

ONCE PUNCTURED DISKS, NON-CONVEX POLYGONS, AND POINTIHEDRA

Hugo Parlier*, Lionel Pournin†

Abstract. We explore several families of flip-graphs, all related to polygons or punctured polygons. In particular, we consider the topological flip-graphs of once-punctured polygons which, in turn, contain all possible geometric flip-graphs of polygons with a marked point as embedded sub-graphs. Our main focus is on the geometric properties of these graphs and how they relate to one another. In particular, we show that the embeddings between them are strongly convex (or, said otherwise, totally geodesic). We also find bounds on the diameters of these graphs, sometimes using the strongly convex embeddings. Finally, we show how these graphs relate to different polytopes, namely type D associahedra and a family of secondary polytopes which we call pointihedra.

1. INTRODUCTION

Triangulations of surfaces are naturally linked to different areas of mathematics including combinatorics, graph theory, geometric topology and anything having to do with surface geometry. They serve as combinatorial models for geometric structures on surfaces, encode surface homeomorphisms and are related to the geometric positioning of points on two dimensional structures. Surface triangulations can be related to one another by *flip* transformations that, in their simplest manifestation, consist in exchanging two arcs (see the next section for a formal definition). Given a set of triangulations, we are interested in the geometry of the associated *flip-graph*. The vertices of this graph are the triangulations and its edges link two triangulations whenever they differ by a single flip.

A well studied example is the graph \mathcal{A}_n of the triangulations of a convex Euclidean n -gon: this graph turns out to be the 1-skeleton of a polytope, the associahedron, and its diameter is by now well understood [10, 11]. Flip-graphs of topological

*Research supported by Swiss National Science Foundation grants PP00P2_128557 and PP00P2_153024

†Research funded by Ville de Paris Émergences project “Combinatoire à Paris”.

Key words: flip-graphs, triangulations of surfaces, combinatorial moduli spaces

surfaces have also been studied: in this case, triangulations are considered up to isotopy, and the underlying flip-graphs are in general infinite graphs. They are related to the large scale geometry of the self-homeomorphism groups of the surface [2]. Somewhere in between the general topological case and the graph of the associahedron lie the triangulations of *filling surfaces* [8, 9]. These are orientable topological surfaces but, like in the case of a polygon, one varies the number of marked points along a privileged boundary curve. When the filling surface is a disk, this gives rise to the graph of the associahedron, because there is a one-to-one correspondence between the triangulations of a disk with n marked points on the boundary, considered up to isotopy, and the Euclidean geodesic triangulations of a fixed convex n -gon. The only other instance that gives rise to a finite flip-graph is when the filling surface is a punctured disk; this particular graph will be studied in the sequel and for future reference we call it \mathcal{T}_n . There is a first natural relationship between these two graphs: considering triangulations that contain a fixed arc from the puncture (which we consider as an interior vertex) to an outer vertex, it is not difficult to see that there are multiple copies of \mathcal{A}_{n+2} inside \mathcal{T}_n . Our first main result concerns the geometry of \mathcal{T}_n : we are able to identify its diameter exactly.

Theorem 1.1. *For every positive integer n , $\text{diam}(\mathcal{T}_n) = 2n - 2$.*

The methods we use to obtain the lower bound on this diameter are similar to those in [8] and [9], but the specific nature of the once-punctured disk provides a particularly nice illustration of them. It is interesting to note that the graph \mathcal{D}_n of the related type D associahedron admits \mathcal{T}_n as a proper subgraph. As shown in [1], the diameter of \mathcal{D}_n is also $2n - 2$. Although neither of these two diameter results imply the other, the similarity is striking. In particular, the diametrically opposite vertices of \mathcal{D}_n exhibited in [1] do not both belong to \mathcal{T}_n so we need to find new pairs of vertices to prove our lower bounds on the diameter of \mathcal{T}_n .

Another natural variant on \mathcal{A}_n is to relax the convexity condition of the underlying polygon: topologically nothing has changed but here we are looking for *geometric* triangulations of the polygon. This gives rise to natural connected subgraphs of \mathcal{A}_n . Note that their connectedness is not a priori obvious [4]. Similarly, one can take a Euclidean convex n -gon P and place a marked point (referred to as a puncture) in its interior asking again for the geometric triangulations of the resulting punctured polygon P^* . This time, the obtained flip-graph $\mathcal{F}(P^*)$ is a natural subgraph of \mathcal{T}_n . In

both graphs, there are copies of flip-graphs of polygons. In a multitude of different variants, we prove *convexity* results about how these subgraphs lie in the larger graphs. We show that certain embedded subgraphs are strongly convex by which we mean that a geodesic path between any two vertices of the subgraph is entirely contained in the subgraph. This allows us to find lower bounds on diameter of the different graphs. As an example of our results we show the following.

Theorem 1.2. *Let P be a convex n -gon. For any placement of the interior marked point within P , the resulting punctured polygon has a flip-graph of diameter at most $2n - 6$. In addition, one can place the interior marked point within P so that this punctured polygon has a flip-graph of diameter at least $2n - 8$.*

The last flip-graphs we consider are slightly different from the above graphs; they are obtained as the 1-skeleton of some secondary polytopes [3], just like \mathcal{A}_n can be obtained as the 1-skeleton of an associahedron. These graphs are closely related to the ones described above. Let P be a convex Euclidean n -gon and P^* the same polygon with an interior marked point. The graph we consider contains both the flip-graph of P (which is isomorphic to \mathcal{A}_n) and the flip-graph of P^* : in particular its vertex set is the union of the two vertex sets. In addition there are edges corresponding to extra flip moves between triangulations of P and triangulations of P^* that consist in inserting the puncture inside a triangulation (a proper definition is given in the next section). The resulting flip-graph, denoted by $\overline{\mathcal{F}}(P^*)$ is the graph of a polytope which we call a pointihedron. Our main results are, again, bounds on the diameter of this graph and convexity results on how \mathcal{A}_n and $\mathcal{F}(P^*)$ lie within it at least for certain placements of the puncture. For instance:

Theorem 1.3. *Let P be a convex n -gon. For certain placements of the puncture, the natural embedding $\mathcal{A}_n \hookrightarrow \overline{\mathcal{F}}(P^*)$ is strongly convex.*

We conjecture that this theorem remains true for any placement of the puncture. However, the methods we use to prove it do not seem to extendable to the general case. Finally, it is shown that, for some placement of the puncture, the natural embedding $\mathcal{F}(P^*) \hookrightarrow \overline{\mathcal{F}}(P^*)$ is not convex.

