THE GEOMETRY OF FLIP GRAPHS AND MAPPING CLASS GROUPS

Valentina Disarlo[†], Hugo Parlier[*]

Abstract. The space of topological decompositions into triangulations of a surface has a natural graph structure where two triangulations share an edge if they are related by a so-called flip. This space is a sort of combinatorial Teichmüller space and is quasi-isometric to the underlying mapping class group. We study this space in two main directions. We first show that strata corresponding to triangulations containing a same multiarc are strongly convex within the whole space and use this result to deduce properties about the mapping class group. We then focus on the quotient of this space by the mapping class group to obtain a type of combinatorial moduli space. In particular, we are able to identity how the diameters of the resulting spaces grow in terms of the complexity of the underlying surfaces.

1. INTRODUCTION

The many relationships between curves, arcs and homeomorphisms of surfaces have provided numerous, rich and fruitful insights into the study of Teichmüller spaces and mapping class groups. In particular, combinatorial structures such as curve, arc and pants complexes have been shown to be closely related to metric structures on Teichmüller spaces and in particular all share the mapping class group as an automorphism group.

Flip graphs are an example of one of these natural combinatorial structures. For a given topological surface with a prescribed set of points, the vertex set of the associated flip graph is the set of maximal multiarcs (which begin and terminate in the prescribed points). Just like the other combinatorial objects, the multiarcs are considered up to isotopy (which preserve the prescribed set of points). As they are maximal, they decompose the surface into triangles and thus we refer to them as triangulations (although they may not be triangulations in the usual sense). Two triangulations share an edge in the flip graph if they

1

are related by a flip - so if they differ by at most one arc. Provided the surface is complicated enough, flip graphs are infinite objects but are always locally finite connected graphs.

Flip graphs can be thought of (and this is the point of view we take in this article) as metric objects by associating length 1 to every edge. As metric spaces they describe how different triangulations are from one another and are a sort of combinatorial analogue of Teichmüller space. With a few exceptions, the mapping class group is again the full automorphism group [16] and as such the finite quotient, which we call a *modular flip graph*, becomes a combinatorial analogue of a moduli space. In contrast to some of the other spaces mentioned before, the action of the mapping class group is proper and, via the Švarc-Milnor lemma, the flip graph is a quasi-isometric model of the mapping class group which makes it an ideal tool for studying its geometry. Mosher [20] implicitly uses the flip graph to study the mapping class group from the combinatorial point of view. This point of view has recently been exploited by Rafi and Tao [29].

Flip graphs of surfaces also appear in a number of other contexts. As hinted at above, flip graphs naturally appear in Teichmüller theory. They appear for instance in Penner's decorated Teichmüller space [25] and in the work of Fomin, Shapiro and D. Thurston ([12] and [13]) in their study of cluster algebras related to bordered surfaces. Flip graphs and some slight variations have been studied in combinatorics and computational geometry by a variety of authors, for instance Negami [22], Bose [6] and De Loera-Rambau-Santos [17].

One of the simplest and most studied flip graphs is the flip graph of a polygon, which turns out to be the 1-skeleton of a polytope called the associahedron [32, 33]. It is a finite graph with a number of remarkable properties including being the graph of a polytope. The celebrated result of Sleator, Tarjan and W. Thurston [30] about the diameter of the associahedron, and proved using 3-dimensional hyperbolic polyhedra, was recently extended by Pournin [26] who also provided a purely combinatorial proof. The diameter of this graph is exactly $2n - 10$ for all $n > 12$. Sleator, Tarjan, and W. Thurston [31] also studied triangulations of spheres up to homeomorphism, which essentially amounts to studying the diameter of a modular flip graph. In this case, they show that the diameter grows like $n \log n$ where n is the number of labelled points on the sphere.

In this article, we study both the geometry of flip graphs and of modular flip graphs. One of the main motivations we have in mind is the study of the mapping class group.

We begin by studying the geometry of flip graphs. Our first main result comes from a very natural question about two triangulations that have an arc a in common. Given any two such triangulations, there is at least one minimal path between them: do all the

triangulations of any minimal path contain the arc a? The answer is yes.

Theorem 1.1. *For every multiarc A, the stratum $\mathcal{F}_A$ is strongly convex.*

In the above result, for any given flip graph, we've denoted $\mathcal{F}_A$ the set of triangulations which contained a prescribed multiarc A. We note that this result for flip graphs of polygons was previously known and an essential tool in [31] and in [26].

We observe that the same question can be asked for the pants graph (where multicurves play the part of multiarcs). For the pants graph, this is known to be true for certain types of multicurves but is in general completely open [4, 5, 3, 34].

We give two applications of this result. It is a recent result of the second author together with Aramayona and Koberda that, under certain conditions, simplicial embeddings between flip graphs only arise *naturally* [2]. By naturally, we mean that the injective simplical map comes from an embedding between the two surfaces. The conditions are on the underlying surface of the flip graph found in the domain. This surface is required to be non-exceptional (or "sufficiently complicated", see Section 3.2 for a precise definition). Now together with the above theorem, this implies the following.

Corollary 1.2. *Suppose Σ is non-exceptional, and let $\mathcal{F}(\Sigma) \to \mathcal{F}(\Sigma')$ be an injective simplicial map. Then $\mathcal{F}(\Sigma)$ is strongly convex inside of $\mathcal{F}(\Sigma')$.*

As geometric properties of flip graphs translate into quasi properties for mapping class groups, we also obtain the following result for mapping class groups. This result also follows by results of Masur-Minsky [18] and Hamenstädt [14].

Corollary 1.3. *For every vertex $T \in \mathcal{F}_A$, there is a commutative diagram:*

$$
\begin{array}{ccc}
\mathcal{F}_A & \lhook\joinrel\longrightarrow & \mathcal{F}(\Sigma) \\
{\scriptstyle \omega_{T|}}\big\uparrow & & \big\uparrow{\scriptstyle \omega_T} \\
\mathrm{Stab}(A) & \lhook\joinrel\longrightarrow & \mathrm{Mod}(\Sigma)
\end{array}
$$

where the inclusion $\mathcal{F}_A \hookrightarrow \mathcal{F}(\Sigma)$ is an isometry and the orbit map $\omega_T : \mathrm{Mod}(\Sigma) \to \mathcal{F}(\Sigma)$ restricts to a quasi-isometry $\omega_{T|} : \mathrm{Stab}(A) \to \mathcal{F}_A$. Moreover, the inclusion $\mathrm{Stab}(A) \hookrightarrow \mathrm{Mod}(\Sigma)$ is a quasi-isometric embedding.

After these results about the geometry of flip graphs and mapping class groups, we shift our focus to the quotient of the former by the latter, namely the geometry of modular flip graphs $\mathcal{MF}(\Sigma)$. In particular, we study their diameter and how it grows in function of the

3

topology of the base surface. Our main results are upper and lower bounds that have the same growth rates in terms of the number of punctures and genus. We summarize them in the following theorem.

Theorem 1.4. *There exist constants $L > 0$ and $U > 0$ such that if Σ is a surface of genus g with n labelled punctures then*

$$L\left(g\log(g+1) + n\log(n+1)\right) \leq \operatorname{diam}\left(\mathcal{MF}(\Sigma)\right) \leq U\left(g\log(g+1) + n\log(n+1)\right).$$

The above result is a combination of results (namely Theorems 4.3, 4.8 and Corollary 4.17) from which the constants L, U can be made explicit. When the punctures are not labeled, we obtain similar results and this time the growth rate is linear in n (Theorem 4.11 and Corollary 4.19).

We note that this result is a generalization of the result of Sleator, Tarjan and W. Thurston mentioned above about the diameters of modular flip graphs of punctured spheres and in fact our lower bounds are obtained using a counting argument and one of their results. Our result also provides a lower bound on the diameters of some slight refinements of the flip graph used in computational geometry and combinatorics. Indeed, it follows that the distance between any two simple triangulations (i.e. not containing multiple edges or loops) of a surface with labelled punctures grows at least like

$$n\log(n) + g\log(g).$$

This can be compared with results of Negami [21, 22] and Cortes et al. [10].

These growth rates are reminiscent of the growth rates of a type of combinatorial moduli space related to cubic graphs. More precisely, one can endow the set of isomorphism types of cubic graphs with m vertices with a metric where one counts the minimal number of *Whitehead moves* (or $\tilde{S}$-transformations in the language of [35]). We refer the reader to [8] or [28] for the definitions of these terms. With this metric, the diameter of this space is also of rough growth $m\log m$ (Cavendish [8], Cavendish-Parlier [9] and Rafi-Tao [28]).

Dual to a triangulation is a cubic graph and flips correspond to specific types of Whitehead moves. One might think in first instance that the two results are in fact the same, but one does not seem to imply the other. On the one hand, flipping only allows for certain moves so the result on flip graphs certainly seems stronger. However, given two cubic graphs with the same number of vertices, there is no guarantee that they are both the dual graph triangulations that lie in the same flip graph.

This article is organized as follows.

In the preliminary section, we provide detailed descriptions of the objects we study and some known results. We also prove a number of preliminary results including for instance a new algorithm to reach a stratum with distance bounded by the intersection number. In particular this provides a new proof that intersection number bounds the flip distance between two triangulations. We also provide a lower bound on distance in terms of intersection number. We conclude the section with two results that are somewhat parallel to the rest of the paper about the mapping class group and flip graphs. As far as we know, although both are known, our proofs are new. We provide these results to illustrate the point that flip graphs can be used to effectively study the mapping class group.

In the third section, we prove Theorem 1.1 stated above. We then provide two applications of this result. The first is about projections to strata and the second is to the large scale geometry of the mapping class group as discussed above.

The final section is about the diameters of modular flip graphs. We begin with upper bounds - first in terms of genus and then in terms of the number of punctures - and we end with the lower bounds.

Acknowledgements.

2. PRELIMINARIES

In this section we describe in some detail the objects we are interested in and introduce tools we use in the sequel. Most of the results we state are already known, although some of the proofs we provide are new (or at least we did not find them in the literature). In particular, at the end of this section we give two quick examples of results one can prove using flip graphs. Neither are essential in the sequel and are just provided for illustrative purposes.

2.1. Definitions and setup

We begin with the basic setup which starts with a topological orientable connected surface Σ and finite set of marked points on it. Unless specifically stated, Σ will be assumed to be triangulable. It is of finite type, has boundary which can consist of marked points, and boundary curves, and each boundary curve must have at least one marked point on it. We make the distinction between *labelled* and *unlabelled* marked points when we look at how homeomorphisms are allowed to act on Σ - this will made explicit in what follows.

Sometimes marked points that do not lie on a boundary curve will be referred to as *punctures*.

To such a Σ one can associate its *arc complex* $\mathcal{A}(\Sigma)$, a simplicial complex where vertices are isotopy classes of simple *arcs* based at the marked points of Σ. Simplices are spanned by *multiarcs* (unions of isotopy classes of arcs disjoint in their interior). We won't explicitly use this complex so we won't describe it in full detail, but we will be interested in the graph which is the 1-skeleton of the cellular complex dual to $A(\Sigma)$: the flip graph of Σ.

The flip graph $\mathcal{F}(\Sigma)$ can be described differently as follows. Vertices of this graph are maximal multiarcs so they decompose Σ into triangles. We refer to these multiarcs as *triangulations* (although they are not always proper triangulations in the usual sense - we apologize any confusion which incurs from this by now quite common terminology).

Two vertices of $\mathcal{F}(\Sigma)$ share an edge if they differ by a so-called *flip*. If a is an arc of a triangulation T which belongs to two triangles which form a quadrilateral, a *flip* is the operation which consists of replacing a by the other diagonal arc a' of the quadrilateral.

Figure 1: A local picture of a flip

Note that certain arcs are not *flippable* - this occurs exactly when an arc is contained in a punctured disc surrounded by another arc.

We denote by $\kappa = \kappa(\Sigma)$ the number of arcs in (any) triangulation of Σ and by $\tilde{\kappa} = \tilde{\kappa}(\Sigma)$ the number of triangles in the complement of any triangulation of Σ. Via an Euler characteristic argument, one obtains $\tilde{\kappa}(\Sigma) = 4g + 4b + 2s + p - 6$ and $\kappa(\Sigma) = 6g + 3b + 3s + p - 6$, where

Figure 2: The central arc is unflippable

g is the genus of Σ, s is the number of punctures, b is the number of boundary components and p is the number of points on the boundary curves of Σ.

Some flip graphs are finite - such as the flip graph of a polygon - but provided the underlying surface has enough topology, $\mathcal{F}(\Sigma)$ is infinite. A simple example of an infinite flip graph is given by the flip graph of a cylinder with a single marked point on each boundary curve. By the above formula, vertices of $\mathcal{F}(\Sigma)$ are of degree 2, and as it is both infinite and connected, it is isomorphic to the infinite line graph.

