

Flip graphs for infinite type surfaces

Ariadna Fossas and Hugo Parlier*

Abstract. We associate to triangulations of infinite type surface a type of flip graph where simultaneous flips are allowed. Our main focus is on understanding exactly when two triangulations can be related by a sequence of flips. A consequence of our results is that flip graphs for infinite type surfaces have uncountably many connected components.

1. Introduction

A variety of simplicial complexes have been used to study surfaces, their homeomorphisms and their geometric structures. For finite type surfaces, arc and curve type graphs have been very useful tools for studying the geometry of different moduli spaces. In particular, flip graphs give a way of measuring distance between triangulations but also provide a coarse model for mapping class groups. In this article, we adapt flip graphs to the setting of infinite type surfaces. As one might expect, the passage to infinite type surfaces requires a little bit of care.

Our starting point will always be a surface Σ obtained by pasting together an infinite collection of triangles, and then removing the vertices who then belong to the space of ends of Σ . We define a graph $\mathcal{F}(\Sigma)$ whose vertices are these triangulations up to isotopy and whose edges come from flipping arcs that lie in quadrilaterals. More precisely, two triangulations are joined by an edge if they are related by any number (possibly infinite) of flips that can be performed simultaneously. This adaptation of the usual flip graph has already been studied in the finite type case [6, 8], and importantly, for infinite type surfaces it allows one to measure distances between a larger set of triangulations. If one only allow single flips, two triangulations can only be related if they differ by finitely many arcs (see [12] for an example of such a flip graph).

Figuring out which triangulations are related by a sequence of flips in our setting is exactly the main result of this paper:

Theorem 1.1. *Let Σ be an infinite type surface. Let S and T be triangulations of Σ . Then, S and T are in the same connected component of $\mathcal{F}(\Sigma)$ if and only if there exists $K \geq 0$ such that for every α arc of S and every β arc of T the intersection numbers $i(\alpha, T)$ and $i(\beta, S)$ are bounded by K .*

*Supported by the Luxembourg National Research Fund OPEN grant O19/13865598.

2020 Mathematics Subject Classification: Primary: 57K20, 05C10. Secondary: 05C12, 32G15, 30F60.

Key words and phrases: triangulations, flip graphs, infinite type surfaces.

The proof of the theorem involves putting together a number of preliminary results, one of them being Proposition 2.7, which uses a technique from [10, 15] which shows that triangulations that share a multiarc form a convex subset of $\mathcal{F}(\Sigma)$. The rest of proof is mainly combinatorial, and relies on a graph coloring argument, and in particular Brooks' theorem.

Using Theorem 1.1, it is straightforward to construct examples of triangulations that are not related by sequences of flip transformations, showing that $\mathcal{F}(\Sigma)$ has multiple connected components. In fact we show:

Corollary 1.2. *For any Σ of infinite type, $\mathcal{F}(\Sigma)$ has uncountably many connected components.*

This is analogous to what happens for hyperbolic structures for infinite type surfaces. The graph $\mathcal{F}(\Sigma)$ can be thought of as a combinatorial analogue of Teichmüller space which, classically, is the space of conformal structures up to quasi-conformal map and also has uncountably many connected components [3]. There is a metric point of view to what we do: if one replaces triangles by ideal hyperbolic triangles, and pastes them with 0 shear, the resulting hyperbolic metric is well defined. A flip now changes the hyperbolic metric, but the two metrics are bi-Lipschitz equivalent. Although this analogy is not concretely used here in any way, it would be interesting to explore to what extent $\mathcal{F}(\Sigma)$ provides a combinatorial model to spaces that arise in the smooth setting.

In a more quantitative direction, our methods give the following upper and lower bounds on flip distance between two triangulations in terms of the maximal intersection between the individual arcs of one triangulation and the other triangulation. The following is the combination of corollaries 2.6 and 3.3.

Corollary 1.3. *Let $K \geq 1$ be a constant and T and T' be triangulations of a surface Σ .*

If every arc α of T and every arc β of T' , the intersection numbers $i(\alpha, T')$ and $i(\beta, T)$ are bounded above by K , then T and T' are related by at most

$$2K^2 \cdot 3^K - K^2$$

flips.

Conversely, if T contains an arc α that satisfies $i(\alpha, T') \geq K$, then the flip distance between T and T' is at least

$$\log_4(3(K+1)).$$

Note that as far as we can tell, the above statement is also new for finite type surfaces.

Organization.

The article is organized as follows. In Section 2, after some definitions and notation, we prove some results that provide the groundwork for the proving the main theorem. In particular, we

show the easier direction of Theorem 1.1, the proof of which is given in Section 3. In the final section (Section 4), we discuss some corollaries of our results, namely about the connected components of flip graphs (Corollary 1.2).