Acknowledgements. We are grateful for conversations with Delphine Milenovic who studied a particular variation of the geometric punctured flip-graph in her

master’s thesis. In particular she coined the term “pointihedra”. Part of this work was done while the first author was visiting the LIPN at the University of Paris XIII and the authors thank the department for its support.

2. PRELIMINARIES

We begin the section by defining the topological flip-graphs of a disk and a once punctured disk and discuss some of their geometric interpretations. We then move on to defining geometric flip-graphs, some of which are intrinsically geometric as, for instance, the flip-graphs of non-convex or punctured Euclidean polygons. The last flip-graph we describe in this section, related to the geometric ones, turns out to be the 1-skeleton of a polytope, the *pointihedron*. Along the way, we explain where the simplicial embeddings come from and we conclude the section with some general remarks on simplicial embeddings between flip-graphs.

2.1. Topological flip-graphs

Let Δ be the disk and Δ^* the punctured disk where the puncture is viewed as a marked point. If Σ is either of these surfaces, we denote by Σ_n the surface obtained from it by placing n marked points on the boundary curve. Two arcs between marked points of Σ_n are non-crossing when they have a common endpoint or disjoint interiors. A triangulation of Σ_n is a maximal set of such non-crossing arcs. We denote by $\mathcal{F}(\Sigma_n)$ the flip-graph of Σ_n . The vertices of this graph are the isotopy classes of the triangulations of Σ_n and its edges connect triangulations that differ by a single arc. Equivalently, two triangulations are connected by an edge of $\mathcal{F}(\Sigma_n)$ when they can be obtained from one another by a single *flip* operation. This operation consists in glueing two neighboring triangles into a quadrilateral, and then split it again into two triangles in the unique other possible way, or equivalently, in exchanging the two diagonals of this quadrilateral.

We denote

$$\mathcal{A}_n = \mathcal{F}(\Delta_n)$$

as this graph is isomorphic to the 1-skeleton of the associahedron [6]. In particular, note that in this case $\mathcal{F}(\Delta_n)$ can be given a geometric interpretation because its vertices are in one-to-one correspondence with the Euclidean triangulations of any convex n -gon.

For future reference we denote

$$\mathcal{T}_n = \mathcal{F}(\Delta_n^*)$$

This graph also has a geometric interpretation, but in terms of hyperbolic geometry. Take a hyperbolic ideal polygon H_n with n ideal points and with an interior cusp. The graph \mathcal{T}_n then can be seen as the set of geodesic ideal triangulations of H_n with the flip relations simply because every isotopy class of a triangulation is realized by a unique ideal geodesic triangulation. Although we will not make specific use of this model for the flip-graph, it is always good to keep in mind that geometric interpretations aren't always Euclidean. However, in the remainder of the article, we will use the word *geometric* for triangulations of Euclidean objects.

We close this discussion on \mathcal{T}_n by observing that it is closely related to the graph of the type D associahedron [1], which we denote by \mathcal{D}_n . We refer to [1] for a precise definition of \mathcal{D}_n , but briefly said, although the vertices of these graphs consist in collections of arcs of the same cardinality, \mathcal{D}_n differs from \mathcal{T}_n in that multiple loops around the puncture are allowed to coexist. There is also a flip operation allowing to pass from one such loop to another. It is interesting to note that \mathcal{T}_n is a proper subgraph of \mathcal{D}_n , a polytopal graph, and as we will discover in Section 3, they both have the same diameter.

2.2. Geometric flip-graphs

Let P be a possibly non-convex simple Euclidean n -gon. As in the topological case, a triangulation of P is a maximal set of pairwise non-crossing arcs, with the only difference that the arcs we choose from are the Euclidean line segments contained in P and connecting two vertices of P . As mentioned above, the triangulations of Δ_n can be viewed as the triangulations of any convex n -gon. This property fails with non-convex polygons: some segment between two vertices of such a polygon P lies outside of it and thus the triangulations of Δ_n containing the corresponding arc cannot be realized as triangulations of P . We denote by $\mathcal{F}(P)$ the geometric flip-graph of P , whose vertices are the triangulations of P and whose edges connect any two triangulations that differ by a single arc. Note here that flips, as operations, consist in exchanging the diagonals of a convex quadrilateral. Convexity makes it sure that the exchanged arcs are both contained in the polygon.

As shown in [4], $\mathcal{F}(P)$ is connected. Note that $\mathcal{F}(P)$ is, in general, a subgraph of

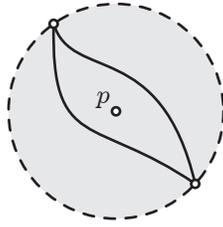


Figure 1: Two non-isotopic arcs with the same vertex set.

\mathcal{A}_n . In other terms, if P is a simple Euclidean n -gon, we have a natural map

$$\mathcal{F}(P) \hookrightarrow \mathcal{A}_n$$

where \hookrightarrow denotes a simplicial injection.

We now turn to the case of punctured polygons. We consider a fixed convex Euclidean n -gon P and denote by P^* a geometric “punctured” version of it. So P^* is really a choice of a $(n + 1)$ -th marked point p in the interior of P . A triangulation of P^* is, again, a maximal set of pairwise non-crossing arcs. In this case, however, the arcs we choose from are the Euclidean line segments whose two endpoints are either two vertices of P , or a vertex of P and the puncture p . Note that, by the maximality of a triangulation as a set of edges, all marked points including p are vertices of it. We denote by $\mathcal{F}(P^*)$ the corresponding geometric flip-graph, that can be defined exactly as in the above two cases. In particular, as an operation, a flip still exchanges the diagonals of convex quadrilaterals.

Triangulations of Δ_n^* cannot all be realized as triangulations of P^* either. In addition, $\mathcal{F}(P^*)$, as a subgraph of \mathcal{T}_n , is always proper: it is impossible to choose P^* in such a way that every topological arc can be realized geometrically. To see this, note that there are two topological arcs between any pair of outer marked points on Δ^* , depending on which side of the puncture they lie (see Fig. 1). However, at most one of them can be realized as an arc of a triangulation of P^* .

In summary, for every choice of P^* , we get a simplicial embedding

$$\mathcal{F}(P^*) \hookrightarrow \mathcal{T}_n$$

which is never onto.

2.3. Pointihedra

The last flip-graph we consider in this article can be obtained as the 1-skeleton of a polytope, just as $\mathcal{F}(\Delta_n)$ can be obtained as the 1-skeleton of an associahedron, and for this reason, it is of special interest.

