

BIG FLIP GRAPHS AND THEIR AUTOMORPHISM GROUPS

ASSAF BAR-NATAN, ADVAY GOEL, BRENDAN HALSTEAD,
PAUL HAMRICK, SUMEDH SHENOY, AND RISHI VERMA

ABSTRACT. In this paper, we study the relationship between the mapping class group of an infinite-type surface and the simultaneous flip graph, a variant of the flip graph for infinite-type surfaces defined by Fossas and Parlier [8]. We show that the extended mapping class group is isomorphic to a proper subgroup of the automorphism group of the flip graph, unlike in the finite-type case. This shows that Ivanov’s metaconjecture, which states that any “sufficiently rich” object associated to a finite-type surface has the extended mapping class group as its automorphism group, does not extend to simultaneous flip graphs of infinite-type surfaces.

1. INTRODUCTION

The *mapping class group* $\text{Map}(S)$ of a surface S is the group of orientation-preserving homeomorphisms of S up to isotopy. Expanding $\text{Map}(S)$ to include orientation-reversing homeomorphisms yields the *extended mapping class group* $\text{Map}^\pm(S)$. The *flip graph* $\mathcal{F}(S)$ of S is the graph whose vertices are triangulations and whose edges are elementary moves, or *flips*. This graph characterizes finite-type surfaces up to homeomorphism, and the automorphisms of the flip graph are precisely those induced by the action of the extended mapping class group [15]. This exemplifies a broader pattern described by the following “metaconjecture” due to Ivanov.

Conjecture 1.1 (Problem 6 of [14]). *Every object naturally associated to a surface and having sufficiently rich structure has the extended mapping class group as its group of automorphisms.*

Here, “sufficiently rich” is intentionally left undefined. The Ivanov metaconjecture has been shown to be true for many objects in addition to the flip graph, such as the curve complex [13], the arc complex [12], and Torelli buildings [7], and in 2019, the Ivanov metaconjecture was shown to be true for an even larger class of complexes on finite-type surfaces [4].

For infinite-type surfaces, i.e., surfaces whose fundamental group is not finitely-generated, it is known that the automorphism group of the curve graph is the extended mapping class group [2, 10]. Therefore, in this paper, motivated by Ivanov’s metaconjecture and the above results, we study the same question in the setting of flip graphs on infinite-type surfaces. We will consider the simultaneous flip graph, a variant of the flip graph defined by Fossas and Parlier [8], which adapts the flip graph for infinite-type surfaces. We show that—unlike in the finite-type case—the symmetries of this graph are no longer fully captured by the extended mapping class group. In particular, we prove the following result.

Theorem 1.2. *The natural action of the extended mapping class group on the flip graph is an isomorphism from $\text{Map}^\pm(\Sigma)$ to a proper subgroup of $\text{Aut}(\mathcal{F}(\Sigma))$.*

To prove this, we first show that the isotopy class of a homeomorphism is determined by its action on $\mathcal{F}(\Sigma)$. For other simplicial complexes on which $\text{Map}^\pm(\Sigma)$ acts, such as the curve

complex and the pants complex, this can be accomplished using the Alexander Method for infinite-type surfaces developed by Hernandez et al. in [11]. To apply this result to ideal arcs, we prove the following lemma, which may be of interest by its own merit.

Lemma 3.1 (Alexander Method for Ideal Arcs). *Let Σ be an infinite-type surface, and let f be a homeomorphism of Σ . For $e \in \text{Ends}(\Sigma)$, let \mathcal{L}_e denote the collection of ideal arcs that are “loops at e ” — those which map both ends of $(0, 1)$ to e . The following are equivalent.*

- (i) *There exists $e \in \text{Ends}(\Sigma)$ such that for all $\lambda \in \mathcal{L}_e$, the ideal arc $f(\lambda)$ is isotopic to λ .*
- (ii) *The homeomorphism f is isotopic to the identity homeomorphism Id_Σ .*

Using the terminology of [11], this means that for any end e , the set \mathcal{L}_e forms an Alexander system for Σ . This also establishes that the mapping class group acts faithfully on the arc graphs discussed in [1].