Acknowledgments.

We thank Stefan Wenger for an inspiring conversation about bi-Lipschitz equivalent planar metrics when we first started working on this project.

2. Setup and preliminary results

Let Σ be a connected orientable surface obtained by the following procedure. We begin with a countable collection of triangles and paste the sides of triangles in pairs to obtain a connected orientable surface $\bar{\Sigma}$ (see Figure 1 for an illustration). The image of the vertices of the triangles under the pasting is a collection of points which we denote by P and call the ideal vertex set. We now set $\Sigma = \bar{\Sigma} \setminus P$. Note that each element of P belongs to an end of Σ , and although they don't belong to Σ , this ideal vertex set is implicit when we use the notation Σ . Although it will also be used just for the topological surface obtained by this procedure, Σ is really a pair consisting in the surface and its ideal vertex set.

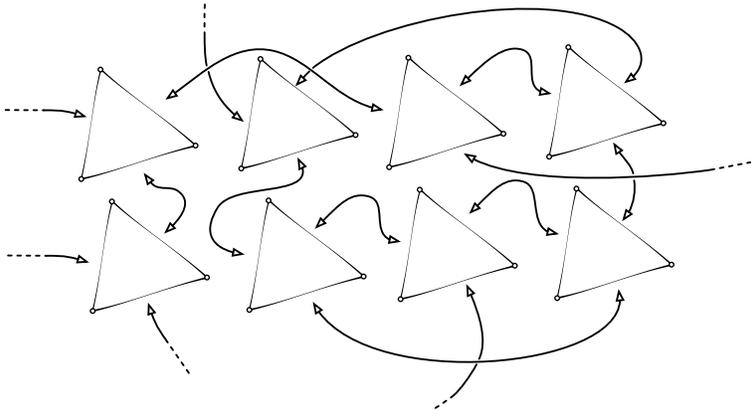


Figure 1: Pasting triangles to obtain Σ

Arcs of Σ are non-trivial simple paths between (non-necessarily distinct) elements of P . We denote by $\mathcal{A}(\Sigma)$ the set of arcs of Σ up to isotopy fixing P pointwise. We are interested in triangulations of Σ by which we mean disjoint collections of arcs between elements of P that decompose Σ into a collection of (open) triangles.

Note that arc and curve type graphs are becoming more understood in the context of infinite type surfaces [5, 9, 11, 1] have been studied in different contexts and for different uses but, to the best of our knowledge, graphs with vertices being triangulations of infinite type surfaces have not yet been studied.

Remark 2.1. An alternative approach would be to consider a surface with a fixed set of marked points and then consider maximal multiarcs (sets of isotopy classes of arcs, disjoint in their interior and maximal for inclusion). For finite type surfaces, these approaches are equivalent because given any surface (of negative Euler characteristic) and non-empty set of marked points, a maximal multiarc always decomposes the surface into triangles. For infinite type surfaces, if a triangulation exists then it is maximal but the converse is not always true. In fact, there are infinite type surfaces with a prescribed set of marked points P which do not admit *any* triangulations with vertex set P . An easy example is the following: consider a one ended infinite genus surface and P consisting of a single marked point. A maximal set of arcs contains an infinite number of arcs, and thus, if you trace a small circle around the marked point and look at the intersection with arcs (if necessary, realize the arcs with a metric) there will be accumulation points. These points correspond to accumulation arcs, which cannot exist in a triangulation.

Surfaces with sets of ideal vertices that can obviously be constructed via such a triangulation include $\mathbb{R}^2 \setminus \mathbb{Z}^2$ with ideal vertex set \mathbb{Z}^2 as portrayed in figure 2. This surface is commonly called the flute surface [2].

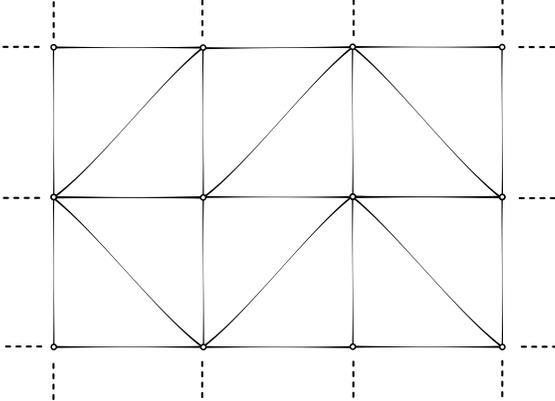


Figure 2: One way to (locally) paste triangles to obtain the flute surface

A more subtle surface is obtained via the standard representation of the Farey graph in the hyperbolic plane with $\mathbb{Q} \cup \{\infty\}$ as ideal vertices. Note that in this case $\Sigma \cong \mathbb{R}^2$, so it only has one end, and all arcs leave and terminate in this end. And yet $P = \mathbb{Q} \cup \{\infty\}$, and there is an order on this set. Surfaces with infinite genus can also be obtained (for instance by adding genus to each triangle in the Farey graph tessellation to obtain the Loch Ness Monster surface as in Figure 3), but, as seen previously, any arbitrary combination of Σ and P is not possible.