Associated to a multiarc A is a *stratum* $\mathcal{F}_A(\Sigma)$ which is the flip graph of triangulations of Σ which contain the multiarc A. We say that a stratum is *strongly convex* if any geodesic between two of its points is entirely contained in the stratum (sometimes this property is referred to as being *totally geodesic*).

Naturally, if A is not separating, $\mathcal{F}_A(\Sigma)$ is isomorphic to the flip graph of $\Sigma \setminus A$ (the surface *cut along* the multiarc A). There is something to be said here - we think of the result of the operation of cutting not as being the deletion of the arcs but by doubling the arcs and then separating them. For instance, if you cut a once punctured torus along an arc, the result is a cylinder with a marked point on each boundary curve. If A is separating, $\mathcal{F}_A(\Sigma)$ is isomorphic to the product of the product of the flip graphs of the connected components of $\Sigma \setminus A$. In the rest of the paper we will denote by $|A|$ the number of arcs in A.

As arcs are thought of as isotopy classes of arcs, the intersection $i(a, b)$ between arcs a and b is defined to be the minimum number of intersection points between their representatives. Generally we assume arcs and multiarcs to be realized in minimal position. If A and B are multiarcs, their *intersection number* is defined as

$$i(A, B) = \sum_{b \in B} \sum_{a \in A} i(a, b).$$

In terms of intersection, two triangulations T, T' are related by a flip if they satisfy

$$i(T, T') = 1.$$

The flip graph is known to enjoy a number of properties.

First of all, for any topological type of Σ, $\mathcal{F}(\Sigma)$ is a connected graph. There are several different proofs of this fact (see for instance Hatcher [15]). We will consider the edges of the flip graph of length 1 and we will endow the flip graph with its shortest path distance. The distance between two triangulations is then equal to the minimum number of flips required to pass from one to the other. In particular there is the following quantitative version relating distance and intersection number which can be deduced from an algorithm described by Mosher in [19] and Penner in [24].

Lemma 2.1. *For any triangulation $S, T \in \mathcal{F}(\Sigma)$ we have $d(S,T) \leq i(T,S)$.*

We will give an alternative proof of this lemma in Section 2.2.1.

The homeomorphisms of Σ considered here always fix pointwise the labelled points of Σ and setwise the unlabelled. Permutations of the unlabelled points are allowed. The *mapping class group* $\mathrm{Mod}(\Sigma)$ of Σ is the group of orientation preserving homeomorphisms of Σ up to isotopy. Isotopies here always fix pointwise the set of the marked points of Σ. The group $\mathrm{Mod}(\Sigma)$ acts simplicially by automorphisms on $\mathcal{F}(\Sigma)$. It is a result of Korkmaz and Papadopoulos [16] that except for some low complexity cases, the automorphism group of $\mathcal{F}(\Sigma)$ is exactly the *extended mapping class group* of Σ - i.e. the group of homeomorphisms up to isotopy (orientation reversing homeomorphisms are also allowed). Related to this result, is a result about subgraphs of flip graphs that are graph isomorphic to other flip graphs. Except for some complexity cases again, such subgraphs only arise in the natural way - as strata associated to a certain multiarc [2].

We will be interested in the geometry of the flip graph as a metric space. We recall a few notions of metric geometry that we will use later in the paper.

Let (X, d_X) and (Y, d_Y) be two metric spaces. A map

$$f : (X, d_X) \rightarrow (Y, d_Y)$$

is a (λ, ϵ)-*quasi-isometric embedding* if for all $x, x' \in X$ we have

$$\frac{1}{\lambda} d_X(x, x') - \epsilon \leq d_Y(f(x), f(x')) \leq \lambda d_X(x, x') + \epsilon.$$

We say that a quasi-isometric embedding f is a *quasi-isometry* if there exists $R \geq 0$ such that the image of f is R-dense in Y. Equivalently, f is a quasi-isometry if there exists a quasi-isometric embedding $g : Y \rightarrow X$ and a constant $K \geq 0$ such that for all $x \in X$ and for all $y \in Y$ we have $d_X(g \circ f(x), x) \leq K$ and $d_Y(f \circ g(y), y) \leq K$. We say that g is a *quasi-inverse* of f.

The following lemma is a classic result in geometric group theory.

8

Lemma 2.2 (Švarc-Milnor). *Let G be a group acting on a metric space (X, d_X) properly and cocompactly by isometries. Then G is finitely generated and for every $x \in X$ the orbit map*

$$\omega_X : G \to X$$
$$g \mapsto g(x)$$

is a quasi-isometry.

The mapping class group acts on the flip graph by isometries. The Švarc-Milnor lemma applies directly to the flip graph and the mapping class group, and of course is only interesting when $\mathcal{F}(\Sigma)$ is of infinite diameter.

Lemma 2.3. *For every triangulation $T \in \mathrm{Mod}(\Sigma)$ the orbit map*

$$\omega_T : \mathrm{Mod}(\Sigma) \to \mathcal{F}(\Sigma)$$
$$g \mapsto g(T)$$

is a quasi-isometry.

Proof. We will use the Švarc-Milnor Lemma. The action of $\mathrm{Mod}(\Sigma)$ on $\mathcal{F}(\Sigma)$ is cocompact since there is only a finite number of ways to glue $\tilde{\kappa}$ triangles to get a surface homeomorphic to Σ. For a triangulation $T \in \mathcal{F}(\Sigma)$, we denote by $\mathrm{Stab}(T)$ its stabilizer in $\mathrm{Mod}(\Sigma)$ (the group of mapping classes that fix T setwise). We will prove that for every T the stabilizer $\mathrm{Stab}(T)$ is finite and this suffices to prove that the action of $\mathrm{Mod}(\Sigma)$ on $\mathcal{F}(\Sigma)$ is proper. Indeed, every mapping class in $\mathrm{Stab}(T)$ induces a permutation of the arcs in T, and there is a short sequence of groups

$$1 \to \mathrm{Stab}(T) \to \mathcal{S}_\kappa$$

where $\mathcal{S}_\kappa$ is the symmetric group on κ elements. The sequence is exact since a mapping class that fixes every arc of a triangulation is the identity by the Alexander lemma (see for instance [11]).

The last assertion follows directly from the Švarc-Milnor lemma. $\qquad\square$

We define the *modular flip graph* $\mathcal{MF}(\Sigma)$ as the quotient of $\mathcal{F}(\Sigma)$ under the action of $\mathrm{Mod}(\Sigma)$. We remark that points in $\mathcal{MF}(\Sigma)$ are triangulations of Σ up to homeomorphisms. By the above lemma $\mathcal{MF}(\Sigma)$ is a connected finite graph that inherits a well-defined distance from $\mathcal{F}(\Sigma)$. We note that an orbit map in Lemma 2.3 is $\mathrm{diam}\,(\mathcal{MF}(\Sigma))$-dense in $\mathcal{F}(\Sigma)$ by the Švarc-Milnor Lemma. We will later investigate the diameter of $\mathcal{MF}(\Sigma)$.

The following result will be a helpful tool in our computation (see Section 4.2). This result was first proved by Sleator-Tarjan-Thuston [30] provided that n is large enough. Recently Pournin [26] provided a combinatorial proof and proved the lower bounds for all $n > 12$.

Theorem 2.4. *If Σ is a disk with $n > 12$ labelled points on the boundary then $\mathcal{F}(\Sigma)$ has diameter $2n - 10$.*

We finally note that an orbit map in Lemma 2.3 is $\mathrm{diam}(\mathcal{MF}(\Sigma))$-dense in $\mathcal{F}(\Sigma)$ by the Švarc-Milnor Lemma.

2.2. Intersection number and distances

2.2.1. An upper bound

In this section we describe an algorithm to get from a triangulation T to a stratum associated to a multiarc A. To do this, we prove that there exists an arc in T that intersects A maximally and such that its flip reduces the number of intersections with A. This provides an alternative proof of Lemma 2.1 above.

Definition 2.5. Let Δ be a triangle in $\Sigma \setminus T$ and A a multiarc. We say that Δ is *terminal* for A if there exists $a \in A$ such that Δ is the first or the last triangle crossed by a (see Figure 3 for an example).


Figure 3: A terminal triangle for a

Definition 2.6. Let T be a triangulation and A a multiarc. Let $t \in T$ be a flippable edge and T' the triangulation obtained by the flip of t. We say that flipping t is *convenient* if $i(T', A) < i(T, A)$ and flipping t is *neutral* if $i(T', A) = i(T, A)$.

Denote by t' the arc obtained by flipping t. Note that flipping t is convenient if and only if $i(t', A) < i(t, A)$. Similarly, flipping t is neutral if and only if $i(t', A) = i(t, A)$. Also note that flipping an arc may be neither convenient nor neutral.

10

Lemma 2.7. *Assume $i(h, A) > 0$. If $h \in T$ is an arc such that $i(h, A) = \max_{t \in T} i(t, A)$, then h is flippable in T.*

Proof. If h is not flippable then there exists an arc $h^\star \in T$ that surrounds h and bounds a once-punctured disk. Hence $i(h^\star, A) \geq 2i(h, A)$ in contradiction with our condition on h. □

Lemma 2.8. *Assume $i(T, A) > 0$. Let $h \in T$ be an arc such that $i(h, A) = \max_{t \in T} i(t, A)$. Let Q be the quadrilateral containing h as a diagonal. If Q contains a terminal triangle for A then flipping h is convenient.*

Proof. Let $a \in A$ be an arc that terminates on Q. We begin by observing that if h is not the first arc of T that a crosses from its terminal point, h is not maximal. Indeed, if $\tilde{h} \in T$ is the first arc crossed, then any arc in A that crosses h is forced to cross $\tilde{h}$ and it has

$$i(\tilde{h}, A) \geq i(h, A) + 1,$$

in contradiction with the maximality of h (see Figure 4).

Figure 4: If h is not the first arc of T crossed by a, it cannot be maximal

We now prove that if h is the first (or the last) edge crossed by an arc in A (as in Figure 5), then flipping h is convenient. Let h' be the arc obtained from flipping h. Let us now compute $i(h', A) - i(h, A)$.

Up to possible symmetries, we may assume that A contains at least an arc that crosses z and h and terminates in the top vertex of Q. Up to isotopy, only three configurations of A and Q are possible, and these are described in Figure 5. We will use the following notation:

- ϵ is the number of arcs in $A \cap Q$ that terminate in the top vertex of Q, crossing z and h. Under our assumption, $\epsilon \geq 1$.

- ϵ' is the number of those that terminate in the top vertex of Q, crossing w and h;

- ϵ'' is the number of those that terminate in the bottom vertex of Q, crossing h and y;

- α is the number of arcs in $A \cap Q$ that wrap around the left endpoint of h crossing z, h and x;

- β is the number of arcs in $A \cap Q$ that wrap around the right endpoint of h crossing w, h and y;

- γ is the number of arcs in $A \cap Q$ that wrap around the bottom vertex of Q, crossing z, h and w;

- η is the number of arcs that cross z, h and y.

Figure 5: The three possible configurations of arcs

We remark that in Figure 5 - (1) we have $\epsilon' = \epsilon'' = 0$, in Figure 5 - (2) we have $\gamma = 0$, in Figure 5 - (3) we have $\eta = 0$. Moreover, in each configuration in Figure 5 the following holds:

$$i(h, A) = \alpha + \beta + \epsilon + \eta + \epsilon' + \epsilon''$$
$$i(h', A) = \gamma + \eta$$
$$i(z, A) = \alpha + \eta + \gamma + \epsilon = \alpha + \epsilon + i(h', A)$$

By definition of h, we have $i(z, A) \leq i(h, A)$. It follows:

$$\alpha + \epsilon + i(h', A) \leq i(h, A)$$
$$i(h', A) - i(h, A) \leq -\alpha - \epsilon \leq -1.$$

We conclude that flipping h is convenient. $\square$

Lemma 2.9. *Assume $i(T, A) > 0$. Let $h \in T$ be an arc such that $i(h, A) = \max_{t \in T} i(t, A)$. Let Q be the quadrilateral containing h as a diagonal, and assume that Q does not contain a*

12

terminal triangle for A. Then flipping h is neutral if and only if, according the notation of Figure 6,

$$i(h, A) = i(y, A) = i(z, A).$$

Moreover, if $i(h, A) \neq i(y, A)$ or $i(h, A) \neq i(z, A)$ then flipping h is convenient.

Proof. We will compute $i(h', A) - i(h, A)$. Since Q is not terminal for A, A and Q look like in Figure 6, up to isotopy.