Acknowledgements. We would like to thank the Canada/USA Mathcamp for providing a welcoming space and setting to meet, discuss, and develop the ideas in this paper. We would also like to thank Kasra Rafi, Hugo Parlier, and Alan McLeay for their support and helpful conversations. Finally, we would like to thank Mark Bell for telling us about the Python package “Bigger” and talking to us about the subtleties of triangulations on infinite-type surfaces.

2. PRELIMINARIES

This section is an introduction to the concepts and terminology regarding infinite-type surfaces that will be used throughout the paper. A *surface* is a connected, orientable topological 2-manifold. A surface Σ is called *finite-type* if $\pi_1(\Sigma)$ is finitely-generated. Otherwise, Σ is called *infinite-type*.

The *end space* of an infinite-type surface Σ , introduced by Freudenthal [9], and denoted by $\text{Ends}(\Sigma)$, is $\varprojlim \pi_0(\Sigma \setminus K)$, where K ranges over all compact subsurfaces of Σ and the discrete spaces $\pi_0(\Sigma \setminus K)$ are ordered by inclusion. We equip $\text{Ends}(\Sigma)$ with the inverse limit topology, and note that it is totally disconnected, compact, and metrizable. The set $\bar{\Sigma} := \Sigma \sqcup \text{Ends}(\Sigma)$, equipped with an appropriate topology, is called the *Freudenthal compactification* of the surface Σ [9]. Concretely, an end can be thought of as a descending chain $e = U_1 \supset U_2 \supset \dots$ of connected open sets such that $\cap_i U_i = \emptyset$. If there exists n such that U_n can be embedded in \mathbb{R}^2 , then e is called *planar*. Otherwise, it is said to be *non-planar* or *accumulated by genus*. We denote by $\text{Ends}_g(\Sigma)$ the subspace of non-planar ends. See [6] for a more detailed exposition.

In 1963, Richards showed that orientable infinite-type surfaces are classified up to homeomorphism by their genus and the pair of spaces $(\text{Ends}(\Sigma), \text{Ends}_g(\Sigma))$. For more information, see [17].

A *simple closed curve* is an embedding of the circle S^1 into Σ . An *ideal arc*, or simply an *arc*, on a surface Σ is an embedding $\alpha : (0, 1) \rightarrow \Sigma$ which extends continuously to a map $\bar{\alpha} : [0, 1] \rightarrow \bar{\Sigma}$ such that the ends of $(0, 1)$ are mapped to ends of Σ . A *multi-arc* is a collection of disjoint arcs. An arc α in Σ is called *peripheral* if $\Sigma \setminus \alpha$ has a connected component that is a disk. An arc is called *essential* if it is not peripheral. A *multi-arc* is a possibly-infinite collection of disjoint arcs, and an isotopy class of a multi-arc is the collection of isotopy classes of its corresponding disjoint arcs.

Let $[\mu]$ and $[\nu]$ be isotopy classes of multi-arcs. The *intersection number* $i(\mu, \nu)$ is the minimum value of $|\mu \cap \nu|$ among all isotopy representatives of μ and ν .

Definition 2.1. An *ideal triangulation* of a surface, denoted T , is a locally-finite maximal collection pairwise disjoint, pairwise non-homotopic, essential ideal arcs whose complementary components are open disks. We also require that T have a tubular neighbourhood homotopic to itself.

There are other alternative definitions of ideal triangulations on infinite-type surfaces in the literature, such as in [8, 16]. We will sometimes consider a triangulation of Σ to be a subspace of Σ rather than a collection of maps into Σ . The closures of complementary components of a triangulation are called *triangles*. It follows from maximality that the boundary of a triangle can consist of either two or three arcs.

If an arc α in a triangulation bounds only one triangle Δ , then $\partial\Delta$ must consist of a single additional arc β [15]. In this situation we call β a *petal arc* for α .

In [16], McLeay and Parlier establish that every infinite-type surface admits an ideal triangulation. Moreover, they determine when exactly a multiarc is a subset of a triangulation:

Lemma 2.2 (Theorem 1.2 of [16]). *An ideal multiarc μ is a subset of an ideal triangulation if and only if for any simple closed curve γ , the intersection number $i(\gamma, \mu) < \infty$.*

If a multiarc μ satisfies the conditions above, then there exists a triangulation T containing μ . We call such a triangulation a *completion* of μ .