Let P be a convex polygon and P^* the same polygon with a (fixed) marked point in its interior, as above. We need to define flips that insert the puncture inside a triangulation of P (or, in an equivalent way, delete the puncture from a triangulation of P^*). Consider a triangulation T of P . If the puncture lies in the interior of a triangle t of T , then one can flip the puncture into T by introducing the three edges that connect it to the vertices of t , resulting in a triangulation of P^* . If the puncture lies on an edge ε of T , then call q the quadrilateral obtained by glueing the two triangles of T incident to ε . The puncture can be flipped into T by first removing ε from T and then re-triangulating q with the four edges that connect the puncture to its vertices, which also results in a triangulation of P^* . We denote by $\overline{\mathcal{F}}(P^*)$ the graph obtained by first taking the disjoint union of $\mathcal{F}(P)$ and $\mathcal{F}(P^*)$, and by then adding an edge between a triangulation of P and a triangulation of P^* whenever the latter can be obtained by flipping the puncture into the former.

As mentioned above, there are polytopes whose graph is isomorphic to $\overline{\mathcal{F}}(P^*)$. We will call any such polytope a *pointihedron*. For instance, consider the set of points made up of p together with all the vertices of P . The secondary polytope of this set is a pointihedron, just as the secondary polytope of the vertex set of P is an associahedron. We refer the interested reader to [3, 7] for more details about secondary polytopes. Note that, while we have two natural embeddings

$$\mathcal{F}(P^*) \hookrightarrow \overline{\mathcal{F}}(P^*),$$

and

$$\mathcal{F}(P) \hookrightarrow \overline{\mathcal{F}}(P^*),$$

only $\mathcal{F}(P)$ will be shown to be a geodesic subgraph of $\overline{\mathcal{F}}(P^*)$ for some placements of the puncture, or in the terminology we now introduce, only the latter embedding is strongly convex for such placements.

If Y is a graph, we say that a subgraph X of Y is *strongly convex in Y* if all the geodesics between any two vertices of X lie entirely in X . Similarly an embedding $X \hookrightarrow Y$ between graphs is said to be strongly convex if the image of X is strongly

convex inside Y . In particular this implies that the intrinsic and extrinsic geometries of X coincide. One of the key ingredients we shall use in the sequel is the strong convexity of many of the embeddings we introduced above.

2.4. Remarks on embeddings between flip-graphs

Although we will not make explicit use of this here, it is interesting to note that embeddings between topological flip-graphs are pretty well understood. The most natural way of getting an embedding is by taking all the triangulations of a surface that contain a fixed set of arcs. This provides a copy of the flip-graph of the surface cut along this set of arcs. Except for some low complexity situations, all flip-graph embeddings arise in this way [5]. This works for both finite and infinite flip-graphs and tells, for example, that whenever there is a copy of the flip-graph of, say a 17-gon inside the flip-graph of a 31-gon, there is a collection of 14 arcs that belong to each and every triangulation in this copy.

In [2], it is shown that, in full generality, these natural copies of flip-graphs are strongly convex. Specifically let μ be a multi-arc on a surface Γ , $\mathcal{F}(\Gamma)$ the flip-graph of Γ and $\mathcal{F}(\Gamma \setminus \mu)$ is the flip-graph of $\Gamma \setminus \mu$. The natural simplicial embedding

$$\mathcal{F}(\Gamma \setminus \mu) \hookrightarrow \mathcal{F}(\Gamma),$$

corresponding to looking at all triangulations containing μ , is strongly convex.

3. THE PUNCTURED DISK

We obtain the diameter of $\mathcal{F}(\Delta_n^*)$ in this section using techniques already used in [8, 9, 10, 11]. We will not need to use these techniques in all their generality here and will only describe the needed notions along the way. The interested reader is referred to the above mentioned articles for details about them. The claimed result will be obtained by giving matching lower and upper bounds on $\text{diam}(\mathcal{F}(\Delta_n^*))$. The proof of the upper bound is straightforward and we give it first.

Lemma 3.1. *Any two triangulations of the once punctured disk Δ_n^* can be transformed into one another using at most $2n - 2$ flips.*

Proof. Consider a triangulation T of Δ_n^* , and recall that it has exactly n interior arcs. If the puncture is not incident to all the interior arcs of T , it is always possible

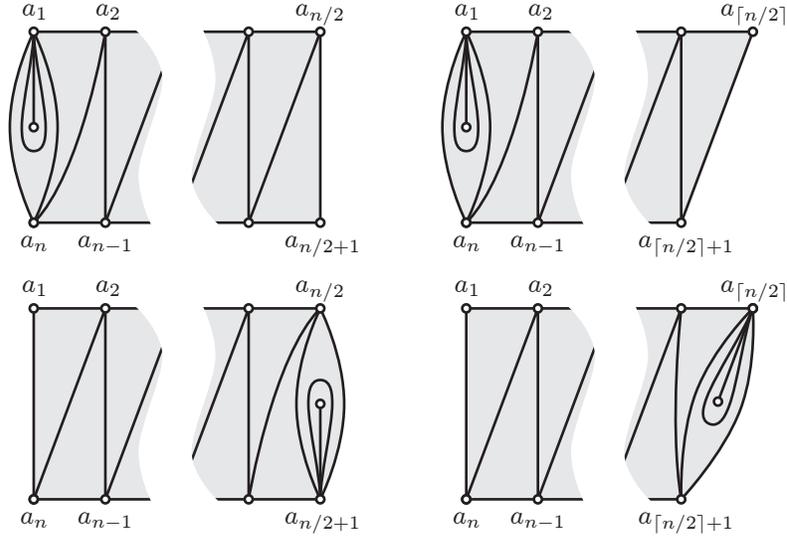


Figure 2: The triangulations A_n^- (top row) and A_n^+ (bottom row) of the once punctured disk, depicted when n is even (left) and odd (right).

to introduce one additional such arc by performing a single flip. Moreover, the puncture is already incident to at least one interior arc of T . Hence, it is possible to transform T into the triangulation of Δ_n^* whose interior arcs are all incident to the puncture by at most $n - 1$ flips. Therefore, any two triangulations of Δ_n^* can be transformed into one another using at most $2n - 2$ flips, as claimed. \square

The lower bound will be obtained by showing that the two triangulations A_n^- and A_n^+ depicted in Fig. 2 have distance at least $2n - 2$ in $\mathcal{F}(\Delta_n^*)$. Most of the interior arcs in these triangulations are arranged as a zigzag. The puncture is placed at one end of the zigzag in A_n^- and at the other end in A_n^+ . It is linked to the boundary by a single interior edge and surrounded by a loop twice incident to the boundary. Note that, when n is equal to 1, 2, or 3, the depiction in Fig. 2 may be ambiguous. For this reason, A_n^- and A_n^+ are shown in Fig. 3 when $1 \leq n \leq 3$.