Now given Σ , we define an associated flip graph $\mathcal{F}(\Sigma)$.

Definition 2.2. Let μ be a (possibly infinite) multiarc of T such that every arc $a \in \mu$ bounds

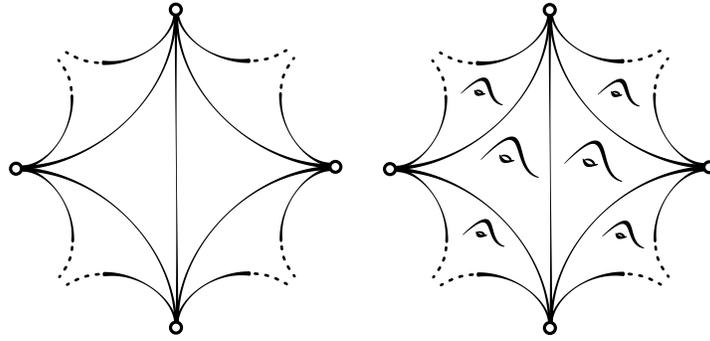


Figure 3: To the right of the Farey triangulation is the Loch Ness Monster: each “triangular” genus 1 sub surface is easily triangulated

two distinct triangles (that form a quadrilateral with a as a diagonal). Suppose further that if $a, b \in \mu$ with $a \neq b$, the quadrilaterals containing a and b as diagonals are distinct. We define the triangulation $f_\mu(T) = T'$ to be the one obtained from T by replacing every arc $a \in \mu$ by the other diagonal arc of the quadrilateral defined by the two triangles containing a . We say that T and T' are related by a *simultaneous flip* or simply a *flip*.

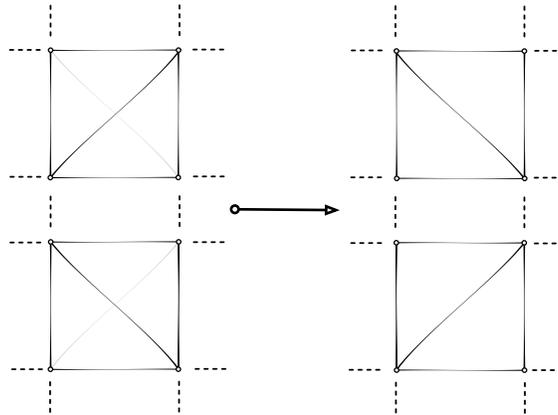


Figure 4: A simultaneous flip is done performing single flips simultaneously on disjoint quadrilaterals

This allows us to define $\mathcal{F}(\Sigma)$: vertices are the set of triangulations of Σ and there is an edge between T and T' if they are related by a flip. We denote the connected components of $\mathcal{F}(\Sigma)$ by $\mathcal{F}^i(\Sigma)$, $i \in I$, where I is an index set.

When Σ is of finite type, $\mathcal{F}(\Sigma)$ is always connected [14] but we will see that when Σ is of infinite type, it is always a graph with infinitely many connected components.

The triangulations lying in a given connected component are formed of arcs of $\mathcal{A}(\Sigma)$. Our

first observations are the following.

Proposition 2.3. *Let $T_\alpha \in \mathcal{F}(\Sigma)$ and $\alpha \in \mathcal{A}(\Sigma)$. Then $i(\alpha, T) < +\infty$.*

Proof. The argument is similar to above: if the arc intersected an infinite number of arcs of T , the intersection points would contain an accumulation point. This accumulation point cannot belong to the interior of a triangle, or to the interior of an edge of the triangle, so does not belong to Σ . \square

Proposition 2.4. *Let $\mathcal{F}^i(\Sigma) \subset \mathcal{F}(\Sigma)$ be a connected component and $\alpha \in \mathcal{A}(\Sigma)$ an arc. Then there exists a triangulation $T_\alpha \in \mathcal{F}^i(\Sigma)$ containing α .*

Proof. For any $T \in \mathcal{F}^i(\Sigma)$, $i(T, \alpha) < +\infty$. Thus there is a finite type subsurface in which α and the subset of T intersected by α both live. It suffices to flip in that subsurface to obtain a triangulation containing α . \square

This implies that for any finite multiarc and any given connected component, there is a triangulation containing that multiarc. This is no longer true for infinite multiarcs. In particular, as mentioned previously, in the sequel there will be examples of triangulations that are not related by any finite number of simultaneous flips.