Definition 2.3. Consider a triangulation T and a disjoint collection of arcs $\mu \subset T$ such that for any $\alpha \in \mu$ bounding distinct triangles Δ_1 and Δ_2 , we have $\Delta_1 \cap \Delta_2 \cap \mu = \alpha$. A *simultaneous flip* sends T to a triangulation T' by replacing each $\alpha \in \mu$ by the other diagonal of its bounding quadrilateral.

Figure 1 depicts an example of a simultaneous flip. We can now define the *flip graph* $\mathcal{F}(\Sigma)$ of a surface Σ , whose vertices are isotopy classes of ideal triangulations of Σ . If Σ is an infinite-type surface, we connect two triangulations by an edge if they are related by a simultaneous flip. If Σ is finite-type, we connect triangulations by an edge only if they are related by a single flip.




FIGURE 1. An example of a simultaneous flip of a triangulation, with the dotted blue arcs before the flip and the red arcs afterward.

Although the flip graph is connected in the finite-type case [5], the flip graph in the infinite-type case has uncountably many connected components. In particular, [8] establishes conditions on when two triangulations are in the same connected component.

Lemma 2.4 (Theorem 1.1 of [8]). *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 arc α of S and every arc β of T , the intersection numbers $i(\alpha, T)$ and $i(\alpha, S)$ are bounded by K .*

Note that there is a natural action of the mapping class group on the flip graph sending an ideal triangulation to its image under the homeomorphism of the surface. Since this action commutes with flips, each mapping class induces an automorphism of the flip graph. In the next section, we show that there exists an automorphism of the flip graph which is not induced by any mapping class.

3. PROOF OF THEOREM 1.2

In this section, we show that the extended mapping class group is a strict subgroup of the automorphism group of the flip graph for infinite-type surfaces. Our methods were first inspired by those used by Branman to show an analogous result for pants graphs [3].

Since the flip graph is disconnected, the connected components can in some sense be manipulated independently of one another. However, every mapping class necessarily affects triangulations in every connected component of the flip graph. This is the essential reason that the symmetries of Σ fail to fully capture the symmetries of $\mathcal{F}(\Sigma)$. After establishing some properties of the action on $\mathcal{F}(\Sigma)$, we will prove this by providing an automorphism of $\mathcal{F}(\Sigma)$ that cannot be induced by a mapping class.

We begin by proving that mapping classes are determined by their action on certain arcs.

Lemma 3.1 (Alexander Method for Ideal Arcs). *Let Σ be an infinite-type surface, and let f be a homeomorphism of Σ . For $e \in \text{Ends}(\Sigma)$, let \mathcal{L}_e denote the collection of ideal arcs that are “loops at e ” — those which map both ends of $(0, 1)$ to e . The following are equivalent.*

- (i) *There exists $e \in \text{Ends}(\Sigma)$ such that for all $\lambda \in \mathcal{L}_e$, the ideal arc $f(\lambda)$ is isotopic to λ .*
- (ii) *The homeomorphism f is isotopic to the identity homeomorphism Id_Σ .*

Proof. Assuming (ii), we see that restricting the isotopy from Id_Σ to f yields an isotopy from λ to $f(\lambda)$. Hence, (ii) implies (i). To show the converse, we claim that any mapping class satisfying (i) must fix every isotopy class of curves on Σ . Then we can use the Alexander Method of [11] to conclude that f must be isotopic to Id_Σ .




FIGURE 2. A diagram showing the main ideas in Lemma 3.1.

Let c be a curve in Σ , and let γ be a path from e to a point p on c such that $|c \cap \gamma| = 1$. Choose a tubular neighborhood U of γ such that $|\partial U \cap c| = 2$. Let $\bar{\gamma}$ be a path in U from a point \bar{p} on c to e such that $\gamma \cap \bar{\gamma} = \emptyset$ and $|c \cap \bar{\gamma}| = 1$. Define $\lambda \in \mathcal{L}_e$ to be the arc that

begins at e , follows γ from e to p , then follows the connected component of $c \setminus \{p, \bar{p}\}$ not contained in U , and finally runs back to e along $\bar{\gamma}$. See Figure 2. Since $c \cup f(c)$ is compact, we can choose a separating curve b such that one component of $\Sigma \setminus b$ contains e and the other contains $c \cup f(c)$. We also require that b intersects γ , $\bar{\gamma}$, $f(\gamma)$, and $f(\bar{\gamma})$ at exactly one point each.