Figure 6: Configuration of arcs

Denote by δ the number of arcs in $A \cap Q$ that wrap around the top corner of Q. In the notation of the proof of Lemma 2.8 we have:

$$i(h, A) = \alpha + \eta + \beta$$
$$i(z, A) = \alpha + \eta + \gamma$$
$$i(y, A) = \eta + \beta + \delta$$
$$i(h', A) = \gamma + \delta + \eta$$
$$i(h', A) - i(h, A) = (\gamma - \beta) + (\delta - \alpha).$$

By hypothesis on h, we have $i(z, A) \leq i(h, A)$, so $\gamma - \beta \leq 0$. Similarly, $i(y, A) \leq i(h, A)$, so $\delta - \alpha \leq 0$. It follows that $i(h', A) - i(h, A) = 0$ if and only if $\gamma = \beta$ and $\delta = \alpha$, that is, if and only if $i(z, A) = i(h, A)$ and $i(y, A) = i(h, A)$. We remark that in all the other cases $i(h', A) - i(h, A) < 0$, and flipping h is convenient. $\square$

Lemma 2.10. *If $i(T, A) > 0$, then there exists an arc $h \in T$ such that $i(h, A) = \max_{t \in T} i(t, A)$ and flipping h is convenient.*

Proof. We will describe a procedure to find h. First pick an arc m such that $i(m, A) = \max_{t \in T} i(t, A)$. By Lemma 2.7 m is flippable. If flipping m is convenient, we set $h = m$ and we are done. Otherwise, by Lemma 2.9 the quadrilateral containing m as a diagonal

contains two more edges, y and z, such that $i(y, A) = i(z, A) = i(m, A) = \max_{t \in T} i(t, A)$. Now set $m = y$ or $m = z$, and repeat this procedure. Lemma 2.8 ensures that the algorithm terminates. Indeed, if the quadrilateral containing m as a diagonal also contains a terminal triangle for A then flipping m is convenient. $\qquad\square$

This now allows to describe a path from T to $\mathcal{F}_A$.

Theorem 2.11. *Let T be a triangulation and A a multiarc. There is a sequence of triangulations $T = T_0 \to \ldots \to T_{i+1} \to \ldots$ iteratively constructed as follows:*

1. *If $i(T_i, A) > 0$, choose $h_i \in T_i$ as in Lemma 2.10. Denote by T_{i+1} the triangulation obtained by flipping h_i.*

2. *If $i(T_i, A) = 0$, terminate.*

Any such sequence constructed this way has at most $i(T, A)$ elements, and if T_n is the terminal triangulation then $T_n \in \mathcal{F}_A$.

Moreover, the above sequence satisfies $\max_{t \in T_{i+1}} i(t, A) \leq \max_{t \in T_i} i(t, A)$ for every $i = 0, \ldots, n$.

Proof. By Lemma 2.10 flipping h_i is convenient, so

$$i(T_{i+1}, A) \leq i(T_i, A) - 1.$$

After at most $i(T, A)$ steps, we have $i(T_n, A) = 0$, that is, $T_n \in \mathcal{F}_A$. $\qquad\square$

From this the following corollary is immediate.

Corollary 2.12. *For every triangulation $T \in \mathcal{F}(\Sigma)$ and for every multiarc A, we have*

$$d(T, \mathcal{F}_A) \leq i(T, A).$$

We remark that if S is a triangulation, $\mathcal{F}_S = S$ and the path described in Theorem 2.11 is a path joining T and S. As such we also have the following corollary.

Corollary 2.13. *The flip graph $\mathcal{F}(\Sigma)$ is connected, and for any two triangulations $T, S \in \mathcal{F}(\Sigma)$ we have $d(T, S) \leq i(T, S)$.*

Our construction also enjoys the following properties.

Corollary 2.14. *The path from T to $\mathcal{F}_A$ described in Theorem 2.11 has the following properties:*

1. *If there exists $a \in A$ such that $a \in T_i$ then $a \in T_j$ for all $j \geq i$;*

2. *If $t \in T_i$ is such that $i(t, A) = 0$ then $t \in T_j$ for all $j \geq i$.*

We will see later that all the geodesic paths between a triangulation and the stratum of a multiarc also have these properties. We use this corollary to deduce the following.

Corollary 2.15. *For every multiarc A, the stratum $\mathcal{F}_A$ is arcwise connected.*

Proof. Let $S, T \in \mathcal{F}_A$ be two triangulations, and consider the path from T to $\mathcal{F}_S = S$ described in Theorem 2.11. We remark that for all $a \in A$ we have $a \in T$ and $i(a, S) = 0$. By Corollary 2.14, for all $a \in A$ we have $a \in T_i$, so $T_i \in \mathcal{F}_A$ for all $i = 0, \ldots, n$. We conclude that the path described is contained in $\mathcal{F}_A$. □

2.2.2. A lower bound on distances

As a complement to our upper bound on distance in terms of intersection number, we now show how intersection also provides a lower bound on distance. We begin with the following observation.

Lemma 2.16. *Let T and T' be two triangulations in Σ that differ by one flip, and let A be a multiarc with $|A|$ components. Then we have:*

$$i(T', A) \geq 2 \cdot \max_{t \in T} i(t, A) - 2 \cdot |A|.$$

Proof. Assume that T and T' differ by a flip on t, we set $T' = T \setminus \{t\} \cup \{t'\}$. Let $h \in T$ be an arc such that $i(h, A) = \max_{t \in T} i(t, A)$. We have:

$$
\begin{aligned}
i(T', A) = i(T, A) - i(t, A) + i(t', A) &\geq i(T, A) - i(t, A) \\
&\geq i(T, A) - i(h, A)
\end{aligned}
\tag{1}
$$

By Lemma 2.7 the arc h is flippable. In the notation of Lemma 2.8, we have

$$
\begin{aligned}
i(h, A) &= \alpha + \eta + \beta + \delta + \epsilon + \epsilon' + \epsilon'' \\
i(x, A) &= \alpha + \delta \\
i(y, A) &= \eta + \beta + \delta + \epsilon'' \\
i(w, A) &= \beta + \epsilon' \\
i(z, A) &= \alpha + \eta + \epsilon
\end{aligned}
\tag{2}
$$

Remark that $\epsilon + \epsilon' + \epsilon'''$ is the total number of arcs that terminates in the quadrilateral containing h, therefore $\epsilon + \epsilon' + \epsilon''' \leq 2|A|$. Combining Equations 2, we have:

$$
\begin{aligned}
i(T', A) &\geq i(T, A) - i(h, A) \\
&\geq (i(x, A) + i(y, A) + i(w, A) + i(z, A) + i(h, A)) - i(h, A) \\
&= (\alpha + \delta) + (\eta + \beta + \delta + \epsilon'') + (\beta + \epsilon') + (\alpha + \eta + \epsilon) \\
&= 2i(h, A) - (\epsilon + \epsilon' + \epsilon'') \\
&\geq 2i(h, A) - 2 \cdot |A|
\end{aligned}
\tag{3}
$$

$\square$

Theorem 2.17. *If T is a triangulation and A is a multiarc such that $\max_{t \in T} i(t, A) \geq 2|A|$, then*

$$
d(T, \mathcal{F}_A) \geq \left\lfloor \frac{\log(i(T, A)) - \log(2|A| - 1)}{\log(\kappa)} \right\rfloor - 2.
$$

Proof. Let $h \in T$ be an arc such that $i(h, A) = \max_{t \in T} i(t, A)$. We have:

$$
i(h, A) \geq \frac{i(T, A)}{\kappa}.
\tag{4}
$$

Let T' be a triangulation that differs from T by one flip. By Lemma 2.16 the following holds when $i(h, A) \geq 2|A|$:

$$
\begin{aligned}
i(T', A) &\geq 2i(h, A) - 2 \cdot |A| \\
&\geq i(h, A) \\
&\geq \frac{i(T, A)}{\kappa}.
\end{aligned}
\tag{5}
$$

We note that the case $i(h, A) \leq 2|A|$ is not very interesting because in this case T is not too far from $\mathcal{F}_A$. Indeed, by Lemma 2.12 we have: $d(T, \mathcal{F}_A) \leq 2\kappa \cdot i(h, A) \leq 2\kappa \cdot |A|$.

Let $d = d(T, \mathcal{F}_A)$, and let $T = T_0 \to \ldots \to T_d \in \mathcal{F}_A$ be a geodesic path from T to $\mathcal{F}_A$. Let $m \leq d$ be the smallest integer such that $\max_{t \in T_{m+1}} i(t, A) \leq 2|A| - 1$ and for every $j \leq m$ we have $\max_{t \in T_j} i(t, A) \geq 2|A|$. We have:

$$
i(T_{m+1}, A) \leq (2|A| - 1) \cdot \kappa.
\tag{6}
$$

By Lemma 5 and the above remark, we also have:

$$
i(T_{m+1}, A) \geq \frac{i(T_m, A)}{\kappa} \geq \ldots \geq \frac{i(T_0, A)}{\kappa^{m+1}}
\tag{7}
$$

16

We have the following inequality that we solve for m:

$$(2|A| - 1) \cdot \kappa \geq \frac{i(T_0, A)}{\kappa^{m+1}}$$

$$\kappa^{m+2} \geq \frac{i(T, A)}{2|A| - 1} \tag{8}$$

$$m \geq \frac{\log i(T, A) - \log(2|A| - 1)}{\log(\kappa)} - 2.$$

We conclude:

$$d(T, \mathcal{F}_A) = d \geq m \geq \frac{\log i(T, A) - \log(2|A| - 1)}{\log(\kappa)} - 2.$$

$\square$

We remark that if $\max_{t \in T} i(t, A) \leq 2|A| - 1$ then by Lemma 2.12

$$d(T, \mathcal{F}_A) \leq i(T, A) \leq (2|A| - 1) \cdot \kappa.$$

Corollary 2.18. *If T and S are two triangulations, then*

$$d(T, S) \geq \left\lfloor \frac{\log(i(T, S))}{\log(\kappa)} \right\rfloor - 4.$$

2.3. Examples of the relationship between flip graphs and the mapping class group

In this section, we provide two examples of how one can use the flip graph to study the mapping class group. They are completely independent from the rest of the paper but are provided to illustrate the variety of ways in which the quasi-isometry between the two objects can be used.

2.3.1. Mapping tori and pseudo-Anosov homeomorphisms

The following proposition follows from a standard construction in 3-dimensional topology known as the *layered triangulation* of the mapping torus of a pseudo-Anosov homeomorphism.

Proposition 2.19. *For every pseudo-Anosov $\phi \in \mathrm{Mod}(\Sigma)$ and for every triangulation $T \in \mathcal{F}(\Sigma)$ we have:*

$$d(T, \phi(T)) \geq \frac{\mathrm{vol}(M_\phi)}{\pi},$$

where $\mathrm{vol}(M_\phi)$ is the volume of the mapping torus M_ϕ of ϕ.

Proof. Consider a geodesic path of flips $T \to \ldots \to \phi(T)$. The number of hyperbolic tetrahedra in the layered triangulation of M_ϕ associated to this path is equal to $d(T, \phi(T))$. For details on the layered triangulation of a mapping torus, we refer to [1]. $\qquad\square$

Our second application is the following.

Corollary 2.20. *For every pseudo-Anosov $\phi \in \mathrm{Mod}(\Sigma)$ the cyclic subgroup $\langle \phi \rangle$ is undistorted in* $\mathrm{Mod}(\Sigma)$.

Proof. We first prove that for every triangulation T and for every ϕ we have:

$$\frac{n \cdot \mathrm{vol}(M_\phi)}{\pi} \leq d(T, \phi^n(T)) \leq n \cdot d(T, \phi(T)).$$

By Lemma 2.3 this suffices to prove the corollary.

The upper bound follows immediately by the triangle inequality. For the lower bound we use Proposition 2.19. We remark that since M_{ϕ^n} is a finite cover of degree n of M_ϕ then

$$\mathrm{vol}(M_{\phi^n}) = n \cdot \mathrm{vol}(M_\phi).$$

It follows

$$d(T, \phi^n(T)) \geq \frac{\mathrm{vol}(M_{\phi^n})}{\pi} = \frac{n \cdot \mathrm{vol}(M_\phi)}{\pi}.$$

$\qquad\square$

2.3.2. The cone construction

Fix a complete finite-area hyperbolic metric M on Σ and a homeomorphism φ between M and Σ. Let $P = \{p_1, \ldots, p_n\}$ be the set of punctures of Σ. It is a classical result of Birman and Series that the set of all simple geodesics on M is nowhere dense on M. We choose a point on the complement of the closure of all the simple geodesics of M and consider its image by φ on Σ. We denote this point p_{n+1} (on both Σ and M). We now set $P' = P \cup \{p_{n+1}\}$ and let Σ' be the punctured surface Σ with an extra marked point at p_{n+1}. This construction is known as *puncturing* (see [27]). Let T be a triangulation of Σ, denote by $\mathcal{G}_M(T)$ the unique M-geodesic representative in its isotopy class (it is an ideal triangulation as the marked points become punctures). Then p_{n+1} is contained in a unique triangle of $\Sigma \setminus \mathcal{G}_M(T)$. We then *cone* the triangle in p_{n+1}: by this we mean add arcs between p_{n+1} and the three vertices of the triangle to obtain a triangulation of Σ' that we denote by $\widehat{T}$. We will also refer to the arcs going to p_{n+1} as the *cone* on p_{n+1}. In the following we will denote by d the flip distance on $\mathcal{F}(\Sigma)$ and by d' the flip distance on $\mathcal{F}(\Sigma')$.