The following proposition is the easy part of Theorem 1.1 and it states that if there is no bound on the intersection between arcs of a triangulation T and an other triangulation T' , then they cannot be related by flips.

Proposition 2.5. *Let $T, T' \in \mathcal{F}(\Sigma)$ be such that for any $K > 0$, there exists $\alpha \in T$ such that $i(\alpha, T') \geq K$, then T and T' cannot be related by a finite number of simultaneous flips.*

Proof. For an integer $N > 0$, consider a sequence of triangulations obtained from T by flipping N times. We denote the sequence by $T_0 = T, \dots, T_N$.

We set $K = \frac{1}{3}(4^{N+1} - 1)$ for reasons that will become apparent in what follows, and let α be an arc with $i(\alpha, T') \geq K$. We will construct a sequence of arcs $\alpha_i \in T_i$ (by induction) for each $i = 1, \dots, N$, such that

$$i(\alpha_i, T') \geq \frac{1}{3}(4^{N+1-i} - 1).$$

This will show that $T_N \neq T'$, as $i(\alpha_N, T') > 0$.

We set α_0 to be α to begin the induction.

At any stage, if the arc α_i does not belong to those that are flipped, we set $\alpha_{i+1} := \alpha_i$ (and in particular its intersection number with T' remains unchanged). If it does belong to an arc

that is flipped, then α_i belongs to a quadrilateral Q_i with four boundary arcs, the collection of which we denote ∂Q_i . Now the key observation is that the quantity

$$\max_{\delta \in \partial Q_i} \{i(\delta, T')\}$$

is bounded below by a linear function of $i(\alpha_i, T')$. This is simply because any arc that intersects α_i must then intersect ∂Q_i in both directions, unless it terminates at a vertex of Q_i . An arc that does not terminate in Q_i contributes 2 to the intersection between ∂Q_i and T' , and an arc that terminates contributes 1, unless it is the other diagonal of Q_i , but there is only such diagonal.

Hence $i(\partial Q_i, T') \geq i(\alpha_i, T') - 1$, where the -1 is to account for the diagonal. Now as Q_i has 4 boundary arcs:

$$\max_{\delta \in \partial Q_i} \{i(\delta, T')\} \geq \frac{i(\alpha_i, T') - 1}{4} \quad (1)$$

as claimed.

We now set α_{i+1} to be an arc of Q_i which realizes this maximal intersection.

We now finish by observing that

$$i(\alpha_{i+1}, T') \geq \frac{i(\alpha_i, T') - 1}{4} \geq \frac{\frac{1}{3}(4^{N+1-i} - 1) - 1}{4} = \frac{1}{3}(4^{N-i} - 1)$$

as claimed. We point out that the choice of K came from the inductive step (inequality 1 above), and the fact that

$$\sum_{j=1}^N 4^j = \frac{1}{3}(4^{N+1} - 4).$$

So after any sequence of N flips leaving from T , the resulting triangulation always has an arc that continues to intersect T' , and hence cannot be T' . As this is true for any N , the two triangulations are never connected by a sequence of flips. \square

The above proof also results in the following quantitative statement.

Corollary 2.6. *If T and T' are such that an arc α of T satisfies*

$$i(\alpha, T') \geq K$$

for $K \geq 1$, then the flip distance between T and T' is at least

$$\log_4(3(K+1)).$$

Proof. The argument above shows that you need at least N flips to relate T and T' if

$$i(\alpha, T') \geq \frac{1}{3}(4^N - 1).$$

By setting $N = \log_4(3(K+1))$, the result follows. \square

One ingredient in our proofs will be the use of projections to natural subgraphs consisting in the triangulations that contain a given multiarc. These projections were at least implicitly studied in [14], and were studied in detail in [10], but in both setups, Σ was of finite type and the flip graphs involved were the "usual" ones, where only a single flip was allowed.

For a given multiarc μ of Σ we define $\mathcal{F}_\mu^i(\Sigma)$ to be the graph of $\mathcal{F}^i(\Sigma)$ spanned by all vertices T containing μ . If μ is finite, then $\mathcal{F}_\mu^i(\Sigma)$ is always non-empty but otherwise, as mentioned above, this might not be the case.

The following result has been shown in [15] for finite type surfaces, and is true in this more general setting as well.

Proposition 2.7. $\mathcal{F}_\mu^i(\Sigma)$ is a convex subgraph of $\mathcal{F}^i(\Sigma)$.

Proof. Note the theorem is true in the finite type setting (for simultaneous flips this is proved in [15] using the strategy for individual flips from [10]).