We now define an isotopy from c to $f(c)$ in terms of the isotopy I from λ to $f(\lambda)$ given by the hypothesis. First, observe that c is isotopic to a curve \tilde{c} contained in $\lambda \cup b$. Likewise, $f(c)$ is isotopic to a curve $\tilde{f(c)}$ contained in $f(\lambda) \cup b$. I extends to an isotopy between \tilde{c} and $\tilde{f(c)}$ by translating $\tilde{c} \cap b$ along b . \square




FIGURE 3. The standard triangulation of the flute surface.

Remark. Initially, one might think that if a mapping class fixes a particular triangulation, it must be the identity. This is not always the case, however. For example, consider the triangulation T of the flute surface depicted in Figure 3. Then, the mapping class that shifts every point on the plane to the right one unit fixes T (not as a set, but as a triangulation). Note that this homeomorphism is not isotopic to the identity. In other words, arbitrary triangulations do not form Alexander systems. This implies that generally, $\text{Map}^\pm(\Sigma)$ does not act freely on $\mathcal{F}(\Sigma)$.

The following technical lemma will be used in the proof of Lemma 3.3 to glean information about the action of $\mathcal{F}(\Sigma)$ on individual arcs from its action on triangulations.

Lemma 3.2. *Let α and β be distinct ideal arcs on Σ , such that α is not a petal arc for β . Then, there exist ideal triangulations T_1, T_2 in distinct connected components of the flip graph that each contain α but not β .*

Proof. We first construct T_1 . If $i(\alpha, \beta) > 0$, then any completion of α can be taken as T_1 . Otherwise, by Lemma 2.2, arcs α and β can be completed into a triangulation T . If β is flippable, we flip β to get a triangulation T' that contains α but not β . Otherwise, β is unflippable, which implies that the triangles bordering both sides of the arcs are the same, and that β is contained within a petal formed by an arc γ . Then, γ must border two distinct triangles, so γ is flippable. After flipping γ , β will thus be flippable, and flipping β will result in a triangulation T_1 that contains α but not β .

Next, we construct T_2 . Since $\Sigma' = \Sigma \setminus (\alpha \cup \beta)$ is still an infinite-type surface, we can choose a countable collection of disjoint simple closed curves δ_i on Σ' . Note that each δ_i , when thought of as a curve on Σ , is disjoint from α and β . We apply the procedure described in Corollary 4.1 of [8] to T_1 using the collection $\{\delta_i\}$ to obtain a triangulation T_2 in a distinct connected component of $\mathcal{F}(\Sigma)$, containing α but not β . \square

Remark. Note that the procedure described above can create uncountably many triangulations, each in a distinct connected component and containing α but not β . This procedure is done precisely in [8]. However, it is not true that every connected component contains such a triangulation. Consider the triangulation T of the flute surface shown in Figure 3. We claim that no triangulation in the connected component of T contains an arc α with an endpoint at the non-isolated end. To see why this is true, observe that none of the arcs of T connect to the non-isolated end, so this end cannot be the vertex of a bounding quadrilateral of an arc. Hence, a simultaneous flip cannot connect any new arcs to this end. Since the resulting triangulation will retain this property, we conclude that no triangulation containing α is reachable from T by a series of simultaneous flips. This implies that the modular flip graph [5] is disconnected for certain infinite-type surfaces, a result which will appear in greater detail in a forthcoming paper by the authors.

The next result shows that mapping classes are determined by their action on the flip graph. In addition, the weaker hypothesis allows us to determine the action of a mapping class on one connected component of $\mathcal{F}(\Sigma)$ by its action on the other components. We use this in the proof of Theorem 1.2 to show that a particular automorphism of $\mathcal{F}(\Sigma)$ is not induced by a mapping class.

