Pournin[1,2,3]

[1] *EFREI, 30-32 avenue de la République, 94800 Villejuif, France*
[2] *LIAFA, Université Paris Diderot - Paris 7, Case 7014, 75205 Paris Cedex 13, France*
[3] *EPFL - École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland*

Abstract. Consider the triangulations of a convex polygon with n vertices. In 1988, Daniel Sleator, Robert Tarjan, and William Thurston have shown that the flip distance of two such triangulations is at most $2n-10$ when n is greater than 12 and that this bound is sharp when n is large enough. They also conjecture that "large enough" means greater than 12. A proof of this conjecture was recently announced by the author. A sketch of this proof is given here, with emphasis on the intuitions underlying the construction of lower bounds on the flip distance of two triangulations.

Résumé. En 1988, Daniel Sleator, Robert Tarjan et William Thurston ont montré que la distance, en nombre de flips, de deux triangulations d'un polygone convexe de n sommets est au plus $2n - 10$ quand n est supérieur à 12. Ils ont également montré que cette borne est atteinte si n est suffisamment grand et ils conjecturent qu'il existe deux triangulations à distance $2n - 10$ dès que n est supérieur à 12. Un preuve de cette conjecture a récemment été proposée par l'auteur. Une ébauche de cette preuve est présentée ici qui explique de manière intuitive les méthodes permettant d'obtenir des bornes inférieures sur la distance, en nombre de flips, de deux triangulations.

Keywords: Triangulations, Flip-graph, Rotation distance, Binary tree, Associahedra

1 Introduction

Consider a convex polygon π with n vertices. For convenience, π will be identified with the set of its vertices. An edge of π is a subset of π with two elements. A triangulation of π is a maximal set of pairwise non-crossing edges of π. The elements of a triangulation will be referred to as *its* edges.

According to this definition, a triangulation T of π contains every boundary edge of π. All the other edges of T will be called its *interior edges*. Note that an interior edge ε of T is the diagonal of a quadrilateral q whose four boundary edges belong to T. Replacing ε within T by the other diagonal, say ς, of quadrilateral q results in another triangulation T/ε of π:

$$T/\varepsilon = [T\backslash\{\varepsilon\}] \cup \{\varsigma\}.$$

This operation is called a *flip*. While working on the dynamic optimality conjecture [7], Daniel Sleator, Robert Tarjan, and William Thurston have shown that, when n is greater than 12, one can transform any triangulation of π into any other triangulation of π by performing at most $2n - 10$ flips [8]. In other words, the flip distance of two triangulations of π is at most $2n-10$ when n is greater than 12. It is also proven in [8] that when n is "large enough", there exist pairs of triangulations of π whose flip distance is precisely $2n-10$. Unfortunately their proof does not give an hint on how large n should be, and they propose the following conjecture: *there exist pairs of triangulations of π whose flip distance is precisely $2n-10$ whenever n is greater than* 12.

While the upper bound obtained by Daniel Sleator, Robert Tarjan, and William Thurston follows from an easy combinatorial argument, they use constructions in hyperbolic space to prove the existence of pairs of triangulations at maximal flip distance. They comment that the role played by hyperbolic geometry in this problem, whose statement is purely combinatorial, may seem mysterious and they ask for a combinatorial proof of their existence result. Three years ago, Patrick Dehornoy made progress toward such a proof by finding a lower bound of the form $2d-O(\sqrt{d})$ on the diameter of the d-dimensional associahedron, using combinatorial arguments [1].

A solution to these two open problems has been announced recently [6]. The purpose of this extended abstract is to sketch this proof with emphasis on the underlying intuitions.

Two triangulations at maximal flip distance are described in Section 3. The proof that these triangulations are indeed maximally distant is sketched in Section 4, and the combinatorial methods used to do so (i.e. general equalities and inequalities on flip distances) are presented in Section 2. In some places, proofs are omitted. In these cases, the interested reader is referred to [6] for the complete argument.

2 Equalities and inequalities on flip distances

In this section, several types of equalities and inequalities on flip-distance are obtained. Consider two triangulations U and V of a same polygon. A *path* of length k between U and V is a sequence of k flips that transforms U into V. A such path is called *minimal* when its length is minimal among all the path between U and V. Hereafter, the length of any minimal path between U and V is denoted by $\delta(\{U,V\})$. Part (a) of Lemma 3 from [8] states that, if an edge of V can be introduced in U by some flip, then there is a minimal path from U to V that begins with this flip. This result is generalized by the following theorem, whose straightforward proof can be found in [6]:

Theorem 1 *Let T, U, and V be three triangulations of a convex polygon. If:*

i. T is found along a minimal path between U and V,

ii. An edge of T can be introduced in U by some flip,

then there is a minimal path from U to V that begin with this flip.

This theorem makes it possible to prescribe the first flip along some minimal path between two triangulations. It can also be thought of as an equality on flip distances. Indeed, if ε is the edge removed by the flip mentioned in the second statement of Theorem 1, then the conclusion of this theorem is:

$$\delta(\{U,V\}) = \delta(\{U/\varepsilon, V\}) + 1.$$

Theorem 1 can further be generalized to sequences flips [6]. This generalization will be replaced here by less complete as in the proof of Theorem 6 thereafter. The proof sketched in Section 4 will also require inequalities that compare flip distances in the case of two polygons with different numbers of vertices. In order to obtain such inequalities, one can use the contraction operation that is now introduced.

