

MODULAR FLIP-GRAPHS OF ONE HOLED SURFACES

Hugo Parlier*, Lionel Pournin[†]

Abstract. We study flip-graphs of triangulations on topological surfaces where distance is measured by counting the number of necessary flip operations between two triangulations. We focus on surfaces of positive genus g with a single boundary curve and n marked points on this curve; we consider triangulations up to homeomorphism with the marked points as their vertices. Our main results are upper and lower bounds on the maximal distance between triangulations depending on n and can be thought of as bounds on the diameter of flip-graphs up to the quotient of underlying homeomorphism groups. The main results assert that the diameter of these quotient graphs grows at least like $5n/2$ for all $g \geq 1$. Our upper bounds grow at most like $[4 - 1/(4g)]n$ for $g \geq 2$, and at most like $23n/8$ for the torus.

1. INTRODUCTION

The set of triangulations of a given surface can be given a geometry by defining the distance between two triangulations as the necessary number of flip operations needed to transform one of them into the other. As stated, this definition is somewhat vague, but if the surface is a (Euclidean) polygon and we think of triangulations as being geometric (realized by line segments), then a flip consists in removing an edge from a triangulation and replacing it by the unique other edge so that the result is still a triangulation of the polygon. Given this geometry, the set of triangulations of a polygon is well studied: it is the graph of the associahedron [3, 8, 9, 10]. In particular its diameter is $2n - 10$ for $n > 12$ [6, 7].

The graph of the associahedron can be defined in purely topological terms: a polygon is a topological disk with n marked points on its boundary and a triangulation is a maximal collection of disjoint isotopy classes of arcs with endpoints in the marked points. (The isotopies are relative to the endpoints and by disjoint it is meant disjoint in their interiors.) Flip transformations are topological; one way of

*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

defining them is that two triangulations are related by a flip if they differ only in one edge. This topological description works when one replaces the disk by any finite type surface with marked points, provided there are marked points on boundary curves. However, provided the surface has enough topology (for instance if it has positive genus), the flip-graph is infinite and has infinite diameter. The mapping class group (self-homeomorphisms up to isotopy) of the underlying surface acts nicely on it: indeed it is basically the isomorphism group of the graph (see Section 2 for precise definitions, statements, and references). As such, we can quotient flip-graphs to obtain well defined finite graphs which represent all possible types of triangulations and we can measure distances between them. We call these graphs *modular flip-graphs* and we are interested in their geometry.

In a previous paper [5], we explored the diameter of these graphs for *filling surfaces*: these surfaces have a privileged boundary curve, but otherwise arbitrary topology. In particular, they were allowed to have interior marked points, more than one boundary, and arbitrary genus. Here we focus our attention to the special case of *one holed surfaces* that have a unique boundary curve and no interior marked points. We investigate the growth of the diameter of the corresponding modular flip-graphs in function of the number of marked points on the boundary, while the genus is fixed. These points are *labelled*: we quotient by homeomorphisms that leave them individually fixed.

The case that we focus on most is the torus; it provides a natural variant on the case of the disk and, as we shall see, is already quite intriguing. For this surface we are able to prove the following.

Theorem 1.1. *Let Σ_n be a one holed torus with n marked points on the boundary and let $\mathcal{MF}(\Sigma_n)$ be its modular flip-graph. Then*

$$\left\lfloor \frac{5}{2}n \right\rfloor - 2 \leq \text{diam}(\mathcal{MF}(\Sigma_n)) \leq \frac{23}{8}n + 8.$$

There is a clear gap between our upper and lower bounds (on the order of $3n/8$ for large n) that we are unable to close; in fact it is even tricky to guess what the correct growth rate might be. It could very well be that the lower bounds can be improved but that the necessary combinatorics are currently out of our reach or, on the contrary, it might be that the upper bounds are considerably improvable (or

both!). We point out that, in the instances where matching upper and lower bounds are known [5, 6, 7] the lower bounds have always been the more difficult ones to obtain.

The methods of the above theorem can be used to deduce more general results about surfaces of genus $g \geq 1$.

Theorem 1.2. *Let Σ_n be a one holed surface of genus $g \geq 1$ with n marked points on the boundary and let $\mathcal{MF}(\Sigma_n)$ be its modular flip-graph. Then*

$$\left\lfloor \frac{5}{2}n \right\rfloor - 2 \leq \text{diam}(\mathcal{MF}(\Sigma_n)) \leq \left(4 - \frac{1}{4g}\right)n + K_g,$$

where K_g only depends on g .

As a direct consequence of this result and the results of [5], we obtain the following.

Corollary 1.3. *Let Σ_n be a filling surface with fixed topology and n marked points on the privileged boundary. If this surface is not homeomorphic to a disk or a once punctured disk, then the diameter of its modular flip-graph grows at least as $5n/2$.*

2. PRELIMINARIES

We begin this section by defining and describing the objects we are interested in.

Our basic setup is as follows. Consider an orientable topological surface of finite type Σ with a single boundary curve. It has no marked points on it but will be endowed with them in what follows. We'll denote by $g \geq 0$ the genus of Σ (so if $g = 0$ then Σ is a disk).

For any positive integer n , from Σ we obtain a surface Σ_n by placing n marked points on the privileged boundary of Σ . These marked points are *labelled* from a_1 to a_n , clockwise around the privileged boundary. We are interested in triangulating Σ_n and studying the geometry of the resulting flip-graphs. An *arc* of Σ_n is an *essential* isotopy class of non-oriented simple paths between two marked points (not necessarily distinct). The isotopies we consider are relative to endpoints and by essential we mean not isotopic to a single marked point. In particular the boundary curve is cut into n essential arcs between subsequent marked points: we call these arcs boundary arcs and denote by α_p the boundary arc between a_p and the next




Figure 1: The flip that exchanges arcs ε and ε'

marked point clockwise around the privileged boundary. All other arcs are *interior* arcs but unless specifically stated, by arc we'll generally mean interior arc.