Lemma 3.3. *If $f \in \text{Map}^\pm(\Sigma)$ acts trivially on all except for one connected component of $\mathcal{F}(\Sigma)$, then f is the identity in $\text{Map}^\pm(\Sigma)$.*

Proof. Let α an arc in Σ . We show that $f(\alpha)$ cannot be β for any arc $\beta \neq \alpha$, which implies that $f(\alpha) = \alpha$. Let $\beta \neq \alpha$ be arbitrary.

First, suppose that α is a petal arc of β . We observe that β must have two distinct endpoints, while α only has one, so $f(\alpha)$ cannot be β . In the case where α is not a petal arc of β , Lemma 3.2 guarantees the existence of triangulations T_1, T_2 containing α but not β such that T_1 and T_2 are in distinct connected components of $\mathcal{F}(\Sigma)$. Then, f must act trivially on either T_1 or T_2 , so $f(\alpha) \neq \beta$. Thus, $f(\alpha) = \alpha$, and by Lemma 3.1, f is the identity in $\text{Map}^\pm(\Sigma)$. \square

Corollary 3.4. *The action of $\text{Map}^\pm(\Sigma)$ on $\mathcal{F}(\Sigma)$ is faithful. In other words, $\text{Map}^\pm(\Sigma)$ embeds into $\text{Aut}(\mathcal{F}(\Sigma))$.*

The preceding result establishes that the flip graph has at least as many symmetries as the surface. But since so little is known about the graph structure of $\mathcal{F}(\Sigma)$, it is difficult to conceive of automorphisms that are not induced by specific mapping classes. The following lemma improves the situation to some extent, allowing us to construct automorphisms of individual connected components via compactly-supported mapping classes.

Lemma 3.5. *Let f be a compactly-supported mapping class. Then, f induces an automorphism on every connected component of $\mathcal{F}(\Sigma)$. In particular, f fixes the connected components of $\mathcal{F}(\Sigma)$.*

Proof. Let $T \in \mathcal{F}(\Sigma)$ be arbitrary. We will show that T and $f(T)$ must lie in the same connected component of $\mathcal{F}(\Sigma)$. Since f has compact support, there exists a finite-type subsurface Σ' of Σ containing every point $x \in T \subset \Sigma$ for which $f(x) \neq x$. It follows that every transverse intersection of arcs of T and arcs of $f(T)$ is contained within Σ' .

Only finitely many arcs of T and $f(T)$ can pass through Σ' , because otherwise, the arcs of T or $f(T)$ would accumulate. Let $\mathcal{A} \subset \Sigma'$ be this set of arcs, and let $K = \max_{\alpha, \beta \in \mathcal{A}} |\alpha \cap \beta|$. This

maximum exists, because intersection points between arcs in \mathcal{A} are discrete and contained in Σ' . Note that for any $\alpha \in T$, we have $i(\alpha, f(T)) \leq K|\mathcal{A}|$, and for any $\beta \in f(T)$, we have $i(\beta, T) \leq K|\mathcal{A}|$. Thus, by Lemma 2.4, T and $f(T)$ lie in the same connected component of $\mathcal{F}(\Sigma)$. Since f induces an automorphism on $\mathcal{F}(\Sigma)$, it must restrict to an automorphism on each connected component of $\mathcal{F}(\Sigma)$. \square

We are now ready to prove Theorem 1.2.

Proof of Theorem 1.2. Let f be a nontrivial mapping class with compact support, and let Γ be a connected component of $\mathcal{F}(\Sigma)$ that f acts nontrivially on. Note that f induces an automorphism on Γ by Lemma 3.5, which we also denote f . We define an automorphism \bar{f} of $\mathcal{F}(\Sigma)$ by

$$\bar{f}(T) = \begin{cases} f(T) & T \in \Gamma \\ T & T \in \mathcal{F}(\Sigma) \setminus \Gamma \end{cases}.$$

By Lemma 3.3, a mapping class inducing \bar{f} would necessarily act trivially on $\mathcal{F}(\Sigma)$, which \bar{f} does not. Thus, \bar{f} cannot be induced by a mapping class. Together with Corollary 3.4, this establishes that the natural action of $\text{Map}^\pm(\Sigma)$ corresponds to a homomorphism $\text{Map}^\pm(\Sigma) \rightarrow \text{Aut}(\mathcal{F}(\Sigma))$ that is injective but not surjective. \square