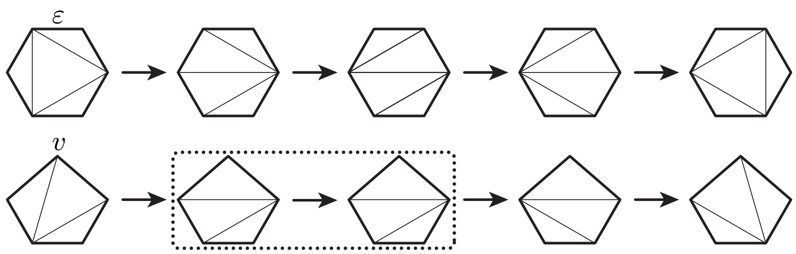

Fig. 1: A (minimal) path between the two triangulations of the hexagon whose interior edges form a triangle (top row), and the sequence of triangulations resulting from the contraction of edge ε to vertex v (bottom row). The framed triangulations in the bottom row are identical, and one of them can be removed from the sequence.

Let ε be some boundary edge of a given polygon π. Contracting ε in a triangulation T of π consists in replacing the two vertices of ε by a single point v within every edge of $T\backslash\{\varepsilon\}$. If v belongs to $\mathrm{conv}(\varepsilon)$, then this operation results in a triangulation of $(\pi\backslash\varepsilon)\cup\{v\}$ (a proof of this is given in [6]). It is assumed in the following that v inherits the labels of the two vertices of ε, and can be referred to indifferently using these two labels. This convention on the labeling of the contracted polygon is slightly different from that in [6]. Denote by $T\curlyvee\varepsilon$ the triangulation obtained by contracting edge ε in a triangulation T of π.

The effect of contractions on a path between two triangulations is now described using an example. Denote by U and V the two triangulations of the hexagon whose interior edges form a triangle. A path between U and V is shown in the top of Figure 1. Contracting the boundary edge ε at the top of the hexagon (as shown in the figure) in every triangulation along this path results in the sequence of triangulations of the pentagon depicted in the bottom of the figure. It can be seen that two consecutive triangulations in this sequence are either identical or connected by a flip. Now observe that the second and the third triangulation in the second row are identical precisely because the two corresponding triangulations in the top row are connected by a flip that modifies the triangle containing edge ε. In other words, the length of the path between $U\curlyvee\varepsilon$ and $V\curlyvee\varepsilon$ resulting from the contraction is less than the length of the path between U and V by the number of flips along this path that modify the triangle containing ε. According to the following theorem, proven in [6], this property holds in general:

Theorem 2 *Let U and V be two triangulations of a same polygon and ε a boundary edge of this polygon. If ψ is a path of length k between U and V, then there exists a path of length $k-j$ between $U\curlyvee\varepsilon$ and $V\curlyvee\varepsilon$, where j is the number of flips along path ψ that modify the triangle containing edge ε.*

Consider a pair P of triangulations of a polygon π and a boundary edge ε of π. Theorem 2 provides a convenient way to obtain inequalities between the flip distances of P and of $P\curlyvee\varepsilon$. Call $\vartheta(P,\varepsilon)$ the maximal number of flips that modify the triangle containing ε along any minimal path between the two elements of P. The following corollary is a direct consequence of Theorem 2:

Corollary 1 *Let P be a pair of triangulations of a polygon π. If ε is a boundary edge of π then:*

$$\delta(P)\geq\delta(P\curlyvee\varepsilon)+\vartheta(P,\varepsilon).$$

This corollary provides lower bounds on $\delta(P)$ that depend on $\vartheta(P,\varepsilon)$. The remainder of the section is dedicated to finding a lower bound on $\vartheta(P,\varepsilon)$.

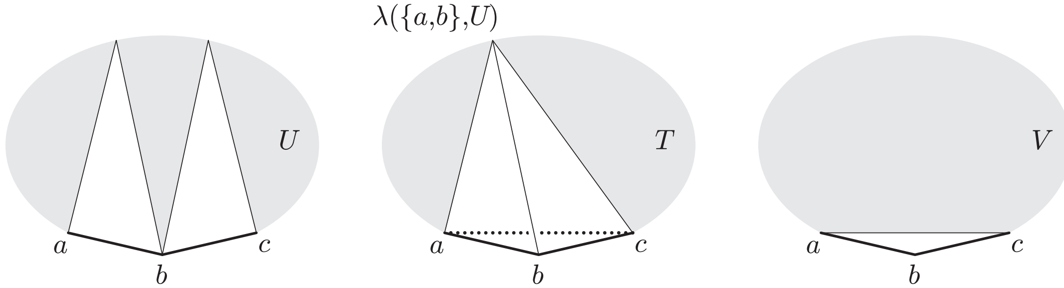

Fig. 2: Triangulations U (left) and V (right) depicted according to the requirements of Theorem 3, and triangulation T (center) used in the proof of this theorem. The dotted line shows that edge $\{b\} \cup \lambda(\{a,b\},U)$ is replaced by edge $\{a,c\}$ in triangulation T by some flip along a minimal path from U to V.

Let T be a triangulation of a polygon π and ε a boundary edge of π. Call a and b the two vertices of ε. Denote by c the vertex of π so that $\{a,c\}$ and $\{b,c\}$ are two edges of T. In other words, a, b, and c are the three vertices of the triangle in T that contains edge ε. The *link* of ε in T, denoted by $\lambda(\varepsilon,T)$ hereafter, is the set $\{c\}$. The following theorem is borrowed from [6]:

Theorem 3 *Let U and V be two triangulations of a polygon π. If a, b, and c three vertices of π so that:*

i. $\{a,b\}$ and $\{b,c\}$ are boundary edges of π and $\{a,c\}$ belongs to V,

ii. $\lambda(\{a,b\},U)$ and $\lambda(\{b,c\},U)$ are distinct subsets of $\pi\backslash\{a,c\}$,

then $\vartheta(\{U,V\},\{a,b\})$ and $\vartheta(\{U,V\},\{b,c\})$ are not both less than 2.