A *triangulation* of Σ_n is a maximal collection of arcs that can be realized in such a way that two of them do not cross. For fixed Σ_n , we call the number of interior arcs of a triangulation the *arc complexity* of Σ_n . Note that by an Euler characteristic argument, it increases linearly in n . Although triangulations are not necessarily “proper” triangulations in the usual sense (triangles can share multiple arcs and arcs can be loops), they do cut the surface into a collection of triangles. This is one reason for using the terminology of *arc* of a triangulation instead of an *edge*, the other being to avoid confusion between the edges of a triangulation and the edges of the flip-graphs we'll now define.

The *flip-graph* $\mathcal{F}(\Sigma_n)$ is constructed as follows. Vertices of $\mathcal{F}(\Sigma_n)$ are the triangulations of Σ_n and two vertices share an edge if they coincide in all but one arc. Equivalently they share an edge if they are related by a *flip*, as shown on Fig. 1. This graph is finite if and only if Σ is of genus 0 but it is always connected and any isotopy class of arc can be introduced into a triangulation by a finite number of flips (see for instance [4]).

In the event that $\mathcal{F}(\Sigma_n)$ is infinite, there is a non-trivial natural action of the group of self-homeomorphisms of Σ_n on $\mathcal{F}(\Sigma_n)$. We require that homeomorphisms fix all marked points individually. We denote $\text{Mod}(\Sigma_n)$ the group of such homeomorphisms up to isotopy. Note that once $n \geq 3$, by the action on the boundary of Σ , all such homeomorphisms are necessarily orientation preserving. As we are interested primarily in large n , we do not need to worry about orientation reversing homeomorphisms. We further note that $\text{Mod}(\Sigma_n)$ is almost the automorphism group of $\mathcal{F}(\Sigma_n)$ and had we allowed homeomorphisms to exchange marked points, it would have been exactly the automorphism group [2]. As we do not allow it, the

automorphism group of the resulting graphs is the cyclic group on n elements.

The combinatorial moduli spaces (the modular flip-graphs) we are interested in are thus

$$\mathcal{MF}(\Sigma_n) = \mathcal{F}(\Sigma_n) / \text{Mod}(\Sigma_n).$$

As described above, these are connected finite graphs which we think of as discrete metric spaces where distance comes from assigning length 1 to each edge. In particular, some of these graphs have loops (a single edge from a vertex to itself) but adding or removing a loop gives rise to an identical metric space on the set of vertices so we think of these graphs as not having any loops. Our main focus is on the vertex diameter of these graphs, which we denote $\text{diam}(\mathcal{MF}(\Sigma_n))$, and how these grow in function of n for fixed Σ .

Our lower bounds on $\text{diam}(\mathcal{MF}(\Sigma_n))$ will be obtained using the operation of deleting a vertex from triangulations of Σ_n already used in [5, 6]. (Equivalently one can think of this operation as the contraction of a boundary arc.) Assuming that n is greater than 1, these operations induce graph homomorphisms from $\mathcal{MF}(\Sigma_n)$ to $\mathcal{MF}(\Sigma_{n-1})$. Consider a triangulation T in $\mathcal{MF}(\Sigma_n)$ and a vertex a_p of T . Some triangle t of T , shown on the left of Fig. 2, is incident to arc α_p . As no vertex of T is interior to Σ and $n \geq 2$, the two other arcs of t are distinct. Let β_p and γ_p denote these arcs. Deleting a_p from T consists in displacing a_p along the boundary to the other vertex a_q of α_p within this triangulation, and by removing either β_p or γ_p from it. In particular, this operation removes vertex a_p and arc α_p from the boundary of Σ . Since displacing a_p to a_q makes β_p and γ_p isotopic, as shown on the right of Fig. 2, the removal of one of them results in a triangulation of Σ_{n-1} .




Figure 2: The triangle incident to α_p in a triangulation of Σ_n (left) and what remains of it when a_p is slid along the boundary to the other vertex a_q of α_p (right).

We now describe how modular flip-graphs are affected by vertex deletions.

First observe that the deletion of a_p sends all triangles of T that are not t to triangles. Hence, it transforms a triangulation T that belongs to $\mathcal{MF}(\Sigma_n)$ into a triangulation in $\mathcal{MF}(\Sigma_{n-1})$ (which appropriately has one less triangle). The resulting triangulation will be denoted by $T \setminus p$ (following the notation introduced in [6]).

The way vertex deletions modify the edges of $\mathcal{MF}(\Sigma_n)$ is described in [6] when Σ is a disk and in [5] in the more general case of filling surfaces. For completeness, we quickly describe the process here, focussing only on what will be strictly needed in the sequel. Let us consider a path of $\mathcal{MF}(\Sigma_n)$ between two triangulations U and V , i.e., a sequence $(T_i)_{0 \leq i \leq k}$ of triangulations in $\mathcal{MF}(\Sigma_n)$ where $T_0 = U$, $T_k = V$, and T_{i-1} can be transformed into T_i by a flip for $0 < i \leq k$. This sequence of triangulations can be thought of as a sequence of k flips that transform U into V , and for this reason, we refer to k as its length. A path between U and V is called a *geodesic* if its length is equal to the distance between U and V in $\mathcal{MF}(\Sigma_n)$, which we denote by $d(U, V)$.

Now, call a flip between U and V *incident to arc α_p* when $U \setminus p$ and $V \setminus p$ are identical. The following theorem is borrowed from [5].