4. DISCUSSION

In this paper, we constructed an automorphism of the flip graph not induced by a mapping class. In other words, there are graph isomorphisms between flip graphs that are not induced by homeomorphisms of the corresponding surfaces. A natural follow-up question to ask is whether simultaneous flip graphs even determine a surface:

Question 4.1. If $\mathcal{F}(\Sigma)$ and $\mathcal{F}(\Sigma')$ are isomorphic as graphs, does this mean that Σ and Σ' are homeomorphic?

The method we used in this paper considers how automorphisms of the flip graph on distinct connected components may be incompatible with mapping classes. However, we might try to restrict our attention to a single connected component of the flip graph. We can ask whether every automorphism of a connected component is induced by a mapping class:

Question 4.2. For a given connected component Γ of the flip graph $\mathcal{F}(\Sigma)$, is $\text{Aut}(\Gamma)$ isomorphic to $\text{Map}^\pm(\Sigma)$?

REFERENCES

- [1] Javier Aramayona, Ariadna Fossas, and Hugo Parlier, *Arc and curve graphs for infinite-type surfaces*, Proceedings of the American Mathematical Society **145** (2017), no. 11, 4995–5006.
- [2] Juliette Bavard, Spencer Dowdall, and Kasra Rafi, *Isomorphisms between big mapping class groups*, International Mathematics Research Notices **2020** (2018), no. 10, 3084–3099.
- [3] Ben Branman, *Spaces of pants decompositions for surfaces of infinite type*, arXiv preprint arXiv:2010.13169 (2020).
- [4] Tara E. Brendle and Dan Margalit, *Normal subgroups of mapping class groups and the metaconjecture of Ivanov*, Journal of the American Mathematical Society **32** (2019), 1009–1070.
- [5] Valentina Disarlo and Hugo Parlier, *The geometry of flip graphs and mapping class groups*, Transactions of the American Mathematical Society **372** (2019), no. 6, 3809–3844.
- [6] Federica Fanoni, Tyrone Ghaswala, and Alan McLeay, *Homeomorphic subsurfaces and the omnipresent arcs*, Annales Henri Lebesgue **4** (2021), 1565–1593.

- [7] Benson Farb and Nikolai V. Ivanov, *Torelli buildings and their automorphisms*, arXiv preprint arXiv:1410.6223 (2014).
- [8] Ariadna Fossas and Hugo Parlier, *Flip graphs for infinite type surfaces*, arXiv preprint arXiv:2011.02324 (2020).
- [9] Hans Freudenthal, *Über die enden topologischer räume und gruppen*, Springer, 1931 (ger).
- [10] Jesús Hernández Hernández, Israel Morales, and Ferrán Valdez, *Isomorphisms between curve graphs of infinite-type surfaces are geometric*, Rocky Mountain Journal of Mathematics **48** (2018), no. 6, 1887–1904.
- [11] ———, *The alexander method for infinite-type surfaces*, Michigan Mathematical Journal **68** (2019), no. 4, 743–753.
- [12] Elmas Irmak and John D. McCarthy, *Injective simplicial maps of the arc complex*, Turkish Journal of Mathematics **34** (2010), no. 3, 339–354.
- [13] Nikolai V. Ivanov, *Automorphisms of complexes of curves and of teichmuller spaces*, International Mathematics Research Notices **1997** (1997), no. 14, 651–666.
- [14] ———, *Fifteen problems about the mapping class groups*, arXiv preprint arXiv:0608325 (2006).
- [15] Mustafa Korkmaz and Athanase Papadopoulos, *On the ideal triangulation graph of a punctured surface* **62** (2012), no. 4, 1367–1382.
- [16] Alan McLeay and Hugo Parlier, *Ideally, all infinite type surfaces can be triangulated*, arXiv preprint arXiv:2102.09531 (2021).
- [17] Ian Richards, *On the classification of noncompact surfaces*, Transactions of the American Mathematical Society **106** (1963), 259–269.