Lemma 2.21. *The cone map*

$$\text{cone}_M : \mathcal{F}(\Sigma) \to \mathcal{F}(\Sigma')$$
$$T \mapsto \widehat{T}$$

is well-defined and 2-Lipschitz.

Proof. If T' is a triangulation of Σ isotopic to T, then $\mathcal{G}_M(T') = \mathcal{G}_M(T)$, so $\widehat{T}' = \widehat{T}$. Figure 7 shows that if two triangulations differ by one flip, their images by the cone map differ by at most 2 flips.


Figure 7: A flip and a corresponding two flip move in the coned triangulation

$\square$

We will now prove that cone_M is a quasi-isometric embedding. Fix a triangulation $H \in \mathcal{F}(\Sigma)$. Denote by $\omega_H : \text{Mod}(\Sigma) \to \mathcal{F}(\Sigma)$ the orbit map of H under $\text{Mod}(\Sigma)$ as in Lemma 2.3. Similarly, denote by $\omega_{\widehat{H}} : \text{Mod}(\Sigma') \to \mathcal{F}(\Sigma')$ the orbit map of $\widehat{H}$ under $\text{Mod}(\Sigma')$. Recall that $\text{Mod}(\Sigma)$ and $\text{Mod}(\Sigma')$ are related by the Birman exact sequence, where the map $f : \text{Mod}(\Sigma') \to \text{Mod}(\Sigma)$ is the *forgetful map*:

$$1 \to \pi_1(\Sigma, p_{n+1}) \to \text{Mod}(\Sigma') \xrightarrow{f} \text{Mod}(\Sigma) \to 1$$

Lemma 2.22. *Let $f : \text{Mod}(\Sigma') \to \text{Mod}(\Sigma)$ be the forgetful map. Let $\omega_{\widehat{H}}^{-1} : \mathcal{F}(\Sigma') \to \text{Mod}(\Sigma')$ be a quasi-inverse of $\omega_{\widehat{H}}$. The following map is quasi-Lipschitz.*

$$F : \mathcal{F}(\Sigma') \to \mathcal{F}(\Sigma)$$
$$T \mapsto \omega_H \circ f \circ \omega_{\widehat{H}}^{-1}(T)$$

19

For all $\psi' \in \mathrm{Mod}(\Sigma')$ we have $F(\psi'(\widehat{H})) = f(\psi')(H)$.

Proof. It is immediate to see that f is 1-Lipschitz with respect to the Humphreis generators of $\mathrm{Mod}(\Sigma)$. The assertion follows by composition with the two quasi-isometries. $\qquad\square$

We remark that the quasi-Lipschitz constants of F depend on the diameter of $\mathcal{MF}(\Sigma)$ and a choice of generators for $\mathrm{Mod}(\Sigma)$ and $\mathrm{Mod}(\Sigma')$.

Lemma 2.23. *For every $\psi \in \mathrm{Mod}(\Sigma)$ there exists $\phi' \in f^{-1}(\psi) \subset \mathrm{Mod}(\Sigma')$ such that*

$$d'(\widehat{\psi H}, \phi'(\widehat{H})) \le 2\tilde{\kappa}$$

Proof. Fix $\psi \in \mathrm{Mod}(\Sigma)$ and choose $\psi' \in \mathrm{Mod}(\Sigma')$ such that $f(\psi') = \psi$, that is, ψ and ψ' are homeomorphisms of Σ isotopic rel P. Let us first compare $\widehat{\psi(H)}$ and $\psi'(\widehat{H})$.

To construct $\psi'(\widehat{H})$ we proceed as follows. Set $\Sigma \setminus \mathcal{G}_M(H) = \bigcup \Delta_i$ where Δ_i is a triangle. We assume $p_{n+1} \in \Delta_1$, so that $\widehat{H}$ is obtained by H coning Δ_1 and $\phi'(\widehat{H})$ is obtained coning $\psi'(\Delta_1)$. To construct $\widehat{\psi(H)}$ we proceed as follows. Set $\Sigma \setminus \mathcal{G}_M(\psi(H)) = \bigcup \Delta_i'$, where Δ_i' is a triangle. Since $\psi' \in f^{-1}(\psi)$, we can assume (up to reordering) that Δ_i' is isotopic to $\psi'(\Delta_i)$ relative to P. We have two cases:

1. $p_{n+1} \in \Delta_1'$;

2. $p_{n+1} \notin \Delta_1'$.

In case (1), we can glue the homeomorphisms $\Delta_i' \to \psi'(\Delta_i)$ in order to construct a homeomorphism $\theta : \Sigma \to \Sigma$ that also fixes p_{n+1}. By construction, θ is an element of $\mathrm{Mod}(\Sigma')$ isotopic to the identity rel P, that is, θ belongs to the kernel of the forgetful map f, and we obtain $\theta(\widehat{\psi H}) = \psi'(\widehat{H})$. Consider the mapping class $\phi' = \theta^{-1} \circ \psi' \in \mathrm{Mod}(\Sigma')$, by construction

$$f(\theta^{-1} \circ \psi') = f(\psi') = \psi \text{ and } \widehat{\psi H} = \phi'(\widehat{H}),$$

and we are done.

In case (2), assume $p_{n+1} \in \Delta_j'$ with $j \ne 1$. We will now see that using at most $2\tilde{\kappa}$ flips we can move the cone on p_{n+1} inside a triangle isotopic to Δ_1'. More precisely, a sequence of two flips as in Figure 8 moves the cone in a triangle adjacent to Δ_j'. Note that this sequence of flips does not change the isotopy class relative to P of the arcs not connected to p_{n+1}. The final triangulation T_1 we obtain has the following properties:

- T_1 has a cone on p_{n+1};

- T_1 agrees with $\widehat{\psi H}$ outside the quadrilateral in Figure 8;

- the arcs of $\widehat{\psi H}$ and T_1 that are not connected to p_{n+1} are pairwise isotopic relative P.

If the triangle of T_1 containing the cone on p_{n+1} is isotopic to Δ'_1 relative to P, then we can proceed as in case (1). Indeed, we construct an homeomorphism $\theta : \Sigma' \to \Sigma'$ such that $\theta(T_1) = \psi'(\widehat{H})$ and θ is isotopic to the identity relative to P. We then set $\phi' = \theta^{-1} \circ \psi'$, and we have $T_1 = \phi'(\widehat{H})$. Otherwise, if the triangle of T_1 containing p_{n+1} is not isotopic to Δ'_1, we keep on performing sequences of flips like in Figure 8 in order to move the cone on p_{n+1}. After at most $\tilde{\kappa}$ sequences of flips, we get to a triangulation $T_{\tilde{\kappa}}$ whose cone on p_{n+1} lies inside a triangle isotopic to Δ'_1. Arguing as above, we get a homeomorphism $\theta : \Sigma' \to \Sigma'$, isotopic relative to P to the identity, such that $\theta(T_{\tilde{\kappa}}) = \psi'(\widehat{H})$. We then set $\phi' = \theta^{-1} \circ \psi'$, we have $T_{\tilde{\kappa}} = \phi'(\widehat{H})$. We conclude as follows:

$$d'(\widehat{\psi(H)}, \phi'(\widehat{H})) \leq d'(\widehat{\psi(H)}, T_{\tilde{\kappa}}) + d'(T_{\tilde{\kappa}}, \phi'(\widehat{H}))$$
$$\leq 2 \cdot \tilde{\kappa}.$$

Figure 8: Passing from $\widehat{\psi H}$ to T_1

$\square$

Theorem 2.24. $\mathrm{cone}_M : \mathcal{F}(\Sigma) \to \mathcal{F}(\Sigma')$ *is a quasi-isometric embedding.*

Proof. By Lemma 2.21, cone_M is 2-Lipschitz. We will now prove that there exists some universal constants $A', B' > 0$ such that for any $S, T \in \mathcal{F}(\Sigma)$ we have $d(S, T) \leq A' \cdot d'(\widehat{S}, \widehat{T}) + B'$. Set $R = \mathrm{diam} \mathcal{MF}(\Sigma)$. Since the orbit of H under $\mathrm{Mod}(\Sigma)$ is R-dense, we can find $\psi, \phi \in \mathrm{Mod}(\Sigma)$ such that

$$d(S, \psi H) \leq R \quad \text{and} \quad d(T, \phi H) \leq R \tag{9}$$

By Lemma 2.21 we have:

$$d'(\widehat{S}, \widehat{\psi H}) \leq 2R \quad \text{and} \quad d'(\widehat{T}, \widehat{\phi H}) \leq 2R \tag{10}$$

21

Finally fix $\psi', \phi' \in \mathrm{Mod}(\Sigma')$ as in Lemma 2.23. We use the notation $a \prec b$ to mean $a \leq kb + h$ for constants k and h.

With this notation:

$$
\begin{aligned}
d(S,T) &\leq d(\psi H, \phi H) + 2R && \text{by 9}\\
&= d(F(\psi'(\widehat{H})), F(\phi'(\widehat{H}))) + 2R \\
&\prec d'(\psi'(\widehat{H}), \phi'(\widehat{H})) && \text{by Remark 2.22}\\
&\prec d'(\psi'(\widehat{H}), \widehat{\psi H}) + d'(\widehat{\psi H}, \widehat{\phi H}) + d'(\phi'(\widehat{H}), \widehat{\phi H}) \\
&\prec d'(\widehat{\psi H}, \widehat{\phi H}) + 4 \cdot \tilde{\kappa} && \text{by Lemma 2.23}\\
&\prec d'(\widehat{\psi H}, \widehat{S}) + d'(\widehat{S}, \widehat{T}) + d'(\widehat{\phi H}, \widehat{T}) \\
&\prec d'(\widehat{S}, \widehat{T}) + 4R && \text{by Equation 10}
\end{aligned}
$$

For F is (A, B) quasi-Lipschitz, and if we keep track of the constants in the above inequalities, we have:

$$d(S,T) \leq A' \cdot d'(\widehat{S}, \widehat{T}) + B',$$

where $A' = A$ and $B' = 4AR + 4A\tilde{\kappa} + B + 2R$. $\qquad \square$

The following result was already proved by Mosher [20] using a different method, and stated by Rafi-Schleimer [27] using the marking graph as a large scale model for $\mathrm{Mod}(\Sigma)$.

Corollary 2.25. *There is a quasi-isometric embedding* $\mathrm{Mod}(\Sigma) \hookrightarrow \mathrm{Mod}(\Sigma')$.

Proof. Consider the following commutative diagram:

$$
\begin{array}{ccc}
\mathcal{F}(\Sigma) & \xrightarrow{\ \mathrm{cone}_M\ } & \mathcal{F}(\Sigma') \\
\omega_{H|} \uparrow & & \uparrow \omega_{\widehat{H}} \\
\mathrm{Mod}(\Sigma) & \longrightarrow & \mathrm{Mod}(\Sigma')
\end{array}
$$

By Lemma 2.3 both ω_H and $\omega_{\widehat{H}}$ are quasi-isometries. The assertion then follows from the above theorem. $\qquad \square$

3. CONVEXITY OF STRATA AND APPLICATIONS

As we saw previously, for any multiarc A the stratum $\mathcal{F}_A$ is connected. We denote by d_A the shortest path distance on $\mathcal{F}_A$. In this section we prove that the natural inclusion $(\mathcal{F}_A, d_A) \hookrightarrow (\mathcal{F}, d)$ is an isometric embedding. Furthermore, we prove that $\mathcal{F}_A$ is strongly convex in $\mathcal{F}(\Sigma)$. The main ingredient in our proof is a 1-Lipschitz retraction of $\mathcal{F}(\Sigma)$ on $\mathcal{F}_A$.