Theorem 2.1. *Let n be an integer greater than 1. Further consider two triangulations U and V in $\mathcal{MF}(\Sigma_n)$. If there exists a geodesic between U and V along which at least f flips are incident to arc α_p , then the following inequality holds:*

$$d(U, V) \geq d(U \setminus p, V \setminus p) + f.$$

This theorem allows to compare distances within $\mathcal{MF}(\Sigma_n)$ with distances within $\mathcal{MF}(\Sigma_{n-1})$. It also provides lower bounds on the former that depend on the number of flips incident to a boundary arc α_p along some geodesic between U and V . It is not too difficult to find configurations such that this number is positive: for instance when the triangles incident to α_p in U and in V are distinct. In some situations, this number is greater than 1. Consider, for instance, the triangulations U and V sketched in Fig. 3. One can see that the two boundary arcs α_p and α_q that share vertex a_q are incident to the same triangle of U . In this case, we say that U has an *ear* in vertex a_q . The triangles of V incident to α_p and α_q are distinct and they do not share an arc. (Note, however, that u and v are possibly equal.) In this situation, we have the following result, also borrowed from [5].




Figure 3: Triangulations U (left) and V (right) from Lemma 2.2.

Lemma 2.2. *Consider two triangulations U and V in $\mathcal{MF}(\Sigma_n)$. Further consider two distinct arcs α_p and α_q on the privileged boundary of Σ_n so that a_q is a vertex of α_p . If U has an ear in a_q and if the triangles of V incident to α_p and to α_q do not have a common arc, then for any geodesic between U and V , there exists $r \in \{p, q\}$ so that at least two flips along this geodesic are incident to α_r .*

3. UPPER BOUNDS

We begin by proving our upper bounds on the diameter of $\mathcal{MF}(\Sigma_n)$ where Σ is a one holed torus.

Theorem 3.1. *Consider a one holed surface Σ of genus 1 with no interior marked points. The following inequality holds:*

$$\text{diam}(\mathcal{MF}(\Sigma_n)) \leq \frac{23}{8}n + 8.$$

Proof. Let T be a triangulation in $\mathcal{MF}(\Sigma_n)$. Call an arc of T *parallel to the boundary of Σ* if it is isotopic to an arc that lies within this boundary. Observe that T contains arcs that are not parallel to the boundary, otherwise Σ would be a planar surface.

Let β be an arc that is not parallel to the boundary of Σ and let us consider $\Sigma_n \setminus \beta$. The resulting surface is of genus 0 with two boundaries, both of which contain a copy of β (see Fig. 4).

The triangulation of $\Sigma_n \setminus \beta$ induced by T must contain an arc between the two boundaries. We choose one and call it γ . Note that the corresponding arc of T is non-parallel to the boundary of Σ . Further observe that, by construction, β and γ are non-homotopic relative to boundary. The surface obtained by cutting Σ_n along arcs β and γ is a polygon, denoted by P in the following of the proof. As the cuts




Figure 4: Surface $\Sigma_n \setminus \beta$.

along β and γ each produce two copies of these arcs, P has $n + 4$ vertices.

Consider the n marked points on the boundary of Σ_n with their clockwise cyclic order. The endpoints of β split the marked points into two sets (we consider the endpoints as being part of both). As β and γ are non-homotopic and non-parallel to boundary, the endpoints of γ are not both contained in exactly one of these sets. Note that several or even all of the endpoints of β and γ could be the same marked point in the boundary of Σ_n . The fact that together, the two pairs of endpoints split the boundary vertices into four (possibly empty) sets will be key in the argument.

Now we consider another triangulation T' in $\mathcal{MF}(\Sigma_n)$ and two arcs β' and γ' of T' , exhibited as β and γ were in the case of T . We denote P' the polygon obtained by cutting Σ_n along β' and γ' . There is a sequence of $\lceil n/8 \rceil$ consecutive marked points on the boundary of Σ_n that does not contain any endpoints of β , γ , β' , or γ' except possibly for the first and last vertices in the sequence. The worst case scenario is when the 8 endpoints are evenly spaced around the boundary. We fix a boundary vertex a_0 lying in this sequence.

Let us flip the arcs of the triangulation of P (resp. of P') induced by T (resp. by T') in order to increase the degree of a_0 as much as possible. This results in “fan” like triangulations of P and P' whose interior arcs are all incident to a_0 . Note that the flips can be chosen so that they each increase the degree of a_0 by one. As both polygons have $n + 4$ vertices, at most $n + 1$ flips are required to do so for each triangulation, and at most $2n + 2$ flips for both of them.

We now flip arcs β , γ , β' , and γ' (two of them within each triangulation) to obtain two triangulations in $\mathcal{MF}(\Sigma_n)$ which we denote by U and U' . As a result of these flips, we have a pair of loops twice incident to a_0 in each of the triangulations U and U' . Up to homeomorphism, the pairs of loops, say (δ, ϵ) in the first triangulation




Figure 5: The triangulation of $\Sigma_n \setminus \{\delta, \epsilon\}$ induced by U .

and (δ', ϵ') in the second are the same. We will not flip these arcs anymore. Instead, we will flip the other arcs until the two triangulations coincide; one way of thinking of this is by looking at the triangulations of $\Sigma_n \setminus \{\delta, \epsilon\}$ and $\Sigma_n \setminus \{\delta', \epsilon'\}$ respectively induced by U and U' . As polygons $\Sigma_n \setminus \{\delta, \epsilon\}$ and $\Sigma_n \setminus \{\delta', \epsilon'\}$ are identical up to homeomorphism, we can flip arcs of these triangulations until they coincide.

Note that each of the loops that have been created within U and U' has two copies on the boundary of respectively $\Sigma_n \setminus \{\delta, \epsilon\}$ and $\Sigma_n \setminus \{\delta', \epsilon'\}$. Each of these copies is incident to a triangle which we refer to as a *hand*, like the hands of a clock that revolve around vertex a_0 (see Fig. 5). The unique vertex of each hand that is not a copy of a_0 will be referred to as its *outer vertex*. Our goal is that both triangulations “tell the same time”: when they do they will be identical up to homeomorphism.