The basic observation that allows one to adapt the proof for the finite type setting here is the following. We suppose $\mathcal{F}_\mu^i(\Sigma)$ is non-empty, that is that there exists at least one triangulation of Σ which contains μ . By Proposition 2.5, this means that for any $T \in \mathcal{F}^i(\Sigma)$ we have

$$\sup_{\alpha \in \mu} i(\alpha, T) < +\infty.$$

We are going to project a triangulation T to a triangulation containing μ by "combing" T along μ (see Figure 5).

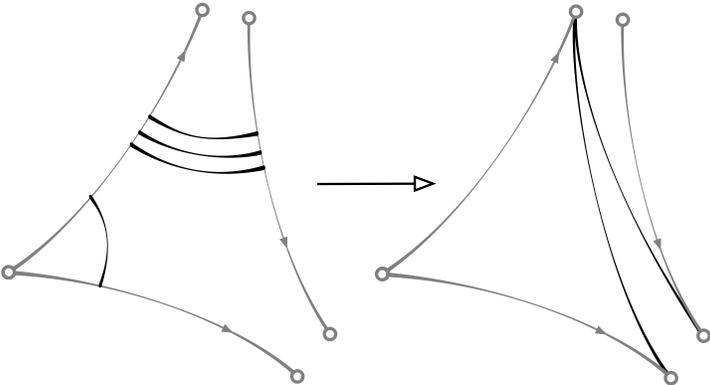


Figure 5: "Combing" a triangulation along an oriented multiarc

The local picture will be the same as in the finite type setting as each arc of μ will intersect a bounded number of arcs of T .

Consider a multiarc μ , and give each of its arcs an orientation (the choice of orientation does not matter).

Now a combing projection works as follows. For a given $T \in \mathcal{F}^i(\Sigma)$, each arc of T is sent to a triangulation containing μ which is defined by “combing” T along μ with the given orientation. Specifically one defines the map as follows: T is intersected by μ and hence results in a collection of subarcs on which are the connected components of $\Sigma \setminus \mu$. For each such subarc a , we define an arc as follows: from an interior point of a (that is a non terminal point), extend it in both directions until the terminal points of a . These are points of μ (possibly a marked point). If these are interior points of an arc α of μ , continue the arc following α with its orientation to its terminal endpoint. This map clearly sends a to an arc, and sends distinct subarcs a, b to (interior) disjoint arcs. Note that different subarcs can be sent to the same arc.

One needs to check that the result is a triangulation, namely that the connected components of the complementary regions to the resulting multiarcs are all triangles. This is relatively straightforward to check: the rough argument is that otherwise there is a complementary region that is of greater (arc) complexity but which must have been intersected by arcs of T or μ . These arcs must have resulted in arcs that continue to intersect the region, hence this cannot be. A detailed argument in the finite type case can be found in both [10] and [15], and the same argument applies here.

Thus the combing map from $\mathcal{F}^i(\Sigma)$ to $\mathcal{F}_\mu^i(\Sigma)$ is well-defined on vertices. Now it suffices to show that edges are sent to edges. Roughly speaking, this can be deduced from the fact that quadrilaterals are sent to a collection of quadrilaterals, triangles or arcs, with at most one quadrilateral in the image. And if multiple quadrilaterals are disjoint, then their images are disjoint. Again, see [15] for a more detailed argument.

Using this projection, convexity can be deduced from the observation that the combing map leaves points of $\mathcal{F}_\mu^i(\Sigma)$ invariant. Given two triangulations $T, T' \in \mathcal{F}_\mu^i(\Sigma)$, the image of geodesic under the combing map between them is a path between them of length at most $d(T, T')$ and entirely contained in $\mathcal{F}_\mu^i(\Sigma)$. \square

This proposition has an immediate consequence. Note that by a metric space of infinite rank we mean the absence of finite rank, meaning that for any positive integer k , it contains a quasi-convex copy \mathbb{Z}^k with the usual metric.

Corollary 2.8. *If Σ is of infinite type, any connected component of the flip graph is of infinite diameter and infinite rank.*

Proof. This follows from the fact that given any triangulation $T \in \mathcal{F}^i(\Sigma)$, you can choose a multiarc $\eta \subset T$ that separates the surface into infinitely many finite type surfaces each with topology:

$$\Sigma \setminus \eta = \dot{\cup}_{k \in \mathbb{N}} \Sigma_k.$$

We can now take a bi-infinite geodesic γ_k on each of the flip-graphs of the finite type subsurfaces Σ_k . These are naturally collection of multiarcs that live on Σ through the inclusion of Σ_k in Σ . This gives rise to a collection of triangulations in $\mathcal{F}^i(\Sigma)$ by choosing a point of γ_k for all $k \in \mathbb{N}$ and taking the union with η . The convexity of $\mathcal{F}_\mu^i(\Sigma)$ means that this is a convex subset giving us quasi-copies of \mathbb{Z}^k for any $k \in \mathbb{N}$ (the "usual" metric and the "diagonal" metric are quasi-isometric). \square

3. The proof of Theorem 1.1

This section is dedicated to proving Theorem 1.1. We begin with a lemma, where we explicitly need a property of the combing projection from Proposition 2.7.