Proof: Consider three vertices a, b, and c of π. Assume that $\{a,b\}$ and $\{b,c\}$ are boundary edges of π and that $\{a,c\}$ is an edge of V. As a consequence, triangle $\{a,b,c\}$ is found in triangulation V. Further assume that $\lambda(\{a,b\},U)$ and $\lambda(\{b,c\},U)$ are distinct subsets of $\pi\backslash\{a,c\}$. It follows that the two edges $\{a,b\}$ and $\{b,c\}$ are contained in two distinct triangles of U whose unique common vertex is b. Triangulations U and V are depicted in the left and in the right of Figure 2 in this case.

Now assume that $\vartheta(\{U,V\},\{a,b\}) \leq 1$. Since the triangles in U containing $\{a,b\}$ and $\{b,c\}$ are distinct, edge $\{a,b\}$ cannot be contained in the same triangle in U and in V. Hence, there is exactly one flip that modifies the triangle containing $\{a,b\}$ along any minimal path from U to V. Consider such a path ψ, and call T the triangulation in which this flip is performed along path ψ. Since there is only one such flip along path ψ, this flip necessarily replaces edge $\{b\} \cup \lambda(\{a,b\},U)$ by edge $\{a,c\}$ as shown in the center of figure 2, where the latter edge is depicted as a dotted line. As $\lambda(\{a,b\},U)$ and $\lambda(\{b,c\},U)$ are distinct, then the triangle containing edge $\{b,c\}$ must have been modified at least once along path ψ before T is reached. Moreover, the flip performed in T also modifies the triangle containing edge $\{b,c\}$. As a consequence, the triangle containing edge $\{b,c\}$ is modified at least twice along path ψ, which shows that $\vartheta(\{U,V\},\{b,c\})$ is not less than 2. $\square$

Theorem 1, Corollary 1 and Theorem 3 are the basic building blocks of the combinatorial method used in Section 4 to obtain sharp lower bounds of flip distances. While they need to be generalized to allow for a rigorous proof of the desired result (see [6]), such generalizations will not be given here. Intuition on how the proof works will be provided instead using the results presented above.

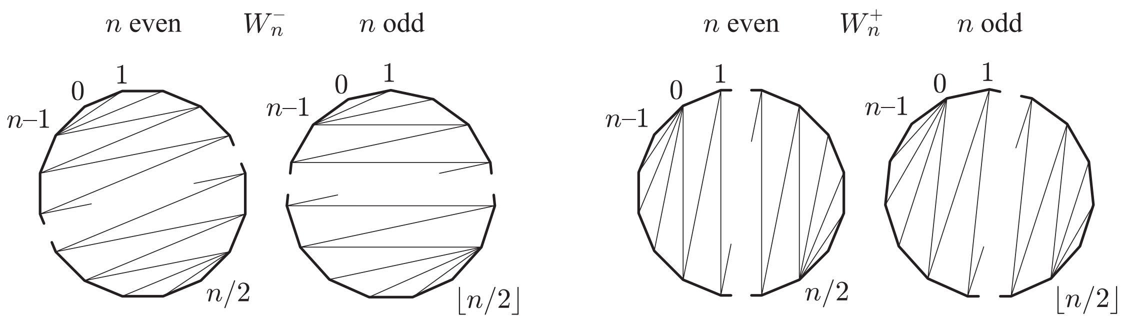

Fig. 3: Triangulations W_n^- (left) and W_n^+ (right) depicted when n is greater than 8.

3 A pair of triangulations at maximal distance

Let π be a convex polygon with n vertices labeled clockwise from 0 to $n-1$. Consider the triangulations W_n^- and W_n^+ of π depicted in Figure 3 depending on the parity of n. It is shown in [6] that these two triangulations have flip distance $2n-10$ when n is greater than 12. As can be seen in the figure, W_n^- contains three interior edges incident to vertex $n-1$. A such set of edges will be referred to as a *comb with three teeth at vertex* $n-1$. Triangulation W_n^- has another comb at vertex $\lfloor n/2 \rfloor - 1$ with three or four teeth depending on the parity of n. Triangulation W_n^+ contains a comb with four teeth at vertex 0, and another comb at vertex $\lfloor n/2 \rfloor$ with three or four teeth depending on the parity of n. The remaining interior edges of W_n^- and W_n^+ form a zigzag (i.e. a simple path whose vertices of edges belong to "opposite" sides of the polygon) that connects the two combs contained in these triangulations. One can see in Figure 3 that the size of these zigzags depends on n. When n is equal to 9, the polygon is precisely large enough to contain the two combs. In this case, the combs have a common tooth, and the zigzag disappears. For this reason, the above description of W_n^- and W_n^+, and their representation in Figure 3 are only valid when n is greater than 8.