By construction, the vertex a_0 belong to a sequence of at least $\lceil n/8 \rceil$ consecutive vertices incident to identical arcs in U and in U' . We now label clockwise the vertices of polygons $\Sigma_n \setminus \{\delta, \epsilon\}$ and $\Sigma_n \setminus \{\delta', \epsilon'\}$ that are not copies of a_0 from a_1 to a_{n-1} . Consider the smallest k so that a_k is the outer vertex of some hand in either U or U' . If a_k is not the outer vertex of a hand in both triangulations, we can move the outer vertex of this hand clockwise using the flip shown on the left of Fig. 6.

We move the hand until either it coincides with a hand in the other triangulation or it is adjacent to another hand in the same triangulation. In the first case we leave both hands as they are; leaving from their common outer vertex, we search clockwise for the next outer vertex either triangulation and repeat the process. If two hands are adjacent, we flip the arc they share as shown on the right of Fig. 6. This flip introduces an arc twice incident to a_0 and allows us to move both hands together. We do so until we either reach the outer vertex of a hand of the other triangulation or if again we meet a third hand on the same triangulation. In the




Figure 6: A flip that moves the outer vertex of a hand by one vertex clockwise in either U or U' (left) and a flip that binds two hands (right). In both cases, the introduced arc is dotted.

first case we flip the arc twice incident to a_0 to unbind the hands and continue with the other hand. In the second case, we introduce an additional arc twice incident to a_0 by a flip, in order to move all three hands together. Again we can repeat the process described above until all four hands have the same outer vertices and the triangulations are the same.

We now count the number of flips that have been necessary to do the whole transformation. We moved the hands around a portion of the boundary spanned by at most $n - \lceil n/8 \rceil$ vertices. In addition, we may have joined or separated hands several times; because we have 4 hands in total, we may have had to perform at most 6 binding or unbinding flips. In total the number of flips did not exceed

$$\frac{7}{8}n + 6.$$

Therefore, together with the first step, we needed no more than

$$\frac{23}{8}n + 8$$

flips to transform T into T' , as claimed. \square

Using the same ideas, one can prove upper bounds for higher genus Σ . According to these, the growth rate of $\text{diam}(\mathcal{MF}(\Sigma_n))/n$ is less than 4 for any genus, improving in this case on the general upper bound of 4 on the growth rate, shown in [5] for any filling surface. The methods are almost identical to the above proof, so we only provide a sketch of the proof noting the differences along the way.

Theorem 3.2. *Let Σ be a one holed surface of genus $g \geq 2$ with no interior marked points.*

The following inequality holds:

$$\text{diam}(\mathcal{MF}(\Sigma_n)) \leq \left(4 - \frac{1}{4g}\right)n + K_g,$$

where K_g is a constant that only depends on g .

Proof. Within any triangulation T in $\mathcal{MF}(\Sigma_n)$, there are $2g$ arcs $\beta_1, \dots, \beta_{2g}$ such that $\Sigma_n \setminus \{\beta_1, \dots, \beta_{2g}\}$ is a polygon. Now given two triangulations, T and T' , we choose such sets of arcs. Cutting along these arcs gives us two triangulated polygons P and P' . We can find a sequence of $n/(8g)$ consecutive vertices on the boundary of Σ_n that do not contain any of the endpoints of the $2g$ arcs on either triangulation. We will call this sequence of vertices the *untouched* vertices. We choose a vertex a_0 in that sequence, and we flip the interior arcs of the triangulations of polygons P and P' induced by T and T' in order to increase the degree of a_0 as much as possible. We further flip the $2g$ arcs along which the cut was performed within each triangulation to get $2g$ loops twice incident to a_0 which generate the topology of Σ_n . The total number of flips that have been performed so far is at most

$$2(n + 4g - 3 + 2g) = 2n + 12g - 6.$$

We can now cut Σ_n along the $2g$ loops within each triangulation to get two new polygons with $n + 4g$ vertices. Each polygon has $4g$ particular boundary arcs obtained as copies of the loops and, therefore, whose endpoints are copies of a_0 . The remaining n boundary arcs of the polygons are the boundary arcs of Σ_n .

Again we have a notion of hands (this time we have $4g$ of them) corresponding to triangles incident to the particular arcs of the polygons. Here we note an important difference with what happened for the torus: in the case of the torus, there was only one type of topological configuration of the loops in Σ_n and there was a self homeomorphism of Σ_n between the sets of $2g$ loops. In general however, this fails to be the case (for instance, one can think of the different ways an octagon can be glued along its boundary arcs in order to get a surface of genus 2). So instead of just requiring that the hands have the same outer vertices, we will perform flips until they all have the same outer vertex. We leave from the last untouched vertex clockwise and, using flips, we move the hands clockwise to the first untouched

vertex. Within each triangulation, this requires

$$n \left(1 - \frac{1}{8g}\right)$$

moving flips and in addition $4g - 1$ binding flips to put the hands together. The two triangulations now only differ on a subsurface of genus g with one boundary and a single marked point on it. They can be made identical by flipping inside this subsurface; this requires at most D_g flips where D_g is the flip graph diameter of this subsurface (and can be made explicit, see [1]).

The total number of flips performed is

$$2n + 12g - 6 + 2n \left(1 - \frac{1}{8g}\right) + 4g - 1 + D_g,$$

which is precisely equal to

$$\left(4 - \frac{1}{4g}\right)n + K_g,$$

where $K_g = 16g - 7 + D_g$. This is the result claimed. \square

4. LOWER BOUNDS

Let Σ be a one holed surface with strictly positive genus g and no interior marked points. Using the techniques from [5, 6], we prove in this section that

$$\text{diam}(\mathcal{MF}(\Sigma_n)) \geq \left\lfloor \frac{5}{2}n \right\rfloor - 2.$$

As in [5, 6] we exhibit two triangulations A_n^- and A_n^+ in $\mathcal{MF}(\Sigma_n)$ whose distance is at least the desired lower bound. These triangulations are built from the triangulation Z_n of Δ_n depicted in Fig. 7, where Δ_n is a disk with n marked vertices on the boundary. The interior arcs of Z_n form a zigzag, i.e., a simple path that alternates between left and right turns. This path starts at vertex a_n , ends at vertex $a_{n/2}$ when n is even, and at vertex $a_{\lceil n/2 \rceil + 1}$ when n is odd. When n is greater than 3, triangulation Z_n has an ear in a_1 and another ear in $a_{\lceil n/2 \rceil + 1}$. When n is equal to 3, this triangulation is made up of a single triangle which is an ear in all three vertices. Observe that Z_n cannot be defined when n is less than 3.