Lemma 3.1. *Let T and T' be triangulations of Σ and let α be an oriented arc of T . Let T'_α be the combing projection of T' onto $\mathcal{F}_\alpha(\Sigma)$. Then*

$$\sup\{i(\beta', T) : \beta' \in T'_\alpha\} \leq \sup\{i(\beta, T) : \beta \in T'\}$$

and for every $\gamma \in T$, $i(\gamma, T'_\alpha) \leq i(\gamma, T)$.

Proof. Note that the above quantities could be infinite. By definition, every arc β' of T'_α is either α , an arc of T' or the concatenation of a subarc of a $\beta \in T'$ and a subarc of α . In the first two cases, the lemma is obvious, and in the final case, the lemma follows as $i(\alpha, T) = 0$ and thus $i(\beta', T) \leq i(\beta, T)$. \square

Lemma 3.2. *There exists a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that the following holds. Let S and T be two triangulations of Σ satisfying*

$$\max\{i(\alpha, T), i(\beta, S) : (\alpha, \beta) \in S \times T\} \leq K$$

for some finite K . Then there exist multiarcs $\mu_1, \dots, \mu_{f(K)}$ of S such that

1. $S = \bigcup_{i=1}^{f(K)} \mu_i$ and
2. if $i(\alpha, \gamma) > 0$ for $\alpha \in \mu_i$ and $\gamma \in T$, then $i(\alpha', \gamma) = 0$ for all $\alpha' \in \mu_i \setminus \{\alpha\}$.

The function can be taken to be

$$f(K) = 2 \cdot 3^K - 1.$$

Proof. Consider the graph G whose vertices are the arcs of S and where two vertices are joined by an edge when the corresponding arcs belong to the boundary of the same triangle. Note

that the degree of all vertices is uniformly bounded by 4. We can define a distance between two arcs of S as the combinatorial distance of the corresponding vertices of G .

Let α_1 and α_2 be two arcs of S both intersected by some arc β of T . The arc β induces a path in Σ from some point in the interior of α_1 to some point in the interior of α_2 . Hence, this path crosses, by hypothesis, at most $K - 2$ other arcs of S . Thus, the distance between α_1 and α_2 in G is at most $K - 1$.

Now take a multiarc μ of S such that the distance in G for every pair of different arcs of μ is at least K . This forces any arc of T to intersect at most one arc in μ .

Finally, consider the $K - th$ power of G (the graph G^K having the same vertex set as G and where two vertices are joined by an edge if their distance in G is at most K). The degree of all vertices of G^K is uniformly bounded by $2 \cdot 3^K - 2$. By Brooks' coloring theorem [7] there exists a $(2 \cdot 3^K - 1)$ -coloring of the vertices of G^K such that any two adjacent vertices are of a different color. We can now define $\mu_1, \dots, \mu_{f(K)}$ to be the monochromatic multiarcs and thus $f(K) \leq 2 \cdot 3^K - 1$. \square

We can now proceed to the proof.

Proof of Theorem 1.1. Suppose first that S and T are two maximal triangulations in the same connected component of $\mathcal{F}(\Sigma)$ and consider a path $S = T_0, T_1, \dots, T_k = T$ joining them. By definition, $i(\alpha, T_0) \leq 1$ for all $\alpha \in T_1$. Note that if α belongs to T_i , every arc of T_0 crossing α either is the other diagonal of the quadrilateral of T_i containing α or it crosses at least one of the boundary arcs of this quadrilateral. Thus, $i(\alpha, T_0) \leq 4 \max\{i(\beta, T_0) : \beta \in T_{i-1}\} + 1$. Hence, for all $\alpha \in T_k$ the intersection with S is at most $\frac{4^k - 1}{3}$. The same procedure applies to the path $T = S_0, T_{k-1} = S_1, \dots, S_k = T_0 = S$ showing that for all $\beta \in S$ the intersection with S is also at most $\frac{4^k - 1}{3}$.