Lemma 3.2. *At least $2n - 2$ flips are required to transform A_n^- into A_n^+ .*

Proof. The proof proceeds by induction on n . When $n = 1$, A_n^- and A_n^+ are identical as shown in Fig. 3. Hence, they have distance 0 in $\mathcal{F}(\Delta_n^*)$ which is equal to $2n - 2$ in this case. Assume that the result holds for some $n \geq 1$.

Consider the triangulation U of Δ_n^* obtained from A_{n+1}^- by displacing vertex a_{n+1}

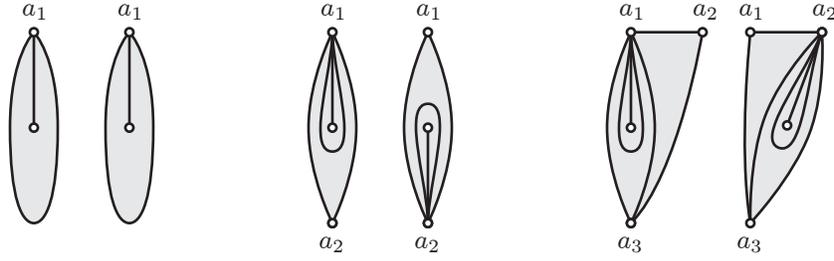


Figure 3: Triangulations A_1^- , A_1^+ , A_2^- , A_2^+ , A_3^- , and A_3^+ (from left to right).

to a_1 within the boundary, while preserving the incidences between a vertex and an arc. Note that, after this operation the loop that was twice incident to a_1 and the interior arc whose endpoints were a_1 and a_{n+1} have become isotopic. Therefore, we assume that one of them has been removed from U . This operation is referred to in [8, 9, 10] as the *deletion of vertex a_{n+1} from A_{n+1}^-* . Indeed, the identification of vertices a_1 and a_{n+1} can be thought of as the removal of the latter. Performing the same operation inside A_n^+ results in a triangulation V of Δ_n^* . Two arcs are made isotopic in this case as well, that cannot both be kept in V : the boundary arc between a_1 and a_2 and the interior arc of A_{n+1}^+ with vertices a_2 and a_n .

Observe that relabelling vertex a_i by a_{n-i+2} within U and V for all i so that $2 \leq i \leq n$ results in A_n^- and A_n^+ respectively. In particular, the number of flips required to transform U into V is also the number of flips required to transform A_n^- into A_n^+ . We now investigate how the deletion of vertex a_{n+1} affects a sequence of flips that changes triangulations A_{n+1}^- and A_{n+1}^+ into one another.

Consider a sequence T_0, \dots, T_k of triangulations of Δ_{n+1}^* such that $T_0 = A_{n+1}^-$, $T_k = A_{n+1}^+$, and whenever $1 \leq i \leq k$, T_{i-1} can be transformed into T_i by a flip. This sequence of triangulations can be thought of as a sequence of k flips that transforms A_{n+1}^- into A_{n+1}^+ . We will assume that k is the least number of flips required to do so. Deleting a_{n+1} from triangulations T_0, \dots, T_k results in a sequence of triangulations that starts with U , ends with V , and so that two consecutive triangulations are either related by a flip, or identical. Indeed, consider the quadrilateral q modified by the flip that changes T_{i-1} into T_i . If this flip affects the triangle incident to the boundary arc α with endpoints a_{n+1} and a_1 (that is, if α is an edge of q), then the deletion changes q into a triangle. In this case, T_{i-1} is identical to T_i . If, on the contrary, α is not an edge of q , then the two triangulations obtained by deleting

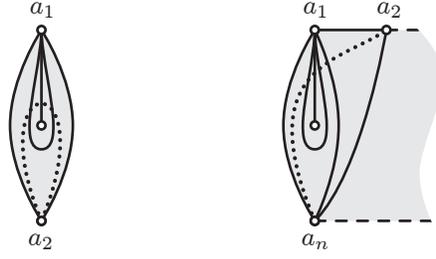


Figure 4: A flip that replaces t^- by t^+ , depicted when $n = 2$ (left) and when $n > 2$ (right). The introduced edge is dotted.

a_{n+1} in T_{i-1} and in T_i are still related by the flip that modifies q . In particular, deleting a_{n+1} from T_0, \dots, T_k provides a sequence of $k - l$ flips that transforms U into V , where l is the number of flips that modify the triangle incident to α in the original sequence. Hence, the number of flips required to transform U into V , or equivalently A_n^- into A_n^+ , is at most $k - l$. This number of flips is, by induction, at least $2n - 2$. As a consequence,

$$2n - 2 \leq k - l. \quad (1)$$

Call t^- and t^+ the triangles incident to α in respectively A_{n+1}^- and A_{n+1}^+ . As t^- and t^+ are distinct, l must be positive. It turns out that l is necessarily greater than 1. Indeed, otherwise, a single flip replaces t^- by t^+ along our sequence of flips, as shown in Fig. 4. As can be seen, the arc introduced by this flip crosses two arcs of the triangulation it is performed in: the arc incident to the puncture and an edge of t^- (which is crossed twice). Hence, such a flip is impossible, and l must be greater than 1. According to inequality (1), the number of flips required to transform A_{n+1}^- into A_{n+1}^+ is then at least $2n$, which completes the proof. \square

The following theorem is a direct consequence of the last two lemmas.

Theorem 3.3. *For every positive integer n , $\text{diam}(\mathcal{F}(\Delta_n^*)) = 2n - 2$.*

Remark 3.4. As mentioned in the preliminaries, \mathcal{T}_n is a proper subgraph of \mathcal{D}_n the type D associahedron graph. In [1] it is shown that

$$\text{diam}(\mathcal{D}_n) = 2n - 2$$

but the vertices used to prove this lower bound are very different from the ones we exhibit. In fact one of them does not belong to \mathcal{T}_n and the other is distance at

most $n - 1$ from all vertices of \mathcal{T}_n . It seems plausible, though, that the diametrically opposite vertices we exhibit are still diametrically opposite in \mathcal{D}_n .

4. GEOMETRIC CASES

In this section, we consider the flip-graphs of Euclidean polygons or Euclidean punctured polygons. We begin with non-convex polygons.

4.1. Flip-graphs of non-convex polygons

If P is a simple, non-convex n -gon, the unique geodesic arc between two different vertices on the boundary is possibly not entirely contained in the polygon. In $\mathcal{F}(P)$ we only consider triangulations made of arcs entirely contained in the polygon and as such, $\mathcal{F}(P)$ is the subgraph of \mathcal{A}_n induced by triangulations that do not contain any forbidden arc. There are many different possible non-convex polygons but we will only look at n -gons with $n - 1$ vertices in convex position.