Assume that $n \geq 3$. A triangulation A_n^- of Σ_n can be built by re-triangulating the ear of Z_n in a_1 as follows: first place an arc with vertices a_1 and a_n within the




Figure 7: The triangulation Z_n of Δ_n depicted when n is even (left) and odd (right).

ear. This arc is isotopic to α_n . Next, place a loop γ twice incident to a_1 between the two isotopic arcs, remove the disk enclosed by this loop, and replace it by a triangulation A_1^- of Σ_1 . The resulting triangulation of Σ_n is shown at the top of Fig. 8. In this representation, the triangle t of A_1^- incident to γ is depicted, but the rest of A_1^- is omitted, as indicated by the hatched surface. Observe that, since a_1 is the only vertex of A_1^- , then t must be bounded by γ and two other loops twice incident to a_1 as shown on the figure.

Another triangulation A_n^+ of Γ_n can be built by re-triangulating the ear of Z_n in $a_{[n/2]+1}$ in the same way, except that γ is chosen incident to $a_{[n/2]+1}$ instead of a_1 . Also, a triangulation A_1^+ of Σ_1 , possibly distinct from A_1^- , replaces the disk enclosed by γ . Triangulation A_n^+ is shown in the bottom of Fig. 8.

In this construction, A_1^- and A_1^+ are arbitrary triangulations in $\mathcal{MF}(\Sigma_1)$. We can also define A_2^- and A_2^+ as the triangulations in $\mathcal{MF}(\Sigma_2)$ that contain a loop arc at respectively vertex a_1 and vertex a_2 , in such a way that these arcs respectively enclose triangulations A_1^- and A_1^+ .

In order to obtain the desired lower bound, one of the main steps will be to show the following inequality for every integer n greater than 2:

$$d(A_n^-, A_n^+) \geq \min(\{d(A_{n-1}^-, A_{n-1}^+) + 3, d(A_{n-2}^-, A_{n-2}^+) + 5\}). \quad (1)$$

This inequality will be obtained using well chosen vertex deletions. Observe that for $n \geq 2$, deleting vertex a_n from both A_n^- and A_n^+ results in triangulations isomorphic to A_{n-1}^- and A_{n-1}^+ . More precisely, once a_n has been deleted, the other vertices need be relabeled in order to obtain A_{n-1}^- and A_{n-1}^+ . If the deleted vertex is a_j , the natural way to do so is to relabel by a_{i-1} every vertex a_i so that $i > j$.




Figure 8: The triangulations A_n^- (top row) and A_n^+ (bottom row) of Σ_n depicted when n is even (left) and odd (right). For simplicity, the triangulations of Σ_1 used to build A_n^- and A_n^+ are mostly omitted here (hatched surface), except for two interior loops twice incident to a_1 (top) and to $a_{\lfloor n/2 \rfloor + 1}$ (bottom).

This relabeling provides a map to the triangulations of Γ_{n-1} . In the following, we call any such map a *vertex relabeling*. The above observation can be reformulated as follows: $A_n^- \setminus \! \setminus n$, resp. $A_n^+ \setminus \! \setminus n$ are isomorphic to A_{n-1}^- , resp. A_{n-1}^+ via the same vertex relabeling. This can be checked using Fig. 8.

It then follows from Theorem 2.1 that, if there exists a geodesic between A_n^- and A_n^+ with at least 3 flips incident to a_n , then

$$d(A_n^-, A_n^+) \geq d(A_{n-1}^-, A_{n-1}^+) + 3, \quad (2)$$

and inequality (1) holds in this case. Now assume that $n \geq 3$ and observe that for any integer i so that $1 \leq i < n$ and any $j \in \{n-i, n-i+1\}$, deleting vertices a_i and a_j from A_n^- and from A_n^+ results in triangulations of Γ_n isomorphic to A_{n-2}^- and A_{n-2}^+ respectively by the same vertex relabeling. Hence, if there exists a geodesic path between A_n^- and A_n^+ with at least 3 flips incident to a_i , and a geodesic path between $A_n^- \setminus \! \setminus i$ and $A_n^+ \setminus \! \setminus i$ with at least 2 flips incident to a_j then, invoking Theorem 2.1 twice yields

$$d(A_n^-, A_n^+) \geq d(A_{n-2}^-, A_{n-2}^+) + 5, \quad (3)$$

and (1) also holds in this case. Observe that (2) and (3) follow from the existence of



Figure 9: The j -th (left) and l -th (right) flips performed along path $(T_i)_{0 \leq i \leq k}$ in the proof of Lemma 4.1. In each case, the introduced arc is dotted and the solid arcs belong to triangulation T_{i-1} .

particular geodesic paths. The rest of the section is devoted to proving the existence of geodesic paths that imply at least one of these inequalities.

Observe that α_n is not incident to the same triangle in A_n^- and in A_n^+ . Therefore, at least one flip is incident to α_n along any geodesic from A_n^- to A_n^+ . We will study these geodesics depending on the arc introduced by their first flip incident to α_n ; this is the purpose of the next three lemmas.

Lemma 4.1. *For some integer n greater than 2, consider a geodesic from A_n^- to A_n^+ . If the first flip incident to arc α_n along this geodesic introduces an arc with vertices a_1 and a_n , then there must be at least two other flips incident to α_n along it.*

Proof. Let $(T_i)_{0 \leq i \leq k}$ be a geodesic from A_n^- to A_n^+ . Assume that the first flip incident to α_n along $(T_i)_{0 \leq i \leq k}$ is the j -th one, and that it introduces an arc β with vertices a_1 and a_n . This flip must then be the one shown on the left of Fig. 9.