3.1. The projection theorem

Let a and t be two arcs. Choose an orientation on a, denote by a^+ the oriented arc, and let $\text{push}_{a^+}(t)$ be the multiarc defined as follows:

- if $i(a,t) = 0$ then $\text{push}_{a^+}(t) = t$;

- if $i(a,t) \neq 0$ then $\text{push}_{a^+}(t)$ is the multiarc obtained by "combing" t following the orientation of a as in Figure 9. Each arc in $\text{push}_{a^+}(t)$ (provided $i(a,t) \neq 0$) has at least one endpoint that coincides with the final endpoint of a.


Figure 9: Combing along an oriented arc

The following lemma follows immediately by the above construction.

Lemma 3.1. *If s and t are arcs, then $i(\text{push}_{a^+}(s), \text{push}_{a^+}(t)) \leq i(s,t)$.*

If $T = (t_1, \ldots, t_\kappa)$ is a triangulation of Σ, we denote by $\text{push}_{a^+}(T)$ the multiarc obtained collecting the isotopy classes of all the arcs $\text{push}_{a^+}(t_i)$: $\text{push}_{a^+}(T) = [\text{push}_{a^+}(t_1), \ldots, \text{push}_{a^+}(t_\kappa)]$. We remark that the set $\{\text{push}_{a^+}(t_i)\}$ may contain isotopic arcs.

Lemma 3.2. *The map*

$$\pi_{a^+} : \mathcal{F}(\Sigma) \to \mathcal{F}_a$$
$$T \mapsto (\text{push}_{a^+}(T), a)$$

is a 1-Lipschitz retraction on $(\mathcal{F}_a, d_a)$.

Proof. We first prove that $\pi_{a^+}(T)$ is also a triangulation of Σ. If a is one of the arcs in T, the assertion is trivial. Suppose $a \cap T \neq \emptyset$.

We consider a parametrization of $a : [0,1] \to \Sigma$ following the orientation of a^+. We suppose that a intersects T minimally and all intersections are transversal so we denote

$$\tau_0 = 0 < \tau_1 < \ldots < \tau_{N+1} = 1$$

23

the values of τ for which $a(\tau) \in T$. Note that $N = i(a, T)$.

For each $\tau' \in [0, 1]$, we consider the following decomposition $D_{\tau'}$ of Σ constructed as follows. To begin, $D_{\tau'}$ contains all arcs of T that do not cross $a|_{\tau=0}^{\tau'}$, contains all vertices of T and has one extra vertex $a(\tau')$. Furthermore, it also contains the arc $a|_{\tau=0}^{\tau'}$. We add arcs iteratively as follows. For each τ_i, $i = 1, \ldots, N - 2$, the point $a(\tau_i)$ will cut a preexisting arc, say b_i, into two subarcs b_i' and b_i''. We add these to the decomposition and they continue to belong to the decomposition for $\tau' > \tau_i$ by concatenating them with the arc $a|_{\tau=\tau_i}^{\tau'}$ in the obvious way. At parameter τ' we denote the resulting arcs $b_i'(\tau')$ and $b_i''(\tau')$. $D_{\tau'}$ is the union of all these arcs up to isotopy fixing the vertices (so any isotopy class is only counted once).

We want to show that $D_{\tau_{N+1}}$ is a triangulation of Σ with the same vertex set as T. Before showing this we claim that for $0 \le i < N + 1$, D_{τ_i} is a set of arcs decomposing Σ into triangles and into one quadrilateral which is simply a triangle with an additional vertex $a(\tau_i)$.

We prove our claim by analyzing the decomposition as τ varies. The key point is that the decomposition only changes for the values τ_i.

For $i = 1$ as all we have added is an arc and a point that splits the first triangle traversed by a into two triangles. As we have added a vertex in $a(\tau_1)$, the following triangle of T traversed by a is now a quadrilateral (see Figure 10).


Figure 10: When the second triangle becomes a quadrilateral

Now suppose by induction that at parameter τ_i with $i < N$ the decomposition D_{τ_i} is as claimed and we now analyze $D_{\tau_{i+1}}$.

To obtain $D_{\tau_{i+1}}$ from D_{τ_i} we have a continuous family D_τ with $\tau \in [\tau_i, \tau_{i+1}]$. Note that $a|_{\tau=\tau_i}^{\tau_{i+1}}$ is a simple path crossing the only quadrilateral of $\Sigma \setminus D_{\tau_i}$. While $\tau' \in [\tau_i, \tau_{i+1}[$, D_τ is

24

just obtained by pushing $a(\tau)$ along the path $a|_{\tau=\tau_i}^{\tau'}$ and thus (up to homeomorphism) is a carbon copy of D_{τ_i}.


Figure 11: The general step

We need to analyse what happens at $\tau = \tau_{i+1}$. Two of the arcs of the quadrilateral become one and as such the quadrilateral collapses to a triangle. More precisely, the point $a(\tau_{i+1})$ lies on an arc of D_{τ_i} so divides this arc into two arcs in $D_{\tau_{i+1}}$; adding this vertex turns the "next" triangle into a quadrilateral. This proves the general step. The above process is illustrated in Figure 11.

What remains to be seen is the final step, when $\tau \in [\tau_N, \tau_{N+1}]$. This final step is very similar to what happens before with the notable difference that the point $a(\tau_{N+1})$ was already a vertex of the decompositions D_τ. So instead of splitting a previous arc into two parts, the quadrilateral containing $a(\tau)$ for $\tau \in [\tau_N, \tau_{N+1}[$ collapses completely leaving only triangles in $D_{\tau_{N+1}}$ (see Figure 12).

Figure 12: The final step

This concludes the proof that $\pi_{a^+}(T)$ is a triangulation.

It is straightforward to see that π_{a^+} is a retraction. In fact, the restriction of π_{a^+} to $\mathcal{F}_a$ is the identity and π_{a^+} is onto by construction.

Let us now prove that π_{a^+} is 1-Lipschitz. Recall that $T_1, T_2 \in \mathcal{F}(\Sigma)$ differ by a flip if and only if $i(T_1, T_2) = 1$. By Lemma 3.1 we have $i(\pi_{a^+}(T_1), \pi_{a^+}(T_2)) \leq 1$. We deduce that

either $\pi_{a^+}(T_1)$ and $\pi_{a^+}(T_2)$ also differ by a flip or they coincide. Let T and T' be two vertices in $\mathcal{F}(\Sigma)$ and $\gamma : T = T_0 \ldots T_m = T'$ is a geodesic path in $\mathcal{F}(\Sigma)$ joining them. By the above argument, $\pi_{a^+}(\gamma) : \pi_{a^+}(T) \ldots \pi_{a^+}(T')$ is a path in $\mathcal{F}_a$ of length at most m, so $d(\pi_{a^+}(T), \pi_{a^+}(T')) \leq d(T, T')$ and π_{a^+} is 1-Lipschitz. $\qquad\square$

Lemma 3.3. *Let A be a multiarc and T be a triangulation. If there exists $t \in T$ such that $i(t, A) = 0$ then every geodesic path from T to $\mathcal{F}_A$ is contained in $\mathcal{F}_t$.*

Proof. Let $\gamma : T = T_0 \ldots T_n$ be a shortest path from T_0 to $\mathcal{F}_A$. We shall prove that for all i, $T_i \in \mathcal{F}_t$. We begin by choosing an orientation on t. Observe that $p_{t^+}(T_n) \in \mathcal{F}_A$ and $T_0 = \pi_{t^+}(T_0)$ by construction, so $\pi_{t^+}(\gamma)$ is also a path from T_0 to $\mathcal{F}_A$. We now argue by contradiction. Let $i \geq 0$ the smallest integer such that $t \in T_i$ and $t \notin T_{i+1}$ (that is, the arc t is flipped). Necessarily we have $i(t, T_{i+1}) = 1$ and by construction

$$\pi_{t^+}(T_i) = \pi_{t^+}(T_{i+1}) = T_i$$

so the length of $\pi_{t^+}(\gamma)$ is at most $n - 1$. This implies that $\pi_{t^+}(\gamma)$ is shorter than γ, in contradiction with the assumption that γ is geodesic. $\qquad\square$

Theorem 3.4. *For every arc a, the stratum $\mathcal{F}_a$ is strongly convex.*

Proof. Let T_0 and T_m be two vertices in $\mathcal{F}_a$ and let $\gamma : T_0 \ldots T_m$ be a geodesic path in $\mathcal{F}(\Sigma)$ joining them. By Lemma 3.2 $\pi_{a^+}(\gamma)$ is a path in $\mathcal{F}_a$ with endpoints T_0 and T_m and we have $d_a(T_0, T_m) \leq m = d(T_0, T_m)$. It follows that the inclusion $\mathcal{F}_a \hookrightarrow \mathcal{F}(\Sigma)$ is an isometric embedding. The strong convexity of $\mathcal{F}_a$ follows from Lemma 3.3 with $A = a$ and $t = a$: for all $i = 0, \ldots, m$ we have $T_i \in \mathcal{F}_a$. $\qquad\square$

Theorem 3.5. *Let $A^\sigma = (a_1^+, \ldots, a_m^+)$ be a multiarc whose m components are enumerated and oriented. The map $\pi_{A^\sigma} = \pi_{a_m^+} \circ \ldots \circ \pi_{a_1^+} : \mathcal{F}(\Sigma) \to \mathcal{F}_A$ is well-defined and a 1-Lipschitz retraction.*

Proof. Since the arcs in A are all disjoint, by Lemma 3.3 we have

$$\pi_{A^\sigma}(\mathcal{F}(\Sigma)) = \mathcal{F}_{a_1} \cap \ldots \cap \mathcal{F}_{a_m} = \mathcal{F}_A.$$

By Lemma 3.2 the map π_{A^σ} is 1-Lipschitz and a retraction. $\qquad\square$

We remark that the map π_{A^σ} does depend on the choice of the orientation and enumeration of the arcs in A. We will study this dependence later.

Theorem 3.6. *For every multiarc A, the stratum $\mathcal{F}_A$ is strongly convex.*

26

Proof. This is a direct corollary of Theorem 3.4. Note that $\mathcal{F}_A = \bigcap_{a \in A} \mathcal{F}_a$ and the intersection of strongly convex subspaces is strongly convex.

$\square$

3.2. Applications

We now focus on some applications of the above results and in particular of Theorem 3.6.

3.2.1. Projections and distances

We begin by looking at some immediate consequences on distances and projection distances to strata.

The following proposition is essentially the definition of distance on $\mathcal{F}_A$ combined with Theorem 3.6.

Proposition 3.7. *Assume that A is a multiarc such that $\Sigma \setminus A = \bigcup_{i=1}^{h} \Sigma_i$ where Σ_i is a connected surface with boundary. Denote by d_i the distance on $\mathcal{F}(\Sigma_i)$. For every $T \in \mathcal{F}_A$ denote by T_i the triangulation of Σ_i induced by T. Then the map*

$$\mathcal{F}_A \longrightarrow \mathcal{F}(\Sigma_1) \times \ldots \times \mathcal{F}(\Sigma_h)$$
$$T \mapsto (T_1, \ldots, T_h)$$

is an isometry between $(\mathcal{F}_A, d)$ and $(\mathcal{F}(\Sigma_1) \times \ldots \times \mathcal{F}(\Sigma_h), d_1 + \ldots + d_h)$

Proof. By definition of $\mathcal{F}_A$, the map is an isometry from $(\mathcal{F}_A, d_A)$. By Theorem 3.5 $d = d_A$ and the assertion follows. $\square$

Proposition 3.8. *Let A be a multiarc. For every choice σ of enumeration and orientation of the arcs in A, we have: $d(\pi_{A^\sigma}(T), \pi_{A^\sigma}(S)) \leq i(T, S)$.*

Proof. It is a straightforward application of Lemmas 2.12 and 3.1. $\square$

Proposition 3.9. *Let A be a multiarc. For every choice σ of enumeration and orientation of the arcs in A, we have $d(T, \mathcal{F}_A) \leq d(T, \pi_{A^\sigma}(T)) \leq 2 \cdot d(T, \mathcal{F}_A)$.*

Proof. Let S be a triangulation in $\mathcal{F}_A$ at minimal distance from T, so that $d(T, S) = d(T, \mathcal{F}_A)$.