We begin by observing the following.

Proposition 4.1. *Consider a simple n -gon P . If there is a unique pair of vertices of P so that the segment between them is not contained in P , then $\text{diam}(\mathcal{F}(P)) \geq 2n - 10$.*

Proof. First note that if n is less than 5, then $2n - 10$ is negative and the result is immediate. We therefore assume that n is at least 5.

Explicit examples of triangulations in \mathcal{A}_n at distance at least $2n - 10$ are given in [10]. It is easy to see that they also belong to $\mathcal{F}(P)$: indeed, $\mathcal{F}(P)$ contains all the triangulations of \mathcal{A}_n with at least one interior edge incident to the non-convex vertex v of P . The two triangulations from [10] each have exactly two vertices that are not incident to any of their interior edges. As $n \geq 5$, one can therefore place them in P such that they each have at least one interior edge incident to v .

As $\mathcal{F}(P)$ is a subgraph of \mathcal{A}_n , the distance of these triangulations in the former graph cannot be less than their distance in the latter. \square

In general we have the following slightly weaker lower bound.

Proposition 4.2. *If P is a simple polygon with $n - 1$ vertices in convex position then*

$$\text{diam}(\mathcal{F}(P)) \geq 2n - 16.$$

Proof. Any such polygon can be cut along an arc α between two outer marked points such that the two resulting polygons are at worst “slightly” non-convex in the sense of the previous theorem. Suppose that the resulting polygons are of size n^- and n^+ : we have $n^- + n^+ = n + 2$. The flip-graph of all triangulations containing α is strongly convex inside \mathcal{A}_n [11]. According to Proposition 4.1, the diameter of the individual flip-graphs are bounded below by $2n^- - 10$ and $2n^+ - 10$, so the diameter of $\mathcal{F}(P)$ is bounded below by

$$2(n^- + n^+) - 20 = 2n - 16,$$

as desired. □

4.2. Punctured Euclidean polygons

We now explore the flip-graph of a convex n -gon P^* with a single marked point p in its interior. As explained previously, for every $n \geq 3$, we can think of this choice as being the choice of p within a fixed regular Euclidean n -gon. We are mainly interested in global properties (that do not depend on the choice of p) but some specific properties for certain placements of p will be given as well.

We recall from the preliminary section that for any choice of P^* there is a natural simplicial embedding $\mathcal{F}(P^*) \hookrightarrow \mathcal{T}_n$ which is never onto. We show that this embedding has a geometric meaning.

Theorem 4.3. *The embedding $\mathcal{F}(P^*) \hookrightarrow \mathcal{T}_n$ is strongly convex.*

Proof. We will define a projection map π from \mathcal{T}_n to $\mathcal{F}(P^*)$ which acts as the identity on $\mathcal{F}(P^*)$, and sends points at distance 1 to points at distance at most 1. Before defining the map formally, let us describe a physical model of it. Suppose you represent your polygon by putting nails on a wooden board and then place an extra nail in place of the interior marked point. Now take a triangulation of Δ_n^* (or, equivalently, a vertex of \mathcal{T}_n). The image of this triangulation by π is what happens when you make each of its arcs a rubber band and string it between the nails that correspond to its vertices.

We now define map π in a formal way. We begin by looking at how it affects arcs. Let γ be an arc between two (not necessarily distinct) marked points a and b of Δ_n^* . If γ can be represented in P^* by a Euclidean line segment with the correct topology

(this is the case, for instance, when either a or b is the interior marked point) then we define $\pi(\gamma)$ to be this segment. If it cannot, we define a unique geometric image in P^* as a multi-arc. Let α be the unique Euclidean arc between a and the interior marked point p of P^* and β be the unique Euclidean arc between b and p . We define $\pi(\gamma) = \alpha \cup \beta$. Note that if $a = b$ then $\alpha = \beta$.

We now need to show that this map sends triangulations to triangulations. To see this, note that the image of a triangle is always a triangle, possibly with an extra arc from the interior marked point to one its vertices on the boundary. In particular, this means the complement of the image of a triangulation cannot contain any polygon of higher complexity, and so $\pi(T)$ is a triangulation for any $T \in \mathcal{T}_n$.

Triangulations at distance 1 are mapped by π to triangulations at distance at most 1 as claimed. This is obvious since this map cannot increase the number of intersections between the arcs of two triangulations. \square

As a consequence of this, we have the following corollary.

Corollary 4.4. *Let T and T' belong to $\mathcal{F}(P^*)$ with $\alpha \in T \cap T'$ an arc. Then all triangulations along a geodesic path between T and T' contain α .*

Proof. As discussed in the preliminaries, this is true for \mathcal{T}_n . As the embedding of $\mathcal{F}(P^*)$ inside \mathcal{T}_n is strongly convex, this remains true for $\mathcal{F}(P^*)$. \square

Consider the subgraph \mathcal{G} induced in $\mathcal{F}(P^*)$ by the triangulations that share a given multi-arc μ . Corollary 4.4 could also be formulated as the strong convexity of the natural embedding $\mathcal{G} \hookrightarrow \mathcal{F}(P^*)$. This convexity property immediately bounds the diameter of $\mathcal{F}(P^*)$ below by that of \mathcal{G} . If μ splits P^* into smaller polygons P_1, \dots, P_k , then the latter is in turn bounded below as:

$$\sum_{i=1}^k \text{diam}(\mathcal{F}(P_i)) \leq \text{diam}(\mathcal{G}).$$

From this observation, we obtain the following general lower bound that does not depend on the placement of the puncture.

Theorem 4.5. *For any punctured convex n -gon P^* ,*

$$\text{diam}(\mathcal{F}(P^*)) \geq 2n - 12.$$

Proof. Call p the puncture of P^* . Let us consider two distinct arcs α and β , both incident to p . These arcs split P^* into two simple polygons, say P^- and P^+ . Respectively call a and b the vertices of α and β that are distinct from p . One can choose α and β in such a way that either a, b , and p are collinear, or P^+ is convex and the only line segment between two vertices of P^- that is not contained in P^+ is that between a and b . The numbers n^- and n^+ of vertices of P^- and P^+ satisfy

$$n^- + n^+ = n + 4. \quad (2)$$

By Proposition 4.2, the diameters of $\mathcal{F}(P^-)$ and $\mathcal{F}(P^+)$ are bounded below by respectively $2n^- - 10$ and $2n^+ - 10$. Now consider the subgraph induced in $\mathcal{F}(P^*)$ by the triangulations that contain both α and β . It follows from Corollary 4.4 that the embedding of this subgraph into $\mathcal{F}(P^*)$ is strongly convex. Therefore,

$$\text{diam}(\mathcal{F}(P^*)) \geq 2n^- - 10 + 2n^+ - 10.$$

Combining this inequality with (2) completes the proof. □

It turns out that Theorem 4.5 is nearly sharp:

Lemma 4.6. *Any two triangulations of any punctured n -gon P^* can be transformed into one another using at most $2n - 6$ flips.*

Proof. Let P^* be a punctured polygon. We can use a a proof similar to that of Lemma 3.1. As the puncture is incident to at least three arcs of a triangulation of P^* , the additive constant in the bound is indeed -6 instead of -2 . □

The gap between the bounds provided by Theorem 4.5 and Lemma 4.6 is small. For some placements of the puncture, this gap can be made even smaller using the result of [10] together with Corollary 4.4.