As the triangles incident to α_n in T_j and in A_n^+ are distinct, at least one of the last $k - j$ flips along our geodesic is incident to this arc. Assume that there is exactly one such flip, say the l -th one. This flip then necessarily introduces an ear in vertex a_1 . Hence, the arc introduced by it must be the one shown as a dotted line on the right of Fig. 9. This is impossible because this arc crosses two interior arcs of T_{l-1} : arc β , and the arc twice incident to a_1 that bounds the triangle incident to α_n in T_{l-1} . Therefore, at least three flips are incident to α_n along $(T_i)_{0 \leq i \leq k}$. \square

The proof of the following lemma is, almost word for word, that of Lemma 4.2 in [5]. We include it for completeness.



Figure 10: The j -th (left) and l -th (right) flip performed along path $(T_i)_{0 \leq i \leq k}$ in the proof of Lemma 4.2. In each case, the introduced arc is dotted and the solid arcs belong to triangulation T_{i-1} .

Lemma 4.2. *Let n be an integer greater than 2. Consider a geodesic from A_n^- to A_n^+ whose first flip incident to α_n introduces an arc with vertices a_1 and a_2 . If α_n is incident to at most 2 flips along this geodesic then α_1 is incident to at least 4 flips along it.*

Proof. Let $(T_i)_{0 \leq i \leq k}$ be a geodesic from A_n^- to A_n^+ whose first flip incident to α_n , say the j -th one, introduces an arc with vertices a_1 and a_2 . This flip must then be the one shown on the left of Fig. 10. Note that it is incident to arc α_1 .

Assume that at most one flip along $(T_i)_{0 \leq i \leq k}$ other than the j -th one is incident to α_n . In this case, there must be exactly one such flip among the last $k - j$ flips of $(T_i)_{0 \leq i \leq k}$, say the l -th one. Moreover, this flip replaces the triangle of T_j incident to α_n by the triangle of A_n^+ incident to α_n . There is only one way to do so, as depicted on the right of Fig. 10. Note that this flip is also incident to α_1 .

Finally, as the arc introduced by the j -th flip along $(T_i)_{0 \leq i \leq k}$ is not removed before the l -th flip, there must be two more flips incident to arc α_1 along this geodesic: the flip that removes the loop arc with vertex a_1 shown on the left of Fig. 10 and the flip that introduces the loop arc with vertex a_2 shown on the right of the figure. This proves that at least four flips are incident to α_1 along $(T_i)_{0 \leq i \leq k}$. \square

The proof of the last of the three lemmas is slightly more involved.

Lemma 4.3. *For $n \geq 4$, consider a geodesic from A_n^- to A_n^+ whose first flip incident to α_n introduces an arc with vertices a_1 and a_p , where $2 < p < n$. Then:*

$$d(A_n^-, A_n^+) \geq \min(\{d(A_{n-1}^-, A_{n-1}^+) + 3, d(A_{n-2}^-, A_{n-2}^+) + 5\}).$$




Figure 11: The j -th flip performed along path $(T_i)_{0 \leq i \leq k}$ in the proof of Lemma 4.3 when $p \leq \lceil n/2 \rceil$ (left) and when $p > \lceil n/2 \rceil$ (right). In each case, the introduced arc is dotted and the solid edges belong to triangulation T_{j-1} .

Proof. Let $(T_i)_{0 \leq i \leq k}$ be a geodesic from A_n^- to A_n^+ whose first flip incident to α_n , say the j -th one, introduces an arc with vertices a_1 and a_p , where $2 < p < n$. This flip is depicted in Fig. 11. As a first step, we will show that T_j has an ear in some vertex a_q where either $2 \leq q < \lceil n/2 \rceil$ or $\lceil n/2 \rceil + 1 < q \leq n$.

Assume that p is not greater than $\lceil n/2 \rceil$. Consider the arc of T_j with vertices a_1 and a_p shown as a solid line on the left of Fig. 11. The portion of T_j bounded by this arc and by arcs $\alpha_1, \dots, \alpha_{p-1}$ is a triangulation of disk Δ_p . If $p > 3$, then this triangulation has at least two ears, and one of them is also an ear of T_j in some vertex a_q where $2 \leq q < p$. As $p \leq \lceil n/2 \rceil$, the desired inequalities on q hold. If $p = 3$, then T_j has an ear in a_2 since the triangulation of Δ_p induced by T_j is made up of a single triangle, and these inequalities on q also hold.

Now assume that $\lceil n/2 \rceil < p < n$. Consider the arc with vertices a_1 and a_p introduced by the j -th flip along $(T_i)_{0 \leq i \leq k}$ and shown as a dotted line on the right of Fig. 11. The portion of T_j bounded by this arc and by arcs $\alpha_p, \dots, \alpha_n$ is a triangulation of Δ_{n-p+2} . Since $n - p + 2$ is at least 3, an argument similar to the one used in the last paragraph shows that T_j has an ear in some vertex a_q where $p < q \leq n$. As p is greater than $\lceil n/2 \rceil$, we obtain $\lceil n/2 \rceil + 1 < q \leq n$, as desired.

Note that, as $2 \leq q < \lceil n/2 \rceil$ or $\lceil n/2 \rceil + 1 < q \leq n$, the triangles incident to arcs α_{q-1} and to α_q in either A_n^- or A_n^+ do not share an arc. Since T_j has an ear in a_q , Lemma 2.2 can be invoked with $U = T_j$ for the two geodesics obtained by splitting $(T_i)_{0 \leq i \leq k}$ at T_j . Doing so, we find that either α_{q-1} and α_q are both incident to exactly 3 flips along this geodesic or one of these arcs is incident to at least 4 flips along it. We will review the two cases separately.