On the other direction, suppose now that there exists $K \geq 0$ such that all arcs α of $S \cup T$ satisfy $i(\alpha, S \cup T) \leq K$. We now apply Lemma 3.2 to decompose the maximal multiarc S into $\mu_1, \dots, \mu_{f(K)}$. To do so we give every arc of μ_i an orientation. For $i \in \{1, \dots, f(K)\}$ we define T_i as the combing projection of T_{i-1} along (oriented) μ_i , with $T_0 = T$. To see that such a simultaneous combing projection exists, let α, α' be two different arcs of μ_i and denote Σ_α (respectively $\Sigma_{\alpha'}$) the subsurface of Σ spanned by all triangles of T_{i-1} intersecting α (respectively α'). Then the interiors of Σ_α and $\Sigma_{\alpha'}$ are disjoint as a consequence of lemmas 3.1 and 3.2. Furthermore, any arc of T_{i-1} intersects at most one of α, α' . In addition the complexities of Σ_α are uniformly bounded because they can be triangulated with at most K triangles. Hence, the total intersection number of $T_{i-1} \cap T_i \cap \Sigma_\alpha$ is bounded above by K^2 and the distance in the corresponding flip graphs is at most K^2 (see for instance Corollary 2.13

in [10]). As these surfaces Σ_α have disjoint interiors, those flips can be done simultaneously, hence there is a path in $\mathcal{F}(\Sigma)$ between T_{i-1} and T_i of length at most K^2 . This completes the proof. \square

As a corollary to the above proof, we get the following quantitative statement (which does not require the surface to be of infinite type in any way).

Corollary 3.3. *For a constant $K \geq 0$, let S and T be triangulations of a surface Σ such that for every α arc of S and every β arc of T the intersection numbers $i(\alpha, T)$ and $i(\beta, S)$ are bounded above by K . Then the simultaneous flip distance between T and S is bounded above by*

$$2K^2 \cdot 3^K - K^2.$$

Proof. By Lemma 3.2 the function $f(K)$ satisfies

$$f(K) \leq 2 \cdot 3^K - 1$$

and this is a bound on the number of "steps" necessary to get from T to S . As argued above, each step requires at most K^2 (simultaneous) flips, hence the result. \square

4. The topology of $\mathcal{F}(\Sigma)$

We finish this paper with some observations about connected components of these flip graphs. Deformation spaces of infinite type surfaces generally have infinite numbers of connected components. For instance, the Teichmüller space of hyperbolic structures up to either quasi-conformal or bi-Lipschitz maps have this property (see [2, 3, 13]). One reason is because these maps only deform the lengths of geodesics by bounded amounts, and on infinite type surfaces, you have a lot of room to construct hyperbolic structures with wildly behaving length spectra [4].

In our setting, we use our result to show that our flip graphs have many connected components.

Corollary 4.1. *For any Σ of infinite type, $\mathcal{F}(\Sigma)$ has uncountably many connected components.*

Proof. On Σ there are infinitely many disjoint simple closed curves, say δ_i , $i \in \mathbb{N}$. Now consider a triangulation $T \in \mathcal{F}(\Sigma)$.

Observation 1. Each δ_i intersects finitely many arcs of T .

We argue as in Proposition 2.3. If not, there would be an accumulation of intersection points between the arcs and δ_i . The accumulation point is a point of δ_i , hence a point that belongs to

either the interior or the interior of an edge of a triangle. But in either case, it cannot be the accumulation point of a collection of arcs.

Observation 2. Any arc of T intersects a finite number of curves among $\delta_i, i \in \mathbb{N}$.

This is because the curves are disjoint, so for instance can be completed into a pants decomposition and as our arcs are two ended, it cannot pass through infinitely many pairs of pants.

So for each δ_i , we have a finite collection of arcs of T that intersect it, which in turn intersect a finite number of other curves. As a consequence we have:

There exists a collection of pairs δ_i, α_j where $j \in \mathbb{N}, \alpha_j \in T, i(\delta_i, \alpha_j) \neq 0$, and $i(\alpha_j, \delta_{i'}) = 0$ if $j \neq j'$.

In other words, by successively choosing pairs of an intersecting curve and arc disjoint from all previous pairs, we get this infinite collection. Up to renumbering and relabelling, we can suppose that we have a collection $\delta_i, \alpha_i, i \in \mathbb{N}$, such that

$$i(\delta_i, \alpha_i) \leq i(\delta_{i+1}, \alpha_{i+1}).$$

Roughly speaking, we'll obtain new triangulations from T by twisting the curves along increasing powers of Dehn twists.

To obtain an uncountable number of connected components of $\mathcal{F}(\Sigma)$, we begin by choosing a collection \mathcal{C} of infinite subsets of \mathbb{N} such that any two subsets contain infinitely many different elements. (Said in other way, the difference sets are always infinite.)