By Theorem 3.5 $\pi_{A^\sigma}(S) = S$, it follows:

$$\begin{aligned}
d(T, \pi_{A^\sigma}(T)) &\leq d(T, S) + d(S, \pi_{A^\sigma}(T)) \\
&\leq d(T, S) + d(\pi_{A^\sigma}(S), \pi_{A^\sigma}(T)) \\
&\leq d(T, S) + d(S, T) \\
&= 2d(T, \mathcal{F}_A).
\end{aligned}$$

$\square$

Corollary 3.10. *Let A be a multiarc. For every choice σ, ϵ of enumeration and orientations of the arcs in A, we have $d(\pi_{A^\sigma}(T), \pi_{A^\epsilon}(T)) \leq d(T, \mathcal{F}_A)$.*

Proof. It follows immediately by Proposition 3.9. $\square$

The next consequence will use a result by Aramayona, Koberda and the second author about simplicial maps between flip graphs. To state the result we require the following notation: we say that a surface Σ is exceptional if it is an essential subsurface of (and possibly equal to) a torus with at most two marked points, or a sphere with at most four marked points. In [2], it is proved that, for surfaces Σ, Σ' with Σ non-exceptional, all injective simplicial maps

$$\phi : \mathcal{F}(\Sigma) \to \mathcal{F}(\Sigma')$$

come from embeddings $\Sigma \to \Sigma'$ (that is Σ is homeomorphic to a subsurface of Σ'). Note that it's obvious that you can construct simplicial maps this way; what's more surprising is that this is, provided your base surface is complicated enough, the only way such maps appear. Together with Theorem 3.6, the following is then immediate.

Corollary 3.11. *Suppose Σ is non-exceptional, and let $\mathcal{F}(\Sigma) \to \mathcal{F}(\Sigma')$ be an injective simplicial map. Then $\mathcal{F}(\Sigma)$ is strongly convex inside of $\mathcal{F}(\Sigma')$.*

3.2.2. On the large scale geometry of the mapping class group

We now turn our attention to the large scale geometry of the mapping class group.

Lemma 3.12. *Let A be a multiarc and $\mathrm{Stab}(A)$ be the subgroup of $\mathrm{Mod}(\Sigma)$ that fixes the isotopy class of each arc in A. Then $\mathrm{Stab}(A)$ has a finite index subgroup isomorphic to $\mathrm{Mod}(\Sigma \setminus A)$.*

Proof. Assume that A has m connected components. Fix an orientation on each arc of A. It is immediate to see that the subgroup $\mathrm{Stab}^+(A) < \mathrm{Stab}(A)$ consisting of the mapping

classes that also fix the orientation of every arc in A is isomorphic to $\mathrm{Mod}(\Sigma \setminus A)$, that is, the subgroup of the surface obtained cutting Σ along A. The assertion follows from the exactness of the short sequence:

$$1 \to \mathrm{Stab}^+(A) \to \mathrm{Stab}(A) \to \mathbb{Z}_2^m \to 1.$$

□

We can now prove the following.

Theorem 3.13. *For every vertex $T \in \mathcal{F}_A$, there is a commutative diagram:*

$$
\begin{array}{ccc}
\mathcal{F}_A & \hookrightarrow & \mathcal{F}(\Sigma) \\
\omega_{T|} \uparrow & & \uparrow \omega_T \\
\mathrm{Stab}(A) & \hookrightarrow & \mathrm{Mod}(\Sigma)
\end{array}
$$

where the inclusion $\mathcal{F}_A \hookrightarrow \mathcal{F}(\Sigma)$ is an isometry and the orbit map $\omega_T : \mathrm{Mod}(\Sigma) \to \mathcal{F}(\Sigma)$ restricts to a quasi-isometry $\omega_{T|} : \mathrm{Stab}(A) \to \mathcal{F}_A$. Moreover, the inclusion $\mathrm{Stab}(A) \hookrightarrow \mathrm{Mod}(\Sigma)$ is a quasi-isometric embedding.

Proof. The inclusion $\mathcal{F}_A \hookrightarrow \mathcal{F}(\Sigma)$ is an isometry by Theorem 3.6. By Proposition 3.7 $\mathcal{F}_A$ is isomorphic and isometric to $\mathcal{F}(\Sigma \setminus A)$. Since the action of $\mathrm{Mod}(\Sigma \setminus A)$ on $\mathcal{F}(\Sigma \setminus A)$ is cocompact, so it is the action of $\mathrm{Stab}(A)$ on $\mathcal{F}_A$ by Lemma 3.12. By the Švarc-Milnor lemma the orbit map $\mathrm{Mod}(\Sigma) \ni \psi \mapsto \psi T \in \mathcal{F}_A$ is a quasi-isometry. By composition the inclusion $\mathrm{Stab}(A) \hookrightarrow \mathrm{Mod}(\Sigma)$ is a quasi-isometric embedding. □

4. THE DIAMETERS OF THE MODULAR FLIP GRAPHS

The goal of this section is to prove upper and lower bounds on the diameters of modular flip graphs in terms of the topology of the surface (namely Theorem 1.4 from the introduction).

Let Σ be a surface of genus g with n labelled points. We assume $g \geq 1$ and $n \geq 2$ (for the case $n = 1$ see Theorem 4.7 and Remark 4.6, for the case $g = 0$ see Theorem 4.13). The case where the points are unlabelled is slightly easier and it will also be treated separately - see Remark 4.10.

We begin with a general observation which allows us to break bounds on $\mathcal{MF}(\Sigma)$ into different parts. The idea is to work with the punctures on one side and genus on the other. To do this we consider triangulations that contain an arc which separates the genus from the punctures: more precisely an arc a which forms a loop based in a puncture and such

that $\Sigma \setminus a = \Omega \cup \Gamma$ where Ω is a disk with $n-1$ punctures and one labelled point on the boundary and Γ is of genus g with a boundary component with a single marked point. Such a loop we call *puncture separating*.

For any choice of puncture on Σ, it is clear that (infinitely many) such loops based in this point exist but up to homeomorphism there is only one such loop (see Figure 13).


Figure 13: A puncture cutting loop

From this we can make the following observation: any two triangulations which are distinct up homeomorphism and both contain a puncture separating loop must be either distinct on Γ or Ω. As such:

$$\operatorname{card}\left(\mathcal{MF}(\Sigma)\right) > \operatorname{card}(\mathcal{MF}(\Gamma))\operatorname{card}(\mathcal{MF}(\Omega)). \tag{11}$$

We will use this for our lower bounds in Section 4.3.

For our upper bounds the following lemma will allow us to introduce a puncture separating arc in a minimal amount of flips.

Lemma 4.1. *For any $T \in \mathcal{MF}(\Sigma)$ and any marked point p of Σ, there exists a puncture separating loop a based in p with*

$$i(a, T) \leq 2(\kappa - n + 1)$$

Proof. We think of T as a graph embedded on Σ and consider a spanning tree of this graph. A regular neighborhood of this tree is a simple closed curve γ which satisfies

$$i(\gamma, T) \leq 2(\kappa - (n-1))$$

as it intersects only half edges that do not belong to the tree and the tree has $n-1$ edges - see Figure 14. (The above inequality is in fact an equality but it is the inequality that we need.)

From γ and given a marked point p, we shall construct an arc as follows: as γ surrounds all punctures, it must pass through a triangle that has p as a vertex. We consider a simple arc c in the triangle between γ and p. Choosing an orientation on c and γ, the concatenation

Figure 14: The curve Γ

of $c\gamma c^{-1}$ gives an isotopy class of arc which is the arc a we are looking for. Notice that by construction it intersects T in at most as many points as γ and we have

$$i(a, T) \leq 2(\kappa - n + 1)$$

as desired.

Figure 15: The curve γ, the path c and the arc a

$\square$

Using this lemma and the upper bound on flip distance in terms of intersection number, we can establish the following.

Lemma 4.2. *For Σ, Ω and Γ as above:*

$$\mathrm{diam}\left(\mathcal{MF}(\Sigma)\right) \leq \mathrm{diam}\left(\mathcal{MF}(\Omega)\right) + \mathrm{diam}\left(\mathcal{MF}(\Gamma)\right) + 2(\kappa - n + 1).$$

The above inequality will allow us to treat the upper bounds by treating $\mathrm{diam}\left(\mathcal{MF}(\Gamma)\right)$ and $\mathrm{diam}\left(\mathcal{MF}(\Omega)\right)$ separately. We begin with the former.

4.1. Upper bounds in terms of genus

As above, Γ is a genus $g \geq 1$ surface with a single boundary curve and a single marked point on the boundary. Our goal here is to show the following result.

Theorem 4.3. *The diameter of the modular flip graph of Γ satisfies*

$$\text{diam}\left(\mathcal{MF}(\Gamma)\right) < Ag \log(g+1)$$

where A can be taken to be 1000.

Before proving the theorem we'll need two topological lemmas.

Lemma 4.4. *Let T be a triangulation of Λ, a genus $g \geq 1$ surface with a single boundary curve and all k marked points on the boundary. Then there exists $a \in T$ such that $\Lambda \setminus a$ is connected and of genus $g - 1$.*

Proof. Observe that for an arc $a \in T$, $\Lambda \setminus a$ being connected and of genus $g - 1$ is equivalent (cutting along a separating arc does not reduce genus). We now claim that T always contains a non-separating arc. As $\Lambda \setminus T$ is a collection of triangles, it is of genus 0. Now as $g(\Lambda) \geq 1$, one of the arcs of T must be non-separating, otherwise $\Lambda \setminus T$ would still have positive genus. □

Lemma 4.5. *Let T be a triangulation of Λ, a genus $g \geq 0$ surface with two boundary curves, both with marked points, and all marked points on the boundary. Then there exists $a \in T$ such that $\Lambda \setminus a$ has only one boundary component.*

Proof. All marked points are on the boundary so it is impossible to triangulate Λ without a triangle that has vertices on both boundary components. To see this we can argue by contradiction. If this is not the case, then we can split the triangles into two non-empty groups depending on whether they have all of vertices on one or the other boundary curve. But as the surface is connected, there must be a triangle of the first group which shares an arc with a triangle of the second. Thus, they must also share vertices, a contradiction. □

We now proceed to the proof of Theorem 4.3.

Proof of Theorem 4.3. Let T be any triangulation of Γ. Denote by a_0 the arc that forms the boundary of Γ.

The first step will be to divide the surface along an arc that has equal genus (or close to equal) on both parts. By Lemma 4.4, there is an arc $a_1 \in T$ such that $\Gamma \setminus a_1$ is of genus $g - 1$. The resulting surface $\Gamma^1 := \Gamma \setminus a_1$ now has two boundary components, one consisting of two arcs and the other of a single arc. Now by Lemma 4.5, there exists $a_2 \in T$ such that $\Gamma_2 := \Gamma^1 \setminus a_2$ has a single boundary curve consisting of 5 arcs. In short, we found two arcs of T such that cutting along these arcs produces a surface of genus $g - 1$ with a single

32

boundary component with 4 more arcs than the original surface Γ. We can iterate the above process at total of $\lfloor \frac{g}{2} \rfloor$ times to obtain a collection of $2\lfloor \frac{g}{2} \rfloor$ arcs such that cutting along these arcs results in a genus $g - \lfloor \frac{g}{2} \rfloor$ surface $\overline{\Gamma}$ with a single boundary curve formed by $1 + 4\lfloor \frac{g}{2} \rfloor$ arcs. One of these is a_0.

Denote p_0 and p'_0 the two vertices of a_0 on $\overline{\Gamma}$. We denote b the unique loop based in p_0 homotopic to the boundary of $\overline{\Gamma}$ and b' the arc from p'_0 to p_0 which forms a triangle with a_0 and b (see Figure 16).


Figure 16: $\overline{\Gamma}$

Both b and b' have a nice property: they don't intersect T too much. More precisely, as there are parallel to the boundary of $\overline{\Gamma}$ which is formed by arcs of T, they intersect each of the the remaining arcs at most twice. Thus

$$i(x, T) \leq 2\left(\kappa(\Gamma) - 2\left\lfloor \frac{g}{2} \right\rfloor\right), \quad x = b, b'.$$

Now $\kappa(\Gamma) = 6g - 2$ so we can deduce that

$$i(b, T) + i(b', T) \leq 20g - 4.$$

Now using the upper bound on the distance to a stratum in function of intersection number, we can introduce the arcs b and b' in at most $20g - 4$ flips.

The reason one might want to do this is that these arcs separate the surface into three canonical surfaces: a triangle containing a_0 and two surfaces with a single boundary curve and of genus $\lfloor \frac{g}{2} \rfloor$ and $g - \lfloor \frac{g}{2} \rfloor$. As such, up to homeomorphism, the pair of arcs b and b' are unique (see Figure 17).

With this in hand, we will prove the bound by induction. We begin by checking the result for $g = 1$. Here we need to check that the diameter is at most $1000 \log(2) > 693 > 5$. But there are at most 5 different possible triangulations. Indeed such a triangulated surface is obtained by pasting four sides of a triangulated 5-gon together. There are $C_3 = 5$ different possible triangulations of the 5-gon and only one to paste together the 5-gon to get a one holed torus.

33

Figure 17: The arcs b and b'

We now suppose $g(\Gamma) \geq 2$.