If α_r is incident to at least 4 flips along $(T_i)_{0 \leq i \leq k}$, where r is equal to $q - 1$ or to q , then Theorem 2.1 yields

$$d(A_n^-, A_n^+) \geq d(A_n^- \setminus r, A_n^+ \setminus r) + 4. \quad (4)$$

Call $s = n - q + 1$ and observe that arc α_s is not incident to the same triangle in $A_n^- \setminus r$ and in $A_n^+ \setminus r$. Hence, some flip must be incident to this arc along any geodesic between $A_n^- \setminus r$ and $A_n^+ \setminus r$. Invoking Theorem 2.1 again, we find

$$d(A_n^- \setminus r, A_n^+ \setminus r) \geq d(A_n^- \setminus r \setminus s, A_n^+ \setminus r \setminus s) + 1. \quad (5)$$

As $A_n^- \setminus r \setminus s$ and $A_n^+ \setminus r \setminus s$ are isomorphic to A_{n-2}^- and A_{n-2}^+ by the same vertex relabeling, the desired result is obtained combining (4) and (5).

Now assume that α_{q-1} and α_q are both incident to exactly 3 flips along $(T_i)_{0 \leq i \leq k}$. If $q = n$, then the result immediately follows from Theorem 2.1 because $A_n^- \setminus n$ and $A_n^+ \setminus n$ are isomorphic to A_{n-1}^- and A_{n-1}^+ by the same vertex relabeling. We will therefore also assume that $q < n$.

Note that at least one of the first j flips and at least one of the last $k - j$ flips along $(T_i)_{0 \leq i \leq k}$ are incident to either α_{p-1} or α_p because these arcs are not incident to the same triangles in T_j and in either A_n^- or A_n^+ . We can in fact assume without loss of generality that exactly one of the first j flips and two of the last $k - j$ flips along $(T_i)_{0 \leq i \leq k}$ are incident to α_{q-1} by, if needed, reversing geodesic $(T_i)_{0 \leq i \leq k}$ and exchanging A_n^- with A_n^+ . In this case, exactly two of the first j flips and exactly one of the last $k - j$ flips along this geodesic are incident to α_q .




Figure 12: The portion of either A_n^- or A_n^+ next to vertex a_q (left), the l -th flip along the geodesic used in the proof of Lemma 4.3 (center), and the first flip incident to α_{q-1} along this geodesic (right). The arc introduced by each flip is dotted.

Call $r = n - q + 1$. The portion of either A_n^- or A_n^+ placed between arcs α_{q-1} , α_q , and α_r is depicted on the left of Fig. 12.

Since only one flip is incident to α_{q-1} along $(T_i)_{0 \leq i \leq j}$, this flip must be the one shown on the right of Fig. 12, and it is also incident to α_q . Now consider the triangle incident to α_q when this flip is performed, labeled t in the figure. This triangle must be introduced by the first flip incident to α_q , earlier along the geodesic. This flip, say the l -th one, is shown in the center of Fig. 12.

Consider a geodesic $(T'_i)_{0 \leq i \leq k'}$ from $A_n^- \setminus q - 1$ to $T_l \setminus q - 1$, and a geodesic $(T''_i)_{j \leq i \leq k''}$ from $T_l \setminus q - 1$ to $A_n^+ \setminus q - 1$. Since three flips are incident to α_{q-1} along $(T_i)_{0 \leq i \leq k}$, it follows from Theorem 2.1 that

$$k' + k'' \leq d(A_n^-, A_n^+) - 3. \quad (6)$$

Observe that the triangles incident to α_r in $A_n^- \setminus q - 1$ and in $T_l \setminus q - 1$ are distinct. Hence, at least one flip is incident to α_r along $(T'_i)_{0 \leq i \leq k'}$ and by Theorem 2.1,

$$k' \geq d(A_n^- \setminus q - 1 \setminus r, T_l \setminus q - 1 \setminus r) + 1. \quad (7)$$

Similarly, the triangles incident to α_r in $T_l^- \setminus q - 1$ and in $A_n^+ \setminus q - 1$ are distinct. Hence, at least one flip is incident to α_r along $(T''_i)_{0 \leq i \leq k''}$ and by Theorem 2.1,

$$k'' \geq d(T_l \setminus q - 1 \setminus r, A_n^+ \setminus q - 1 \setminus r) + 1. \quad (8)$$

By the triangle inequality, (7) and (8) yield

$$k' + k'' \geq d(A_n^- \setminus q - 1 \setminus r, A_n^+ \setminus q - 1 \setminus r) + 2. \quad (9)$$

Since $A_n^- \setminus q - 1 \setminus r$ and $A_n^+ \setminus q - 1 \setminus r$ are isomorphic to A_{n-2}^- and A_{n-2}^+ by the same vertex relabeling, the desired inequality is obtained combining (6) and (9). \square

We can now prove the main inequality.

Theorem 4.4. *For every integer n greater than 2,*

$$d(A_n^-, A_n^+) \geq \min(\{d(A_{n-1}^-, A_{n-1}^+) + 3, d(A_{n-2}^-, A_{n-2}^+) + 5\}).$$