Note that \mathcal{C} is uncountable: it is of the same cardinality as infinite binary sequences which pairwise differ in infinitely many indices. And infinite binary sequences, without the difference condition, have the cardinality of \mathbb{R} . And for any given infinite binary sequence, there are at most a countable number of elements that only differ for a finite number of indices, hence our restricted set is also uncountable.

Now we give each δ_i an orientation, and for $N \in \mathcal{C}$, we consider the homeomorphism φ_N obtained by Dehn twisting δ_i m_i times if $i \in N$ where m_i is taken to be such that

$$i(\varphi_N(\alpha_i), \alpha_i) \geq i.$$

We now observe that if $M, N \in \mathcal{C}$ and $M \neq N$, then $\varphi_M(T)$ and $\varphi_N(T)$ belong to different connected components. Indeed, consider the sequence of integers $M \setminus N$, which is infinite by, and denote them $i_k, k \in \mathbb{N}$.

Thus the homeomorphism φ_M acts on α_{i_k} by twisting m_{i_k} times along δ_{i_k} but φ_N leaves α_{i_k} invariant. Hence:

$$i(\varphi_N(\alpha_{i_k}), \varphi_M(\alpha_{i_k})) = i(\alpha_{i_k}, \varphi_M(\alpha_{i_k})) \geq i_k.$$

And so there are arcs of $\varphi_M(T)$ with arbitrarily large intersection with arcs of $\varphi_N(T)$, and thus by Theorem 1.1, lie in different connected components. \square

References

- [1] Aramayona, Javier, Fossas, Ariadna and Parlier, Hugo. Arc and curve graphs for infinite-type surfaces. *Proc. Amer. Math. Soc.* 145 (2017), no. 11, 4995–5006.
- [2] Basmajian, Ara. Hyperbolic structures for surfaces of infinite type. *Trans. Amer. Math. Soc.* 336 (1993), no. 1, 421–444.
- [3] Basmajian, Ara. Large parameter spaces of quasiconformally distinct hyperbolic structures. *J. Anal. Math.* 71 (1997), 75–85.
- [4] Basmajian, Ara and Kim, Youngju. Geometrically infinite surfaces with discrete length spectra. *Geom. Dedicata* 137 (2008), 219–240.
- [5] Bavard, Juliette. Hyperbolicité du graphe des rayons et quasi-morphismes sur un gros groupe modulaire. *Geom. Topol.* 20 (2016), no. 1, 491–535.
- [6] Prosenjit Bose, Jurek Czyzowicz, Zhicheng Gao, Pat Morin, and David R. Wood. Simultaneous diagonal flips in plane triangulations. *J. Graph Theory* 54 (2007), no. 4, 307–330.
- [7] Brooks, R. Leonard. On colouring the nodes of a network. *Proc. Cambridge Philos. Soc.* 37 (1941), 194–197.
- [8] Disarlo, Valentina and Parlier, Hugo. Simultaneous flips on triangulated surfaces. *Michigan Math. J.* 67 (2018), no. 3, 451–464.
- [9] Durham, Matthew Gentry, Fanoni, Federica and Vlamiš, Nicholas G. Graphs of curves on infinite-type surfaces with mapping class group actions. *Ann. Inst. Fourier (Grenoble)* 68 (2018), no. 6, 2581–2612.
- [10] Disarlo, Valentina and Parlier, Hugo. The geometry of flip graphs and mapping class groups. *Trans. Amer. Math. Soc.* 372 (2019), no. 6, 3809–3844.
- [11] Fanoni, Federica, Ghaswala, Tyrone and McLeay, Alan. Homeomorphic subsurfaces and the omnipresent arcs. Preprint, 2020.
- [12] Fossas, Ariadna and Nguyen, Maxime. Thompson’s group \mathcal{T} is the orientation-preserving automorphism group of a cellular complex. *Publ. Mat.* 56 (2012), no. 2, 305–326.

- [13] Liu, Lixin and Papadopoulos, Athanase. Some metrics on Teichmüller spaces of surfaces of infinite type. *Trans. Amer. Math. Soc.* 363 (2011), no. 8, 4109–4134.
- [14] Mosher, Lee. Mapping class groups are automatic. *Ann. of Math. (2)* 142 (1995), no. 2, 303–384.
- [15] Parlier, Hugo and Weber, Ashley. Geometric simplicial embeddings of arc-type graphs. *J. Korean Math. Soc.*, to appear.

Addresses:

GAP Nonlinearity and Climate, Institut des Sciences de l'Environnement, Université de Genève, Suisse

Department of Mathematics, FSTM, University of Luxembourg, Esch-sur-Alzette, Luxembourg

Emails:

ariadna.fossastenas@unige.ch

hugo.parlier@uni.lu