Given two different triangulations S and T in $\mathcal{MF}(\Gamma)$, we flip both triangulations to obtain triangulations S' and T' with arcs as above. These triangulations now both belong to a stratum of $\mathcal{MF}(\Gamma_{b,b'})$ where b and b' are as above. We denote Γ_1 and Γ_2 the two non-triangular surfaces in $\Gamma \setminus \{b \cup b'\}$. Denote (for $k = 1, 2$) S'_k, resp. T'_k, the restrictions of S', resp. T', to Γ_k. We shall now flip S'_k and T'_k inside $\mathcal{MF}(\Gamma_k)$ for $k = 1, 2$. Once the triangulations coincide on both Γ_1 and Γ_2. they will coincide on Γ.

By induction for $k = 1, 2$:

$$d(S'_k, T'_k) \leq \operatorname{diam}(\mathcal{MF}(\Gamma_k)) \leq A \frac{g+1}{2} \log \frac{g+3}{2}.$$

By induction (here we take into account that g can be odd in the bound of $g - \lfloor \frac{g}{2} \rfloor$):

$$d(S'_1, T'_1) \leq \operatorname{diam}(\mathcal{MF}(\Gamma_1)) \leq A \frac{g}{2} \log \frac{g+2}{2}$$

and

$$d(S'_2, T'_2) \leq \operatorname{diam}(\mathcal{MF}(\Gamma_2)) \leq A \frac{g+1}{2} \log \frac{g+3}{2}.$$

Putting this all together:

$$
\begin{aligned}
d(S, T) &\leq d(S, S') + d(T, T') + d(S'_1, T'_1) + d(S'_2, T'_2) \\
&\leq 40g - 8 + A \frac{g}{2} \log \frac{g+2}{2} + A \frac{g+1}{2} \log \frac{g+3}{2} \\
&\leq A g \log(g+1).
\end{aligned}
$$

34

The last inequality can be checked via a computation using $A = 1000$ and $g \geq 2$. □

Remark 4.6. In light of Lemma 4.2, in the above theorem we've treated the case where the boundary of Γ is a loop. The above proof however applies verbatim to the case where Γ has a single puncture and no other boundary. The resulting theorem is the following.

Theorem 4.7. *If Γ is a surface with genus g and one puncture, then the diameter of the modular flip graph of Γ satisfies*

$$\mathrm{diam}(\mathcal{MF}(\Gamma)) < Ag\log(g+1)$$

where A can be taken to be 1000*.*

4.2. Upper bounds in terms of number of punctures

We now focus our attention on the flip graph of Ω, a disk with $n-1$ interior punctures and one marked point on the unique boundary curve of Ω. Our goal is to prove the following upper bound which is very similar to the upper bound for Γ.

Theorem 4.8. *If Ω has $n-1$ labelled punctures then the diameter of the modular flip graph of Ω satisfies*

$$\mathrm{diam}\left(\mathcal{MF}(\Omega)\right) < An\log(n+1)$$

where A can be taken equal to 400*.*

Before proceeding to the proof, we state a preliminary lemma.

Lemma 4.9. *Let T be a triangulation of Λ, a $m \geq 1$-punctured disk with $k \geq 1$ marked points on the boundary. Then T contains an arc a between an interior puncture and marked point on the boundary.*

Proof. If not, then a simple curve parallel to boundary does not intersect T and hence $\Lambda \setminus T$ contains an embedded annulus. □

With that observation in hand, we now proceed to the proof of Theorem 4.8.

Proof of Theorem 4.8. Let T be a triangulation of Ω where we suppose that $n \geq 2$ (if $n = 1$ then there the flip graph has a single triangulation). We denote the boundary arc of Ω a_0, the boundary marked point p_0 and the remaining punctures p_j, $j = 1, \dots, n$. Our goal will be to flip our triangulation to a canonical triangulation and argue by induction on the distance to this canonical triangulation. The upper bound on distance between arbitrary

35

Figure 18: The canonical triangulation

triangulations is then at most twice this distance. Our canonical triangulation S is the following.

The triangulated surface is formed of layers. Each layer except for the last one is a cylinder with two boundary arcs a_{k-1} and a_k with punctures $p_{k-1} \in a_{k-1}$ and $p_k \in a_k$ for $k = 1, \ldots, n-1$. The cylinders all contain a single interior arc from the triangulation as in the figure. The last layer is a disk with boundary a_{n-1} and puncture $p_{n-1} \in a_{n-1}$ and interior puncture p_n. There is an arc in the triangulation between p_{n-1} and p_n.

To reach this triangulation from T we proceed as follows. We begin by finding arcs that will divide the surface into punctured disks with the same (or close to the same) number of punctures in each disk. By Lemma 4.9, there is an arc $c \in T$ such that $\Omega \setminus c$ is a disk with 3 boundary arcs: the arc a_0 and the two copies of c. We reiterate the above process $\lfloor \frac{n}{2} \rfloor$ times cutting along $\lfloor \frac{n}{2} \rfloor$ arcs to obtain a disk $\overline{\Omega}$ with $1 + 2\lfloor \frac{n}{2} \rfloor$ boundary arcs. On this boundary, a_0 joins two vertices: p_0' and another, say p_0'', both copies of p_0. Consider the arc b which forms a loop in p_0' parallel to the boundary of $\overline{\Omega}$. Similarly, consider b' which forms a triangle with a_0 and b: b' is an arc between p_0' and p_0'' which runs parallel to the boundary of $\overline{\Omega}$.

36

Both b and b' have a nice property: they don't intersect T too much. More precisely, as there are parallel to the boundary of $\overline{\Omega}$ which is formed by arcs of T, they intersect each of the the remaining arcs at most twice. Thus

$$i(x, T) \leq 2(\kappa(\Omega) - \left\lfloor \frac{n}{2} \right\rfloor), \quad x = b, b'.$$

Now $\kappa(\Omega) = 3n - 2$ and $-2\left\lfloor \frac{n}{2} \right\rfloor \leq -n + 1$ so we can deduce that

$$i(b, T) + i(b', T) \leq 10n - 10.$$

Now using the upper bound on the distance to a stratum in function of intersection number, we can introduce the arcs b and b' in at most $10n - 10$ flips.

The resulting triangulation now has an arc surrounding $\left\lfloor \frac{n}{2} \right\rfloor$ punctures, another surrounding $n - \left\lfloor \frac{n}{2} \right\rfloor$ punctures and the two arcs form a triangle with a_0 (see Figure 19). We now argue by
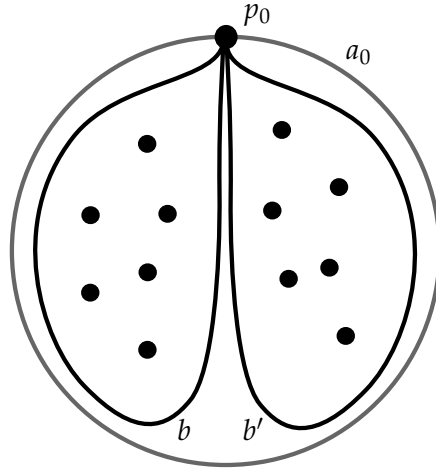

Figure 19: The arcs b and b'

induction on the two subsurfaces Ω_b and $\Omega_{b'}$ surrounded by b and b' to flip them towards their canonical triangulations. We have no control over which punctures are found in Ω_b and $\Omega_{b'}$ but the punctures do inherit an order from Ω and their canonical triangulations are meant with respect to that order. The number of flips inside each of the two subsurfaces, by induction, is at most

$$A(\left\lfloor \frac{n}{2} \right\rfloor + 1) \log(\left\lfloor \frac{n}{2} \right\rfloor + 2).$$

Denote the resulting triangulation T'. We now need to merge the two subtriangulations of T' to obtain the canonical one. To do this we proceed by steps where each step in the process will be to add a cylinder bounded by arcs a_{k-1} and a_k with punctures p_{k-1} and p_k.

We begin with the first step. Puncture p_1 is either found in Ω_b (the lefthand subsurface) or in $\Omega_{b'}$ (the righthand subsurfaces). In either event it shares an arc with p_0 as both sub triangulations are canonical (and thus ordered). If p_1 on the left, we flip as in Figure 20, and similarly if p_1 is on the right. As illustrated in the figures, the process takes 6 flips. We've constructed the first ring of the canonical triangulation. This ring surrounds a divided subsurface and we are in the same situation as above, where p_1 and a_1 play the part of p_0 and a_0 and with one less interior puncture. We can iterate the process a total of $n-1$ times (the last step is automatic) and arguing by induction we have reached the canonical triangulation in $6(n-1)$ steps from T'.


Figure 20: Merging step

Putting this all together we have that for any $T \in \mathcal{MF}(\Omega)$

$$d(T,S) \leq 10n - 10 + A(\left\lfloor \frac{n}{2} \right\rfloor + 1) \log(\left\lfloor \frac{n}{2} \right\rfloor + 2) + 6(n-1).$$

Arguing like in the genus case (see the proof of Theorem 4.3) we obtain that

$$d(T,S) \leq An \log(n+1).$$

This shows that any two triangulations are at distance at most $2An \log(n+1)$ where A can be taken equal to 200. □

Remark 4.10. The upper bound on $\mathcal{MF}(\Omega)$ is much easier if the punctures are unlabelled. Indeed, given a vertex p, if a triangulation contains arcs that are not incident to p, you can always find a flip that increases the incidence in p. Let S and T be two triangulations. After

38

at most $4\kappa - 2n$ valence-increasing flips, both T and S look like in Figure 21, that is, up to homeomorphisms they differ only in the shaded area. The shaded area can be thought as a triangulated n-agon. By Theorem 2.4 T and S differ by at most $4\kappa - 2n + 2n = 4\kappa$ flips.


Figure 21: The shaded area is triangulated.

Theorem 4.11. *If Ω has $n - 1$ unlabelled punctures then the diameter of the modular flip graph of Ω satisfies*

$$\mathrm{diam}(\mathcal{MF}(\Omega)) < An$$

where A can be taken equal to 12.

Remark 4.12. The above proof however applies verbatim to the case where Ω is a punctured sphere (in this case the arcs b and b' in Figure 19 coincide.)We thus have the following.

Theorem 4.13. *If Ω is a sphere with n labelled punctures then*

$$\mathrm{diam}(\mathcal{MF}(\Omega)) < An \log(n)$$

where A can be taken equal to 410.

Theorem 4.14. *If Ω is a sphere with n unlabelled punctures then*

$$\mathrm{diam}(\mathcal{MF}(\Omega)) < An$$

where A can be taken equal to 22.

4.3. Lower bounds via counting arguments

We now focus on lower bounds. They will essentially follow from a theorem of Sleator, Tarjan and Thurston [31] and a counting argument.

We begin with the following general lemma which follows from a theorem on grammars on graphs [31].

Lemma 4.15. *Let Λ be a surface with n punctures and $\mathcal{MF}(\Lambda)$ its modular flip graph. Then for a fixed triangulation $T_\mu \in \mathcal{MF}(\Lambda)$ we have:*

$$\mathrm{card}\{T \in \mathcal{MF}(\Lambda) \mid d(T, T_\mu) \leq m\} \leq 4^{10m} 4^{\tilde{\kappa}(\Lambda)}.$$

39

Proof. This is a direct consequence of Theorem 2.3 of [31] and the discussion in Section 5 in [31]. For any triangulation T one can construct its dual graph $G(T)$ (see Figure 22).


Figure 22: The graph dual to a triangulation

The graph $G(T)$ is a trivalent graph that has exactly $\tilde{\kappa}(\Lambda)$ vertices. We consider the three half-edges incident to a vertex labelled by the integers $1, 2, 3$ in clockwise order. Changing T by one flip is equivalent to evolve $G(T)$ according the grammar in Figure 23.


Figure 23: The grammar of a flip

This grammar has two productions: one for doing the flip and the other for preparing the half-edges labels to allow the flip. Indeed, one flip on T corresponds to perform at most 5 productions on $G(T)$: two to prepare the half-edge labels on the first vertex, two to prepare the half-edge labels on the second vertex, and one for the flip. It follows that the number of triangulations that can be obtained from T_μ in at most m flips is bounded above by the number of graphs that can be derived by $G(T_\mu)$ with at most $10m$ productions. The latter is bounded above by $4^{10m}4^{\kappa(\tilde{\Lambda})}$ by a straightforward application of Theorem 2.3 [31] to the grammar we described. The same proof works verbatim for T_ν. □

Remark 4.16. Setting $m = \operatorname{diam}(\mathcal{MF}(\Sigma))$ in the lemma above, and then solving for m using Inequality 4, one obtains the following result:

Corollary 4.17. *Let Σ be a surface of genus g with n marked points, Γ be a surface of genus g with one boundary component and exactly one marked point on it, and Ω be a disk with $n - 1$ interior*

punctures. We have:

$$\text{diam}\left(\mathcal{MF}(\Sigma)\right) > \frac{\log(\text{card}(\mathcal{MF}(\Gamma))) + \log(\text{card}(\mathcal{MF}(\Omega))) - \tilde{\kappa}(\Sigma)\log(4)}{10\log(4)}$$

We now count vertices of our combinatorial moduli spaces.