Theorem 4.7. *For any convex n -gon P , one can find a placement of the puncture so that the resulting punctured polygon P^* satisfies $\text{diam}(\mathcal{F}(P^*)) \geq 2n - 8$.*

Proof. Let P be a convex n -gon. If n is equal to 3 or 4, then the result is immediate as $\text{diam}(\mathcal{F}(P^*))$ is non-negative. We therefore assume that n is at least 4. According to [10], one can find two triangulations U and V of P whose distance in $\mathcal{F}(P)$ is not

less than $2n - 10$. It can be required that some triangles u and v of respectively U and V have non-disjoint interiors and exactly one common vertex a . Indeed, the triangulations at distance at least $2n - 10$ given in [10] (see Figs. 5 and 6 therein) have such triangles whenever $n \geq 4$. Now place the puncture, which we denote by p , both in the interior of u and in that of v . Call U' and V' the triangulations of P^* obtained by flipping p into respectively U and V .

Consider a sequence T_0, \dots, T_k of triangulations of P^* such that $T_0 = U'$, $T_k = V'$, and whenever $1 \leq i \leq k$, T_{i-1} can be transformed into T_i by a flip. We will use an operation similar to the vertex deletions performed in the proof of Lemma 3.2. By Corollary 4.4, any T_i contains the arc α with vertices a and p . Hence, we can remove this arc from T_i by displacing p to a while keeping the incidences between a vertex and an arc unaffected. Note that two pairs of arcs of T_i are made isotopic by this operation. Removing one arc from each of these pairs results in a triangulation of P . We refer to this operation as the *contraction of T_i along arc α* .

By construction contracting U' and V' along α respectively results in U and V . As in the proof of Lemma 3.2, the triangulations obtained from T_{i-1} and T_i by this contraction are either identical or related by a flip. In fact they are identical when the flip that transforms T_{i-1} into T_i affects one of the two triangles incident to α . Hence, contracting T_0, \dots, T_k along α provides a sequence of $k - l$ flips that transform U into V , where l is the number of flips that affect a triangle incident to α . As at least $2n - 10$ flips are required to change U into V ,

$$2n - 10 \leq k - l. \quad (3)$$

Now recall that triangles u and v have a unique common vertex. It follows that a triangle of U' incident to α cannot be a triangle of V' . As a single flip cannot affect both the triangles incident to α , there must be at least two flips that affect a triangle incident to α along any sequence of flips that transforms U' and V' . As a consequence, $l \geq 2$ and, assuming that k is the least number of flips needed to transform U' into V' , the result follows from inequality (3). \square

5. THE GEOMETRY OF THE POINTIHEDRON

Let P^* be a punctured polygon obtained by placing a puncture within a polygon P . Recall that $\overline{\mathcal{F}}(P^*)$ is obtained by considering both $\mathcal{F}(P)$ and $\mathcal{F}(P^*)$ and connecting

them with the flips that introduce or remove the puncture. It turns out that these additional flips allow for an upper bound of $2n - 7$ on the diameter of $\overline{\mathcal{F}}(P^*)$, which is less by one than our upper bound on $\text{diam}(\mathcal{F}(P^*))$:

Lemma 5.1. *For any punctured n -gon P^* ,*

$$\text{diam}(\overline{\mathcal{F}}(P^*)) \leq \text{diam}(\mathcal{A}_n) + 3.$$

Proof. Let P^* be a punctured polygon obtained by placing a puncture p within a n -gon P . Call F the triangulation of P^* whose all arcs are incident to p . We will exhibit sequences of flips that transform two triangulations of P or P^* into one another. We review three cases depending on whether each of them is a triangulation of P or P^* .

First consider two triangulations of P^* . When the combined degree of p within them is at least 7, then these triangulations can be transformed into one another by at most $2n - 7$ flips via F , by the same procedure as in the proofs of Lemmas 3.1 and 4.6. Since $\text{diam}(\mathcal{A}_n) \geq 2n - 10$, the result follows. If the combined degree of p in the two triangulations is 6 (it cannot be less), then they each can be transformed into a triangulation of P by the flip that removes the puncture. Hence, their distance in $\overline{\mathcal{F}}(P^*)$ is at most $\text{diam}(\mathcal{A}_n) + 2$, and the result holds.

Now consider a triangulation T of P^* . We will transform it into an arbitrary triangulation of P . If the degree of p in T is at least 5, then it can be changed into F by at most $n - 5$ flips. Any triangulation of P can be transformed into F by first flipping p into it and then by introducing all the missing arcs incident to p using at most $n - 3$ flips. This shows that T can be transformed into any triangulation of P by at most $2n - 7$ flips. Since $\text{diam}(\mathcal{A}_n) \geq 2n - 10$, the result follows. If the degree of p in T is at most 4, then it can be transformed into a triangulation of P by at most 2 flips. Hence the distance of T in $\overline{\mathcal{F}}(P^*)$ to any triangulation of P is at most $\text{diam}(\mathcal{A}_n) + 2$, which implies the desired result.

Finally, two triangulations of P can, by definition, be transformed into one another by at most $\text{diam}(\mathcal{A}_n)$ flips, which completes the proof. \square

Remark 5.2. As \mathcal{A}_n has diameter $2n - 10$ whenever n is at least 13, it is a direct consequence of Lemma 5.1 that the diameter of $\overline{\mathcal{F}}(P^*)$ is at most $2n - 7$ for any punctured polygon P^* with at least 13 boundary vertices.

Let us turn our attention to lower bounds. For some placements of the puncture, we obtain a lower bound on the diameter of the pointihedron that matches our lower bound on the diameter of $\mathcal{F}(P^*)$.