The definition of W_n^- and W_n^+ given in [6] is more general. In particular, it is valid whenever n is greater than 2. Obviously, when $3 \leq n \leq 8$, the combs found in these triangulations lose teeth, or even disappear. For these small values of n, though, W_n^- and W_n^+ keep a number of important properties that are needed in the proof. Triangulations W_n^- and W_n^+ are sketched in Figure 4 when $4 \leq n \leq 8$. If n is equal to 3, both triangulations shrink to a single triangle, and for this reason, they are omitted in the figure. For any n greater than 2, denote:

$$A_n = \{W_n^-, W_n^+\}.$$

When $3 \leq n \leq 12$, the flip distance of pair A_n can be easily found using a computer program or a mathematical proof. In particular, the flip distances reported in the following table are obtained in [6] using the methods presented in Section 2, and their generalizations:

n	3	4	5	6	7	8	9	10	11	12
$\delta(A_n)$	0	1	2	4	5	7	8	10	12	14

One can see in this table that $\delta(A_n)$ is greater than $2n-10$ when $3 \leq n \leq 8$ and equal to this value when $9 \leq n \leq 12$. Using the results of the previous section, a recursive lower bound on $\delta(A_n)$ will be

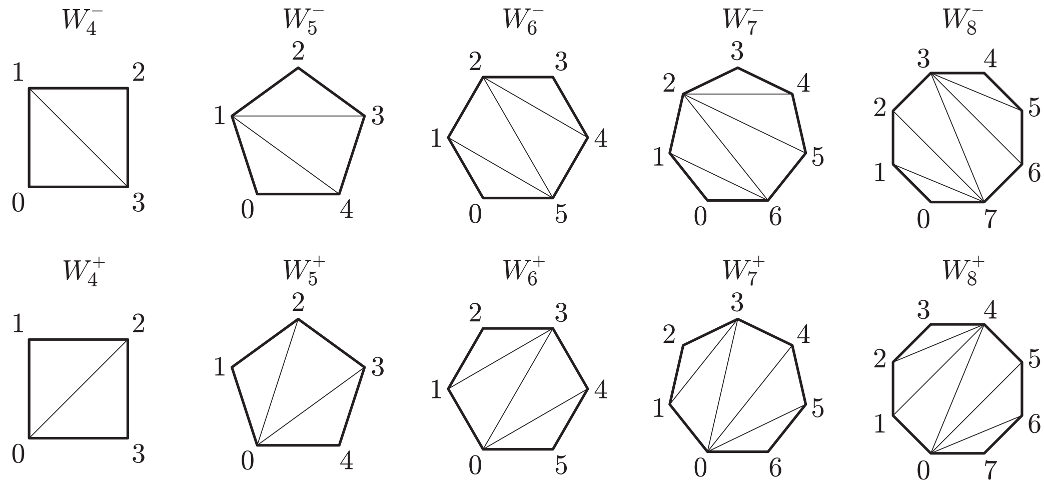

Fig. 4: Triangulations W_n^- (left) and W_n^+ (right) depicted when $4 \leq n \leq 8$.

obtained in the next section for any $n > 12$. This lower bound together with the values of $\delta(A_n)$ reported in the above table will produce the desired result.

4 Sketch of the proof

It can be seen in Figures 3 and 4 that contracting edge $\{n-2, n-1\}$ in triangulations W_n^- and W_n^+ results in pair A_{n-1}, up to a relabeling of the vertices. Note that this observation holds whenever n is greater than 3. Hence, one obtains the following result as a consequence of Corollary 1:

Theorem 4 *Let n be an integer greater than 3, if $\vartheta(A_n, \{n-2, n-1\}) \geq 2$,*

$$\delta(A_n) \geq \delta(A_{n-1}) + 2.$$

Assuming that the inequality $\delta(A_n) \geq \delta(A_{n-1}) + 2$ holds in general for any n greater than 3, the desired result immediately follows. Unfortunately, this inequality is only obtained under a rather strong condition, that is the existence of a minimal path between W_n^- and W_n^+ that modifies at least twice the triangle containing edge $\{n-2, n-1\}$. In fact, such a minimal path does not necessarily exist. In order to overcome this difficulty, other contractions will be considered.

Observe, for instance, that the contraction considered in Theorem 4 can be iterated. If n is greater than 4, contracting edge $\{n-2, n-1\}$ in A_n and then edge $\{0, 1\}$ will result in pair A_{n-2}, up to a relabeling of the vertices. This observation remains true if one exchanges the order of the two contractions: first contracting edge $\{0, 1\}$ and then edge $\{n-2, n-1\}$ in A_n still results in a pair of triangulations whose flip-distance is $\delta(A_{n-2})$. It can be seen in Figures 3 and 4 that the triangle containing edge $\{n-2, n-1\}$ is not the same in triangulations $W_n^- \curlyvee \{0, 1\}$ and $W_n^+ \curlyvee \{0, 1\}$. Note that this observation holds whenever n is greater than 5. Hence, at least one flip modifies this triangle in any minimal path between these two triangulations and, as a consequence, Corollary 1 yields:

$$\delta(A_n \curlyvee \{0, 1\}) \geq \delta(A_{n-2}) + 1.$$

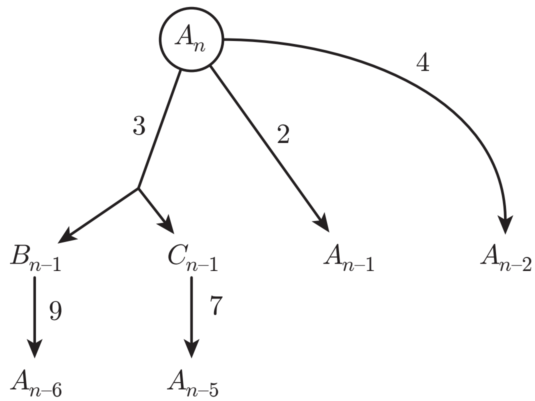

Fig. 5: Sketch of the proof of Theorem 9. An arc of weight w from a pair P to a pair Q represents the inequality $\delta(P) \geq \delta(Q) + w$ obtained under some condition. These conditions are omitted in this sketch but can be found in the statements of the corresponding theorems.