Lemma 4.18. *Let Γ be a surface of genus $g \geq 2$ with a single boundary loop and one marked point on the boundary. Let Ω be a disk with a single boundary component with a marked point on the boundary and with $n - 1$ interior labelled points. Then*

$$\text{card}\{\mathcal{MF}(\Gamma)\} \quad \geq \quad \frac{g-1}{2}(g-1)! \tag{12}$$

$$\text{card}\{\mathcal{MF}(\Omega)\} \quad \geq \quad C_{n-2}(n-1)! \tag{13}$$

where C_k is the k-th Catalan number.

Proof. We begin with Inequality 12. For a given triangulation $T \in \mathcal{MF}(\Gamma)$, if we collapse the triangle which contains the boundary arc by cutting the triangle and pasting the two loose arcs together, we obtain a triangulated surface of genus g with a single marked point. If you perform this on two triangulations $S, T \in \mathcal{MF}(\Gamma)$ and obtain different triangulations up to homeomorphism, then the triangulations we necessarily different to begin with. As such, there are at least as many triangulations in $\mathcal{MF}(\Gamma)$ then triangulations of a genus g surface with a single marked point. It is a result of Penner [23] that there are at least

$$\frac{g-1}{2}(g-1)!$$

such triangulations and so the inequality is proved.

For Inequality 13 we argue as follows. Denote by p_0 the marked point on the boundary curve and a_0 the boundary loop. We begin by considering only triangulations where each interior puncture is surrounded by a single loop based at p_0 (see Figure 21).

For two triangulations to be the same, they must coincide on the exterior of these loops. Cutting along the loops, one obtains an n-gon with one privileged side a_0. As such, we are in the classical case of counting triangulations of a polygon with an order on the sides and there are C_{n-2} such triangulations. Any permutation of the vertex labelling gives a different polygon and thus we obtain the stronger lower bound

$$(n-1)!\,C_{n-2}.$$

$\square$

From this we obtain the following lower bound.

Corollary 4.19. *Let Σ be a surface with n labelled punctures. We have*

$$\operatorname{diam}\left(\mathcal{MF}(\Sigma)\right) > B\left(n\log(n+1) + g\log(g+1)\right),$$

where B can be taken equal to $2 \cdot 10^{-5}$.

Proof. We will use the following inequalities:

1. $\log(C_n) > n$;

2. $\log n! > n\log(n) - n$.

Assume $n \geq 3$ and $g \geq 3$. From Lemma 4.18 we get:

$$\log(\operatorname{card}(\mathcal{MF}(\Gamma))) > \log(g-1)! > (g-1)\log(g-1) - g \qquad (14)$$
$$\log(\operatorname{card}(\mathcal{MF}(\Omega))) > \log(n-1)! > (n-1)\log(n-1) - n \qquad (15)$$

Assume that the punctures of Σ are labelled. Plugging in the inequality in Corollary 4.17 we have:

$$\operatorname{diam}\left(\mathcal{MF}(\Sigma)\right) > \frac{(g-1)\log(g-1) - g + (n-1)\log(n-1) - n - \tilde{\kappa}(\Sigma)\log(4)}{10\log(4)}$$
$$> \frac{(g-1)\log(g-1) - g + (n-1)\log(n-1) - n - (4g + 2n - 6)\log(4)}{10\log(4)}$$
$$> \frac{(g-1)\log(g-1) - (4\log(4)+1)g}{10\log(4)} + \frac{(n-1)\log(n-1) - (2\log(4)+1)n}{10\log(4)}$$
$$> B\left(g\log(g+1) + n\log(n+1)\right)$$

where B can be taken to be $2 \cdot 10^{-5}$ and for $g \geq 705$ and $n \geq 50$. It is immediate to verify that

$$\operatorname{diam}\left(\mathcal{MF}(\Sigma)\right) > B\left(g\log(g+1) + n\log(n+1)\right)$$

also holds in the remaining cases ($g \leq 704$ or $n \leq 49$). □

We note that we can improve the constant B by conditioning g and n (giving them both lower bounds) but our principle interest is in the order of growth.

We obtain a similar result on lower bounds for unlabeled marked points.

Corollary 4.20. *If Σ has $n \geq 511$ unlabelled punctures and is of genus g then*

$$\operatorname{diam}\left(\mathcal{MF}(\Sigma)\right) > B(g\log(g+1) + n)$$

where B can be taken to be 10^{-3}.

Proof. The graph grammar described in Lemma 4.15 can be refined (see [31] for details) so that

$$\operatorname{card}\mathcal{MF}(\Sigma) \leq 3^{\tilde{\kappa}(\Sigma)}8^m.$$

We have:

$$m \geq \frac{\log(\operatorname{card}(\mathcal{MF}(\Sigma))) - \tilde{\kappa}(\Sigma)\log(3)}{\log(8)}.$$

Let $\tilde{\Omega}$ be a disk with a single boundary component with a marked point on the boundary and with $n-1$ interior unlabelled points. As in Lemma 4.15 we have

$$\operatorname{card}\mathcal{MF}(\Sigma) \geq \operatorname{card}\mathcal{MF}(\Gamma)\operatorname{card}\mathcal{MF}(\tilde{\Omega}).$$

Now we use a result of Brown [7] that provides lower bounds on the cardinality of $\mathcal{MF}(\tilde{\Omega})$:

$$\operatorname{card}\mathcal{MF}(\tilde{\Omega}) > \frac{2(4n-7)!}{(n-1)!(3n-4)!}.$$

An explicit computation shows that, for $n \geq 511$, the following holds:

$$\operatorname{card}\mathcal{MF}(\tilde{\Omega}) > (9.1)^n > 3^{2n}.$$

From this we can conclude that

$$\begin{aligned}
\operatorname{diam}\left(\mathcal{MF}(\Sigma)\right) &> \frac{(g-1)\log(g-1) - g - \tilde{\kappa}(\Sigma)\log(3)}{\log(8)} \\
&> \frac{(g-1)\log(g-1) - g + \log(9.1)n - (4g+2n-6)\log(3)}{\log(8)} \\
&> \frac{(g-1)\log(g-1) - (4\log(3)+1)g}{\log(8)} + \frac{\log(9.1)n - \log(9)n}{\log(8)} \\
&> B(g\log(g+1) + n)
\end{aligned}$$

where the latter inequality holds for $g \geq 705$ and B can be taken to be equal to 10^{-3}. The final assertion can be checked directly for the cases $g < 705$. $\qquad\square$

As before, we note that by putting lower bounds on g and n, the constant B can be improved.

References

[1] AGOL, I. Ideal triangulations of pseudo-Anosov mapping tori. In *Topology and geometry in dimension three*, vol. 560 of *Contemp. Math.* Amer. Math. Soc., Providence, RI, 2011, pp. 1–17.

[2] ARAMAYONA, J., KOBERDA, T., AND PARLIER, H. Injective maps between flip graphs. *Preprint* (2014).

[3] ARAMAYONA, J., LECUIRE, C., PARLIER, H., AND SHACKLETON, K. J. Convexity of strata in diagonal pants graphs of surfaces. *Publ. Mat. 57*, 1 (2013), 219–237.

[4] ARAMAYONA, J., PARLIER, H., AND SHACKLETON, K. J. Totally geodesic subgraphs of the pants complex. *Math. Res. Lett. 15*, 2 (2008), 309–320.

[5] ARAMAYONA, J., PARLIER, H., AND SHACKLETON, K. J. Constructing convex planes in the pants complex. *Proc. Amer. Math. Soc. 137*, 10 (2009), 3523–3531.

[6] BOSE, P., AND VERDONSCHOT, S. *Computational Geometry.* Springer Berlin Heidelberg, 2012, ch. A history of flips in combinatorial triangulations, pp. 29–44.

[7] BROWN, W. G. Enumeration of triangulations of the disk. *Proc. London Math. Soc. (3) 14* (1964), 746–768.

[8] CAVENDISH, W. Growth of the diameter of the pants graph modulo the mapping class group. *Preprint* (2011).

[9] CAVENDISH, W., AND PARLIER, H. Growth of the Weil-Petersson diameter of moduli space. *Duke Math. J. 161*, 1 (2012), 139–171.

[10] CORTÈS, C., GRIMA, C. I., HURTADO, F., A., M., SANTOS, F., AND VALENZUELA, J. Transforming triangulations on nonplanar surfaces. *SIAM J. Discrete Math. 24* (2010), 821–840.

[11] FARB, B., AND MARGALIT, D. *A primer on mapping class groups*, vol. 49 of *Princeton Mathematical Series.* Princeton University Press, Princeton, NJ, 2012.

[12] FOMIN, S., SHAPIRO, M., AND THURSTON, D. Cluster algebras and triangulated surfaces. I. Cluster complexes. *Acta Math. 201*, 1 (2008), 83–146.

[13] FOMIN, S., AND THURSTON, D. Cluster algebras and triangulated surfaces. part II: Lambda lengths. *Preprint* (2012).

[14] HAMENSTÄDT, U. Geometry of the mapping class group II: A biautomatic structure. *Preprint* (2009).

[15] HATCHER, A. On triangulations of surfaces. *Topology Appl. 40*, 2 (1991), 189–194.

[16] KORKMAZ, M., AND PAPADOPOULOS, A. On the ideal triangulation graph of a punctured surface. *Ann. Inst. Fourier 4*, 62 (2012), 1367–1382.

[17] LOERA, J. A. D., RAMBAU, J., AND SANTOS, F. *Triangulations: structures for algorithms and applications*. No. 20 in Algorithms and Computation in Mathematics. Springer, 2010.

[18] MASUR, H. A., AND MINSKY, Y. N. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal. 10*, 4 (2000), 902–974.

[19] MOSHER, L. Tiling the projective foliation space of a punctured surface. *Trans. Amer. Math. Soc. 306*, 1 (1988), 1–70.

[20] MOSHER, L. Mapping class groups are automatic. *Annals of Mathematics 142*, 2 (1995), pp. 303–384.

[21] NEGAMI, S. Diagonal flips in triangulations on closed surfaces, estimating upper bounds. *Yokohama Mathematical Journal 45* (1998), 113–124.

[22] NEGAMI, S. Diagonal flips in pseudo-triangulations on closed surfaces. *Discrete Math. 240* (2001), 187–196.

[23] PENNER, R. C. Weil-petersson volumes. *J. Differential Geom. 35* (1992), 559–608.

[24] PENNER, R. C. Universal constructions in Teichmüller theory. *Adv. Math. 98*, 2 (1993), 143–215.

[25] PENNER, R. C. *Decorated Teichmüller theory*. Eur. Math. Soc., Zürich, 2012.

[26] POURNIN, L. The diameter of associahedra. *Adv. Math. 259* (2014), 13–42.

[27] RAFI, K., AND SCHLEIMER, S. Covers and the curve complex. *Geometry & Topology 2009* (13), 2141–2162.

[28] RAFI, K., AND TAO, J. Diameter of the thick part of moduli space and simultaneous Whitehead moves. *Duke Mathematical Journal 162* (2013), 1833–1876.

[29] RAFI, K., AND TAO, J. Uniform growth rate. *Preprint* (2014).

[30] SLEATOR, D. D., TARJAN, R. E., AND THURSTON, W. P. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc. 1*, 3 (1988), 647–681.

[31] SLEATOR, D. D., TARJAN, R. E., AND THURSTON, W. P. Short encodings of evolving structures. *SIAM J. Discrete Math. 5*, 3 (1992), 428–450.

[32] STASHEFF, J. D. Homotopy associativity of *H*-spaces. I, II. *Trans. Amer. Math. Soc. 108* (1963), 293–312.

[33] TAMARI, D. *Monoïdes préordonnés et chaînes de Malcev*. Thèse, Université de Paris, 1951.

[34] TAYLOR, S. J., AND ZUPAN, A. Products of farey graphs are totally geodesic in the pants graph. *Preprint* (2013).

[35] TSUKUI, Y. Transformations of cubic graphs. *J. Franklin Inst. B 333*, 4 (1996), 565–575.

Addresses:

Department of Mathematics, University of Fribourg, Switzerland

Indiana University, Bloomington IN, USA

Emails: hugo.parlier@unifr.ch, vdisarlo@indiana.edu