Theorem 5.3. *For any convex n -gon P , one can find a placement of the puncture so that the resulting punctured polygon P^* satisfies $\text{diam}(\overline{\mathcal{F}}(P^*)) \geq 2n - 8$.*

Proof. The proof will be similar to that of Theorem 4.7, except that Corollary 4.4 cannot be invoked in this case as the puncture can appear or disappear along a sequence of flips. This will further restrict the placement of the puncture. Consider a convex n -gon P . If n is equal to 3 or 4, then the result is immediate as $\text{diam}(\overline{\mathcal{F}}(P^*)) \geq 0$. We therefore assume that n is at least 4.

As in the proof of Theorem 4.7, we build two triangulations U' and V' of P^* by flipping the puncture into two triangulations U and V of P whose distance in $\mathcal{F}(P)$ is not less than $2n - 10$. We require that some ear u of U (i.e. a triangle with two edges in the boundary of P) and some triangle v of V have non-disjoint interiors and exactly one common vertex a . Note that a is then necessarily the vertex that is incident to both the edges of u in the boundary of P . This requirement is satisfied, for instance, by the triangulations at distance at least $2n - 10$ given in [10] (see Figs. 5 and 6 therein) as soon as $n \geq 4$. Place the puncture, which we denote by p , both in the interior of u and in that of v . Call U' and V' the triangulations of P^* obtained by flipping p into respectively U and V .

Note that, as u is an ear of U , then the arc α between a and p belongs to all the triangulations of P^* . This makes it possible to contract them along α as in the proof of Theorem 4.7, even though we cannot invoke Corollary 4.4.

Consider a sequence T_0, \dots, T_k of triangulations of P or P^* such that $T_0 = U'$, $T_k = V'$, and whenever $1 \leq i \leq k$, T_{i-1} can be transformed into T_i by a flip. We will contract T_i along α whenever T_i is a triangulation of P^* . In this case α belongs to T_i and one can perform this contraction in exactly the same way as in the proof of Theorem 4.7. By construction, this contraction transforms U' and V' in respectively U and V . When T_i is already a triangulation of P , we assume as a convention that the contraction along α does not affect it. The triangulations obtained from T_{i-1} and T_i by the deletion are either identical or related by a flip. As in the proof of Theorem 4.7, they are identical when the flip that transforms T_{i-1} into T_i affects

one of the two triangles incident to α . In particular, this is the case when this flip either introduces the puncture or removes it. Hence, contracting T_0, \dots, T_k along α provides a sequence of $k - l$ flips that transforms U into V , where l is the number of flips that affect a triangle incident to α . As at least $2n - 10$ flips are required to change U into V , the following inequality holds:

$$2n - 10 \leq k - l. \quad (4)$$

As argued in the proof of Theorem 4.7, if the puncture belongs to all triangulations T_0, \dots, T_k , then l is at least 2. Now, if the puncture is removed by some flip between two of these triangulations, then some other flip must reintroduce it later because V' admits p as a vertex. These two flips affect a triangle incident to α , and l must be at least 2 in this case as well. As a consequence, assuming that k is the least number of flips needed to transform U' into V' , the result follows from inequality (4). \square

It seems the above can be improved. In fact we conjecture the following.

Conjecture 5.4. *Consider a convex n -gon P so that $n \geq 13$. For some placement of the puncture in the interior of P , the resulting punctured polygon P^* satisfies*

$$\text{diam}(\overline{\mathcal{F}}(P^*)) = 2n - 7.$$

We point out that we do not know whether it would be reasonable to state this conjecture for any placement of the puncture. We now focus on the geometry of the natural embedding of the associahedron graph inside the pointihedron graph. Some of the arguments in the proof of Theorem 5.3 can be used to show that this embedding is strongly convex when the puncture is placed within a triangle that shares at least two edges with the considered convex polygon. In fact, we can do better. A boundary quadrilateral of a convex polygon is a quadrilateral that shares at least three of its edges with this polygon. We have the following convexity result when the puncture is placed within the interior of the intersection of two boundary quadrilaterals as sketched in Fig. 5.

Theorem 5.5. *Consider a convex n -gon P . Consider two boundary quadrilaterals q^- and q^+ of P that share exactly three vertices. For any placement of the puncture in the interior of $q^- \cap q^+$, the resulting punctured polygon P^* is such that the natural embedding*

$$\mathcal{A}_n \hookrightarrow \overline{\mathcal{F}}(P^*)$$

is strongly convex.

Proof. Assume the puncture p is placed in the interior of $q^- \cap q^+$. Call u, v , and w the three common vertices of q^- and q^+ such that v is between u and w , as shown in Fig. 5. Denote by α the Euclidean arc connecting v to p and by β the one connecting u to w . Consider a sequence T_0, \dots, T_k of triangulations of P or P^* such that T_{i-1} and T_i are related by a flip for $1 \leq i \leq k$. We also assume that both T_0 and T_k are triangulations of P and that k is their distance in $\overline{\mathcal{F}}(P^*)$. We will prove that T_i is a triangulation of P whenever $0 \leq i \leq k$. To do so we will need to use a contraction operation along arc α , very much like in the proof of Theorem 5.3.

We review two cases.

First assume that p belongs to the interior of the triangle with vertices u, v , and w . Note that this triangle is actually a proper subset of $q^- \cap q^+$. Because of this particular placement of the puncture, the arcs that share their vertices with P never cross α . Therefore, α belongs to every triangulation of P^* . In this case, the contraction along α can be defined as in the proof of Theorem 5.3. By construction, the triangulations T'_{i-1} and T'_i obtained by contracting T_{i-1} and T_i along α are either identical or related by a flip. More precisely, T_{i-1} and T_i are identical when the flip that transforms them into one another affects a triangle incident to arc α . Since k is the least number of flips needed to transform T_0 into T_k , two consecutive triangulations in the sequence T'_0, \dots, T'_k cannot be identical because this sequence of triangulations would otherwise provide a sequence of less than k flips that transforms T_0 into T_k . As any flip that introduces the puncture into a triangulation

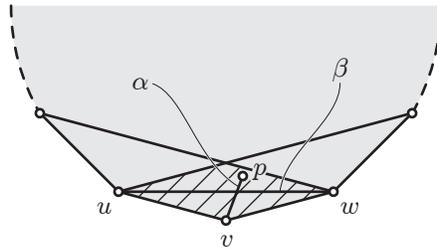


Figure 5: The intersection (hatched) of two boundary quadrilaterals of a larger convex polygon. These quadrilaterals share three vertices u, v , and w . In this representation, the puncture p is placed so that arcs α and β are crossing.

affects a triangle incident to arc α , then such a flip cannot be the one that transforms T_{i-1} into T_i and the result follows.