Combining this inequality with the inequality obtained invoking Corollary 1 for the contraction of $\{0, 1\}$ in pair A_n results in the following theorem:

Theorem 5 *Let n be an integer greater than 5, if $\vartheta(A_n, \{0, 1\}) \geq 3$,*

$$\delta(A_n) \geq \delta(A_{n-2}) + 4.$$

As was the case with Theorem 4, the inequality provided by Theorem 5 is subject to an (even stronger) condition involving the existence of a particular path between W_n^- and W_n^+. Again, nothing tells us that such a path exists. For this reason, the condition complementary to the requirements of Theorems 4 and 5 will have to be investigated. Before going any further, consider Figure 5, where the main proof is sketched using a tree. In this figure, each arc corresponds to an inequality obtained under some condition. The inequality corresponding to an arc of weight w from a pair P to a pair Q is:

$$\delta(P) \geq \delta(Q) + w.$$

Hence, the arcs from A_n to A_{n-1} and from A_n to A_{n-2} respectively correspond to Theorems 4 and 5. Since these two inequalities are obtained under the condition that particular paths exist between W_n^- and W_n^+, the third arc originating at pair A_n in Figure 4 will be considered under the complementary condition. Observe that this arc has two ends. This is due to the fact that the inequality corresponding to this arc relates $\delta(A_n)$ with either $\delta(B_{n-1})$ or $\delta(C_{n-1})$. The two pairs of triangulations B_n and C_n can be defined for any integer n greater than 8 as:

$$B_n = \{Y_n^-, Y_n^+\} \text{ and } C_n = \{X_n^-, Y_n^+\},$$

where Y_n^-, Y_n^+, and X_n^- are depicted in Figure 6. Triangulation Y_n^- is depicted using solid lines in the left of the figure. It contains two combs at vertices 3 and $\lceil n/2 \rceil - 1$. The comb at vertex 3 always has three

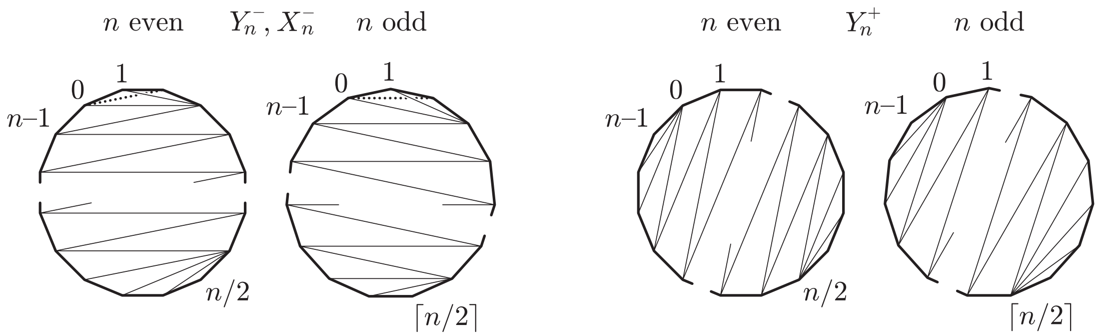

Fig. 6: Triangulation Y_n^- (left) and Y_n^+ (right) shown in solid lines when n is greater than 8. Triangulation X_n^-, obtained by flipping edge $\{1,3\}$ in Y_n^- is depicted using dotted lines (left).

teeth, and the other comb has four teeth if n is even, and only three is n is odd. The remaining interior edges of the triangulation form a zigzag that connects the two combs. Note that when n is equal to 9 or to 10, the zigzag disappears as the two combs become adjacent. Triangulation X_n^- is obtained by flipping edge $\{1,3\}$ in triangulation Y_n^-, which is sketched using dotted lines in the left of Figure 6. Observe that this flip removes the comb at vertex 3. Triangulation Y_n^+ is shown in the right side of the figure. As can be seen, Y_n^+ can be built by appropriately relabeling the vertices of W_n^-. In particular Y_n^+ contains a comb with three teeth at vertex 0 and another comb at vertex $\lceil n/2 \rceil$ whose number of teeth (three or four) depends on the parity of n.

Observe that pairs B_n and C_n are also defined when $7 \leq n \leq 8$ in [6]. This is needed for the computer-free proof that $\delta(A_n) \geq 2n-10$ when $3 \leq n \leq 12$. This result has already been given above without a proof (see [6] for one). It is therefore not necessary to define B_n and C_n when $7 \leq n \leq 8$ here. For the same reason, the following theorem is stated here when n is greater than 9 rather than when n is greater than 7. It gives the inequality corresponding to the leftmost arc originating at pair A_n in Figure 5:

Theorem 6 *Let n be an integer greater than 9. If $\vartheta(A_n,\{n-2,n-1\}) \leq 1$ and $\vartheta(A_n,\{0,1\}) \leq 2$, then there exists $P \in \{B_{n-1}, C_{n-1}\}$ so that $\delta(A_n) = \delta(P)+3$ and $\vartheta(P,\{0,1\}) \leq 1$.*

Proof (sketch): Assume that $\vartheta(A_n,\{n-2,n-1\}) \leq 1$ and that $\vartheta(A_n,\{0,1\}) \leq 2$. Observe that $\{n-2,n-1\}$ is not contained in the same triangle in W_n^- and W_n^+. Hence, there is exactly one flip that modifies the triangle containing this edge along any minimal path from W_n^- to W_n^+. Consider such a path ψ, and call T the triangulation resulting from the flip along this path that modifies the triangle containing $\{n-2,n-1\}$. Since there is only one such flip along path ψ, this flip necessarily replaces edge $\{3,n-1\}$ by edge $\{0,n-2\}$. In particular, T contains edges $\{0,n-2\}$ and $\{0,3\}$. As a first consequence, all the boundary edges of quadrilateral $\{0,1,2,3\}$ are contained in T and one diagonal of this quadrilateral (i.e. $\{0,2\}$ or $\{1,3\}$) must also be contained in T. Denote by ε this diagonal.