Proof. Assume that $n \geq 3$ and consider a geodesic $(T_i)_{0 \leq i \leq k}$ from A_n^- to A_n^+ . If at least 3 flips are incident to α_n along it, then Theorem 2.1 yields

$$d(A_n^-, A_n^+) \geq d(A_{n-1}^-, A_{n-1}^+) + 3.$$

Indeed, as mentioned before, $A_n^- \setminus n$ and $A_n^+ \setminus n$ are respectively isomorphic to A_{n-1}^- and A_{n-1}^+ via the same vertex relabeling. Hence the result holds in this case and we assume in the remainder of the proof that at most two flips are incident to α_n along $(T_i)_{0 \leq i \leq k}$. By Lemma 4.1, the first flip incident to α_n along $(T_i)_{0 \leq i \leq k}$, say the j -th one, must introduce an arc with vertices a_1 and a_p where $2 \leq p < n$. If p is equal to 2 then, by Lemma 4.2 and Theorem 2.1,

$$d(A_n^-, A_n^+) \geq d(A_n^- \setminus 1, A_n^+ \setminus 1) + 4. \quad (10)$$

Observe that arc α_{n-1} is not incident to the same triangle in $A_n^- \setminus 1$ and in $A_n^+ \setminus 1$. Therefore, there must be at least one flip incident to α_{n-1} along any geodesic between these triangulations. Therefore, invoking Theorem 2.1,

$$d(A_n^- \setminus 1, A_n^+ \setminus 1) \geq d(A_n^- \setminus 1 \setminus n - 1, A_n^+ \setminus 1 \setminus n - 1) + 1. \quad (11)$$

As $A_n^- \setminus 1 \setminus n - 1$ and $A_n^+ \setminus 1 \setminus n - 1$ are isomorphic to A_{n-2}^- and A_{n-2}^+ via the same vertex relabeling, the result is obtained combining (10) and (11).

Finally, if the j -th flip introduces an arc with vertices a_1 and a_p , where $2 < p < n$, then n must be greater than 3 and the result follows from Lemma 4.3. \square

We conclude the following.

Theorem 4.5. *Let Σ be a one holed surface of genus $g > 0$ with no interior marked points. The following inequality holds:*

$$\text{diam}(\mathcal{MF}(\Sigma_n)) \geq \left\lfloor \frac{5}{2}n \right\rfloor - 2.$$

Proof. By Theorem 4.4, the result is easily proven by induction on n , provided it holds when n is equal to 1 or to 2. In n is equal to 1, the it is obtained immediately because $\lfloor 5n/2 \rfloor = 2$ and $\text{diam}(\mathcal{MF}(\Sigma_1)) \geq 0$.

Assume that $n = 2$. In this case, we can prove that at least 3 flips are incident to arc




Figure 13: The first (left) and last (right) flips incident to α_2 along any geodesic from A_2^- to A_2^+ . In both cases, the introduced arc is dotted.

α_2 along any geodesic from A_2^- and A_2^+ . Indeed, consider the triangle incident to α_2 in A_2^- . The only arc of this triangle that can be flipped is the loop twice incident to a_1 . As a consequence, the first flip incident to α_2 along any geodesic from A_2^- to A_2^+ is as shown on the left of Fig. 13. Similarly, the last flip incident to α_2 along a such geodesic must be the one depicted on the right of the figure. Note that these flips are distinct as they do not remove the same arc. Now consider the triangles incident to α_2 in the triangulation obtained from the first of these flips and in the triangulation within which the second is performed. These triangles are distinct because one of them has an arc twice incident to a_1 and the other an arc twice incident to a_2 . Hence, there must be a third flip incident to α_2 along any geodesic from A_2^- to A_2^+ . It therefore follows from Theorem 2.1 that $d(A_2^-, A_2^+) \geq 3$, which is precisely the desired inequality when $n = 2$. \square

The additive constant in the lower bound provided by this theorem can be easily improved to take the genus into account. Indeed, recall that A_1^- and A_1^+ have been chosen in an arbitrary way at the beginning of the section. However, as shown in [1] (see Corollary 4.19 therein) the diameter of $\mathcal{MF}(\Sigma_1)$ is bounded below by $K g \log(g)$, for some positive constant K and large enough g . Therefore, triangulations A_1^- and A_1^+ can be chosen in such a way that, from the argument used in the proof of Theorem 4.5, the additive constant in our lower bound is $K g \log(g) - 2$ instead of -2 , provided g is large enough.

Note, however, that this additive constant cannot be improved in the case of the one holed torus, i.e., when $g = 1$. Indeed, in this case, a triangulation of Σ_1 is made up of three triangles: one of them is incident to the boundary arc and adjacent to the other two triangles t^- and t^+ . As $g = 1$, the way t^- and t^+ are glued to

one another is prescribed. It follows that Σ_1 has exactly one triangulation up to homeomorphism and that the diameter of its modular flip-graph is 0 (which is exactly the lower bound provided by Theorem 4.5 when $n = 1$).

REFERENCES

- [1] V. Disarlo and H. Parlier, *The geometry of flip graphs and mapping class groups*, arXiv:1411.4285 (2014).
- [2] M. Korkmaz and A. Papadopoulos, *On the ideal triangulation graph of a punctured surface*, Ann. Inst. Fourier **62** (2012), no. 4, 1367–1382.
- [3] C. W. Lee, *The associahedron and triangulations of the n -gon*, Eur. J. Comb. **10** (1989), 551–560.
- [4] L. Mosher, *Tiling the projective foliation space of a punctured surface*, Trans. Am. Math. Soc. **306** (1988), no. 1, 1–70.
- [5] H. Parlier and L. Pournin, *Flip-graph moduli spaces of filling surfaces*, arXiv:1407.1516 (2014).
- [6] L. Pournin, *The diameter of associahedra*, Adv. Math. **259** (2014), 13–42.
- [7] D. Sleator, R. Tarjan, and W. Thurston, *Rotation distance, triangulations, and hyperbolic geometry*, J. Am. Math. Soc. **1** (1988), 647–681.
- [8] J. Stasheff, *Homotopy associativity of H -spaces*, Trans. Am. Math. Soc. **108** (1963), 275–312.
- [9] ———, *How I ‘met’ Dov Tamari*, Associahedra, Tamari Lattices and Related Structures, Progress in Mathematics, vol. 299, Birkhäuser, 2012, pp. 45–63.
- [10] D. Tamari, *Monoides préordonnés et chaînes de Malcev*, B. Soc. Math. Fr. **82** (1954), 53–96.

Addresses:

Department of Mathematics, University of Fribourg, Switzerland

LIPN, Université Paris 13, Villetaneuse, France

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