Now assume that p does not belong to the interior of the triangle with vertices u , v , and w . In this case, we define the contraction of T_i along α as above when T_i is a triangulation of P or when T_i contains α . If T_i is a triangulation of P^* that does not contain α , this operation is still well defined because of our choice of placement of p . As p lies in the interior of $q^- \cap q^+$, if T_i does not contain α , it must contain an arc that crosses α . The arc β is the only possible arc. Note that T_i also contains the two arcs connecting p to u and to w because any arc that crosses one of the two must also cross β and so cannot belong to T_i .

Therefore, when T_i is a triangulation of P^* that contains β , we will define the contraction of T_i along α as the operation that first performs the flip that replaces β by α within T_i , and then contracts the resulting triangulation along α .

By construction, the triangulations T'_{i-1} and T'_i obtained by contracting T_{i-1} and T_i along α are still either identical or related by a flip. In fact, when p does not lie in β , T'_{i-1} and T'_i are identical exactly when the flip that transforms T_{i-1} into T_i affects a triangle incident to arc α , as in the proof of Theorem 5.3. When p lies in β , this is still true, except if T_{i-1} (resp. T_i) contains β and the flip that transforms T_{i-1} into T_i introduces (resp. removes) the puncture. It will be important to remark that, if T_i is obtained from T_{i-1} by a flip that introduces the puncture, and neither triangulations contain β , then T'_{i-1} and T'_i are identical because this flip must introduce α , and because we are not in the exception just mentioned.

As k is the least number of flips required to transform T_0 into T_k , the sequence of triangulations T'_0, \dots, T'_k still provides a geodesic between T_0 and T_k . Hence T'_{i-1} and T'_i are never identical. Assume that some triangulation among T_0, \dots, T_k is a triangulation of P^* . Let j be the least index so that T_j is a triangulation of P^* and l be the index so that T_i is a triangulation of P^* when $j \leq i < l$ and T_l is a triangulation of P . As p does not belong to the interior of the triangle with vertices u , v , and w , contracting a triangulation of P^* along α results in a triangulation that cannot contain β . This is the case for triangulations T'_j to T'_{l-1} . As a consequence, T_{j-1} and T_l cannot both contain β . Otherwise, by Lemma 3 from [11], T'_j, \dots, T'_{l-1} would contain β . Assume without loss of generality that this flip is the one that introduces the puncture within T_{j-1} . In this case, neither T_{j-1} , nor T_j contains β , and by the

above remark, T'_{j-1} is identical to T'_j , a contradiction. □

In light of the result above, we suggest the following conjecture.

Conjecture 5.6. *Let P^* be a punctured convex n -gon. The natural embedding*

$$\mathcal{A}_n \hookrightarrow \overline{\mathcal{F}}(P^*)$$

is strongly convex regardless of the placement of the puncture.

The method we use to prove Theorem 5.5 is highly dependent on the placement of the puncture. In order to prove Conjecture 5.6, it will probably be required to use a different approach which relies less on the local combinatorics.

In addition to \mathcal{A}_n , the other natural embedded subgraph of $\overline{\mathcal{F}}(P^*)$ is $\mathcal{F}(P^*)$ and we could ask about the geometry of its embedding. We shall show that the situation is very different by considering the two triangulations of a heptagon shown in Fig. 6. Using these, we'll show that the natural embedding $\mathcal{F}(P^*) \hookrightarrow \overline{\mathcal{F}}(P^*)$ is generally not strongly convex.

Theorem 5.7. *Consider a convex n -gon P where $n \geq 7$. One can find a placement of the puncture so that, for some triangulations U and V of the resulting punctured polygon P^* , no geodesic between U and V is entirely contained in $\mathcal{F}(P^*)$.*

Proof. We just need to prove that the result holds for the triangulations shown in Fig. 6. In order to extend it to triangulations of a larger punctured n -gon P^* , it suffices to glue the triangulations of Fig. 6 along one boundary edge to a fixed triangulation of a $(n - 5)$ -gon. By Corollary 4.4, the geodesics between the two resulting triangulations of P^* will keep the desired property.

First observe that the triangulations shown in Fig. 6 differ by 6 interior arcs. Hence a geodesic between them that never removes the puncture must have length at least



Figure 6: Two triangulations of a punctured Euclidean heptagon.

6. In fact, the length of such a geodesic is at least 7 because none of the flips that can be performed in either of these triangulations introduces an edge of the other triangulation. Now observe that flipping the puncture out of the triangulation shown on the left of Fig. 6 results in a triangulation of the heptagon whose four interior edges are incident to the same vertex. Such a triangulation is distant by at most 4 flips from any other triangulation of the heptagon. Hence, the two triangulations in Fig. 6 are at most 6 flips away if the removal and the insertion of the puncture are allowed. \square

REFERENCES

- [1] C. Ceballos and V. Pilaud, *The diameter of type D associahedra and the non-leaving-face property*, *Eur. J. Combin.* **51** (2016), 109–124.
- [2] V. Disarlo and H. Parlier, *The geometry of flip graphs and mapping class groups*, arXiv:1411.4285 (2014).
- [3] I. M. Gel'fand, M. M. Kapranov, and A. V. Zelevinsky, *Discriminants of polynomials of several variables and triangulations of Newton polyhedra*, *Leningrad Math. J.* **2** (1991), 449–505.
- [4] F. Hurtado, M. Noy, and J. Urrutia, *Flipping edges in triangulations*, *Discrete Comput. Geom.* **22** (1999), 333–346.
- [5] J. Aramayona, T. Koberda, and H. Parlier, *Injective maps between flip graphs*, arXiv:1409.7046 (2014).
- [6] C. W. Lee, *The associahedron and triangulations of the n -gon*, *Eur. J. Comb.* **10** (1989), 551–560.
- [7] J. A. De Loera, J. Rambau, and F. Santos, *Triangulations: structures for algorithms and applications*, *Algorithms and Computation in Mathematics*, vol. 25, Springer, 2010.
- [8] H. Parlier and L. Pournin, *Flip-graph moduli spaces of filling surfaces*, arXiv:1407.1516 (2014).
- [9] ———, *Modular flip-graphs of one holed surfaces*, arXiv:1510.07664 (2015).

- [10] L. Pournin, *The diameter of associahedra*, Adv. Math. **259** (2014), 13–42.
- [11] D. Sleator, R. Tarjan, and W. Thurston, *Rotation distance, triangulations, and hyperbolic geometry*, J. Am. Math. Soc. **1** (1988), 647–681.

Addresses:

Department of Mathematics, University of Fribourg, Switzerland

LIPN, Université Paris 13, Villetaneuse, France

Emails: hugo.parlier@unifr.ch, lionel.pournin@univ-paris13.fr