It can be seen in Figure 3 that edges $\{0,2\}$, $\{0,3\}$, and $\{0,n-2\}$ can be introduced in W_n^- by successively flipping $\{1,n-1\}$, $\{2,n-1\}$, and $\{3,n-1\}$ in this order. Reversing the order of the first two flips, this sequence will introduce edges $\{1,3\}$, $\{0,3\}$, and $\{0,n-2\}$ instead. This proves that edges ε, $\{0,3\}$, and $\{0,n-2\}$ can be introduced in W_n^- by performing three consecutive flips. As shown above, these three edges are contained in T. Hence, it follows from Theorem 1 that this sequence of three flips is

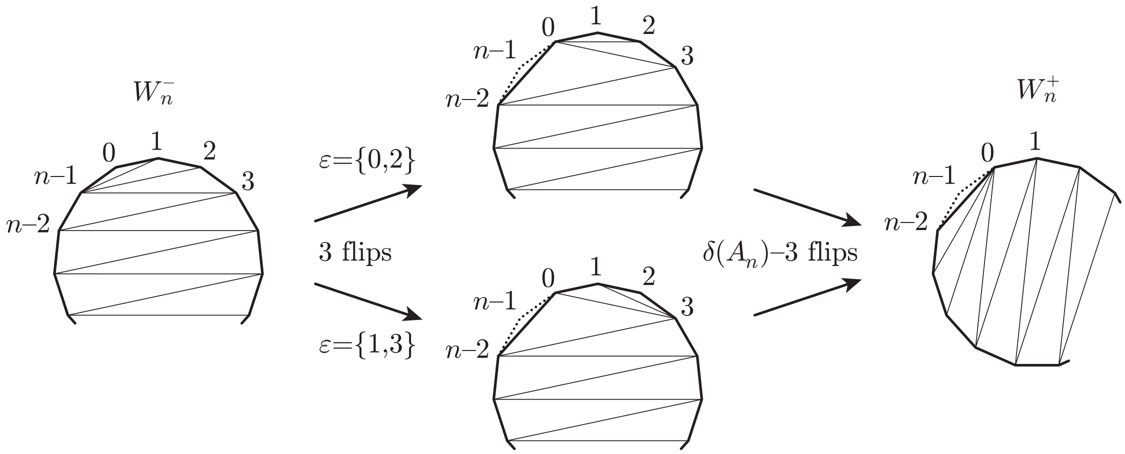

Fig. 7: Triangulation U (center) used in the proof of Theorem 6, and shown here depending on ε. In this proof, triangulation U is reached after three flips along a path ϕ from W_n^- to W_n^+.

the beginning of some minimal path ϕ from W_n^- to W_n^+.

The triangulation U obtained after the three first flips along path ϕ is shown in Figure 7 depending on whether ε is equal to $\{0,2\}$ or to $\{1,3\}$. One can see that U and W_n^+ both contain the triangle with vertices 0, $n-1$, and $n-2$. Hence, it follows from part (b) of Lemma 3 from [8] that this triangle will be contained in every triangulation found along any minimal path between U and W_n^+. As a consequence, one can remove the two boundary edges of π incident to vertex $n-1$ from U and from W_n^+ without changing the flip distance. The pair of triangulations thus obtained is either equal to B_{n-1} (if $\varepsilon=\{1,3\}$) or to C_{n-1} (if $\varepsilon=\{0,2\}$). Observe in particular that relabeling the vertices of π is unnecessary here.

Denote $P=B_{n-1}$ if $\varepsilon=\{1,3\}$ and $P=C_{n-1}$ if $\varepsilon=\{0,2\}$. It has been proven that:

$$\delta(A_n)=\delta(\{U,W_n^+\})+3 \text{ and } \delta(\{U,W_n^+\})=\delta(P).$$

Combining these two equalities yields $\delta(A_n)=\delta(P)+3$. Now observe that exactly one of the first three flips along path ϕ modifies the triangle that contains in $\{0,1\}$. Since a minimal path between W_n^- and W_n^+ can be built from the three first flips along path ϕ and from any minimal path between the two elements of P, one obtains:

$$\vartheta(A_n,\{0,1\})\geq\vartheta(P,\{0,1\})+1.$$

As $\vartheta(A_n,\{0,1\})\leq 2$, this proves that $\vartheta(P,\{0,1\})$ is not greater than 1. □

Two more inequalities have to be obtained, corresponding to the arcs originating at pairs B_{n-1} and C_{n-1} in Figure 5. According to the statement of Theorem 6, these two inequalities can be proven under the assumption that the image by $\vartheta(P,\{0,1\})$ is at most 1, where P is the pair at the origin of the arc. The proofs of these inequalities rely on Theorem 3 and its generalizations. Here, only the simplest of these inequalities will be proven. Moreover, instead of using the generalization of Theorem 3 found in [6], the proof will be sketched in an alternative, more intuitive way.

Theorem 7 *Let n be an integer greater than 11. If $\vartheta(C_n,\{0,1\})\leq 1$, then $\delta(C_n)\geq\delta(A_{n-4})+7$.*

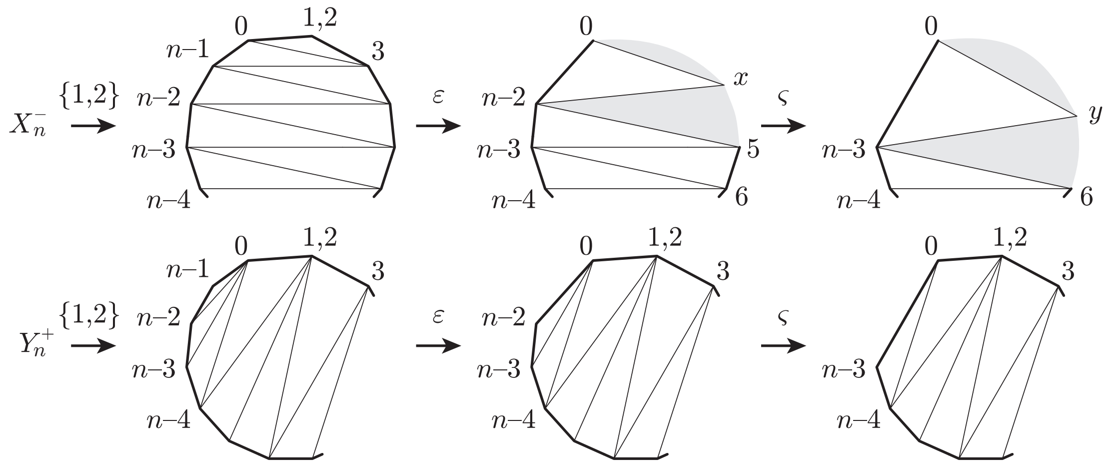

Fig. 8: The first three contractions used in the proof of Theorem 7. The arrows indicate the successive contractions of edges $\{1,2\}$, ε, and ς in triangulations X_n^- and Y_n^+. The vertex resulting from the contraction of edge $\{1,2\}$ is labeled by 1 and by 2 following the convention introduced in Section 2. The vertex labeled x is either equal to 3 or to 4, and the vertex labeled y is either equal to 3, to 4, or to 5 depending on ε and ς.

Proof (sketch): Assume that $\vartheta(C_n,\{0,1\}) \leq 1$. First observe that successively contracting edges $\{1,2\}$, $\{n-3,n-2\}$, $\{n-2,n-1\}$, and $\{0,n-1\}$ in pair C_n results in pair A_{n-4} up to a renumbering of the vertices. It can be seen in Figure 6 that, in pair C_n, vertices 0, 1, and 2 satisfy the conditions on a, b, and c required by Theorem 3. As $\vartheta(C_n,\{0,1\}) \leq 1$, then according to this theorem,

$$\vartheta(C_n,\{1,2\}) \geq 2. \tag{1}$$

The two elements of pair $C_n \curlyvee \{1,2\}$ are depicted in the left of Figure 8. It can be seen in this figure that the links of edges $\{n-3,n-2\}$, $\{n-2,n-1\}$, and $\{0,n-1\}$ in triangulation $X_n^- \curlyvee \{1,2\}$ are respectively $\{5\}$, $\{4\}$, and $\{3\}$. This is a consequence of n being greater than 11. Indeed, recall that X_n^- has a comb at vertex $\lceil n/2 \rceil - 1$ (see Figure 6). Since n is greater than 11 then $5 \leq \lceil n/2 \rceil - 1$ and, among the three edges $\{n-3,n-2\}$, $\{n-2,n-1\}$, and $\{0,n-1\}$, only the first one is possibly placed between two teeth of the comb at vertex $\lceil n/2 \rceil - 1$. Hence, the three links are necessarily distinct.

Now observe that contracting edges $\{n-3,n-2\}$, $\{n-2,n-1\}$, and $\{0,n-1\}$ in any order in this pair always result in A_{n-4} (recall that when an edge is contracted to a vertex, this vertex inherits the labels of the two vertices of the contracted edge). It can be proven, using Theorem 3, that one of these orders provides the desired inequality. Indeed, it can be seen in the left of Figure 8 that, in pair $C_n \curlyvee \{1,2\}$, vertices 0, $n-1$, and $n-2$ satisfy the conditions on a, b, and c in the statement of this theorem. Hence, the following inequality holds, either with $\varepsilon = \{0,n-1\}$ or with $\varepsilon = \{n-2,n-1\}$:

$$\vartheta(C_n \curlyvee \{1,2\}, \varepsilon) \geq 2. \tag{2}$$

The two elements of pair $C_n \curlyvee \{1,2\} \curlyvee \varepsilon$ are depicted in the center of Figure 8. In this figure, x is either equal to 3 (if $\varepsilon = \{n-2,n-1\}$) or to 4 (if $\varepsilon = \{0,n-1\}$). It can be seen that, in pair $C_n \curlyvee \{1,2\} \curlyvee \varepsilon$, vertices 0, $n-2$, and $n-3$ satisfy the conditions on a, b, and c in the statement of Theorem 3. Hence, the following inequality holds, with $\varsigma = \{0, n-2\}$ or with $\varsigma = \{n-3, n-2\}$:

$$\vartheta(C_n \curlyvee \{1,2\} \curlyvee \varepsilon, \varsigma) \geq 2. \tag{3}$$

The two elements of pair $C_n \curlyvee \{1,2\} \curlyvee \varepsilon \curlyvee \varsigma$ are depicted in the right of Figure 8. Here, y is equal to 3, to 4, or to 5 depending on the values of ε and ς. Now recall that n is greater than 11. As a consequence, $y < n-4$. It can be seen in Figure 8 that, in this case, the triangle containing edge $\{0, n-3\}$ is not the same in the two elements of pair $C_n \curlyvee \{1,2\} \curlyvee \varepsilon \curlyvee \varsigma$. Hence, every minimal path between these two triangulations modifies at least once the triangle containing edge $\{0, n-3\}$, and therefore:

$$\vartheta(C_n \curlyvee \{1,2\} \curlyvee \varepsilon \curlyvee \varsigma, \{0, n-3\}) \geq 1. \tag{4}$$

Invoking four times Corollary 1 and using (1), (2), (3), and (4) yields the desired inequality. □

Observe that Theorem 7 can be proven using a generalization of Corollary 1, which avoids invoking this corollary four times (see [6]). The above alternative proof explains how this generalization works, and may help the interested reader in understanding the argument used in [6]. Note in particular, that the sequence of contractions used in the above proof is built so that every contraction "eats up" at least two flips, except for the last one.

Using a similar argument, one obtains the following result, that corresponds to the arc originating at pair B_{n-1} in Figure 5. The reader will find a rigorous proof of this result in [6]. Note that this proof requires five successive contractions instead of just four.

Theorem 8 *Let n be an integer greater than 11. If $\vartheta(B_n, \{0,1\}) \leq 1$ then $\delta(B_n) \geq \delta(A_{n-5}) + 9$.*

This theorem provides the last of the inequalities corresponding to the arcs shown in Figure 5. Recall that every inequality is obtained under some condition. The conditions associated to the different arcs originating at a given pair in Figure 5 exhaust all possibilites, though, and one can read this figure as the following recursive lower bound on $\delta(A_n)$:

Theorem 9 *For any integer n greater than 12,*

$$\delta(A_n) \geq \min(\delta(A_{n-1}) + 2, \delta(A_{n-2}) + 4, \delta(A_{n-5}) + 10, \delta(A_{n-6}) + 12).$$

This result can be used to prove inductively that $\delta(A_n) \geq 2n + O(1)$ for every integer n greater than 2. Now recall that, as mentioned in Section 3, the flip-distance of W_n^- and W_n^+ is at least $2n-10$ when $3 \leq n \leq 12$. Hence, one obtains that $\delta(A_n) \geq 2n + 10$ for every integer n greater than 2. Further recall that, as shown in [8] using a combinatorial argument, the flip distance of two such triangulations is at most $2n-10$ when n is larger than 12. One therefore obtains the following result:

Theorem 10 *For any integer n greater than 12,*

i. The flip distance of triangulations W_n^- and W_n^+ is exactly $2n-10$,

ii. The flip graph of a polygon with n vertices has diameter $2n-10$.

5 Conclusion

The method given in [6] to obtain sharp lower bounds on flip distances has been presented here with emphasis on the underlying intuitions, sometimes at the expense of completeness. The reader is referred to [6] for a complete and rigorous argument. This method provides a way to solve the two open problems formulated in [8]. In other words, it can be used to obtain a combinatorial proof that the flip-graph (i.e. the graph whose vertices are the triangulations, and whose edges are the flips) of a polygon with n vertices has diameter $2n - 10$ when n is greater than 12. A direct consequence of this result is that the diameter of the d-dimensional associahedron [4] is $2d - 4$ for all integers d greater than 9.

It is natural to ask whether this method, or a variation of it, could be applied to the several generalizations of triangulations and flips that can be found in the literature. One of these generalizations consists in considering the regular triangulations of a finite, but otherwise arbitrary d-dimensional set of points (see for instance [2] and [3]). In this case, though, the operation of edge contraction may not be possible [6]. Another way to generalize the problem is to consider multitriangulations and their flips [5]. It has been shown that edge contraction is possible in this case. Unfortunately, Lemma 3 from [8] may not carry over to multitriangulations, and one may need alternative arguments.

Finally, it was suggested by a referee of this paper that the methods presented above could also be used to obtain the maximal distance of centrally symmetric triangulations, under the modified flip operation that either exchanges two diameters of the polygon or simultaneously exchanges the diagonals of two centrally symmetric quadrilaterals. In this case, contractions should affect two opposite edges of the polygon. Solving this problem would, in addition, provide the diameters of cyclohedra.

References

[1] P. Dehornoy. On the rotation distance between binary trees. *Advances in Mathematics*, 223:1316–1355, 2010.

[2] I. M. Gel'fand, M. M. Kapranov, and A. V. Zelevinsky. Discriminants of polynomials of several variables and triangulations of Newton polyhedra. *Leningrad Mathematical Journal*, 2:449–505, 1990.

[3] I. M. Gel'fand, M. M. Kapranov, and A. V. Zelevinsky. *Discriminants, resultants, and multidimensional determinants*. Birkhäuser, 1994.

[4] C. W. Lee. The associahedron and triangulations of the n-gon. *European Journal of Combinatorics*, 10:551–560, 1989.

[5] V. Pilaud and F. Santos. Multitriangulations as complexes of star polygons. *Discrete & Computational Geometry*, 41:284–317, 2009.

[6] L. Pournin. The diameters of associahedra. arXiv:1207.6296, submitted, 2012.

[7] D. D. Sleator and R. E. Tarjan. Self-adjusting binary search trees. *Journal of the ACM*, 32:652–686, 1985.

[8] D. D. Sleator, R. E. Tarjan, and W. P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *Journal of the American Mathematical Society*, 1:647–681, 1988.