

Curves intersecting exactly once and their dual cube complexes

Tarik Aougab and Jonah Gaster

Abstract. Let S_g denote the closed orientable surface of genus g . We construct exponentially many mapping class group orbits of collections of $2g + 1$ simple closed curves on S_g which pairwise intersect exactly once, extending a result of the first author [1] and further answering a question of Malestein, Rivin, and Theran [10]. To distinguish such collections up to the action of the mapping class group, we analyze their dual cube complexes in the sense of Sageev [12]. In particular, we show that for any even k between $\lfloor g/2 \rfloor$ and g , there exists such collections whose dual cube complexes have dimension k , and we prove a simplifying structural theorem for any cube complex dual to a collection of curves on a surface pairwise intersecting at most once.

Keywords. Curves on surfaces, curve systems.

Mathematics Subject Classification (2010). 57M05, 57M07, 57M20, 57M50, 57M60.

Contents

1	Introduction	1062
2	Background	1065
3	Sageev’s construction	1067
4	Mapping class group orbits of collections of curves	1068
5	Recognizing n -cubes dual to a 1-system	1070
6	A family of maximum complete 1-systems	1079
7	Restricting mapping class group orbits via $C(\Gamma(\epsilon))$	1085
8	The labeled polygon $P(\epsilon)$ associated to $\Gamma(\epsilon)$	1088
9	Stabilizing 1-systems	1097
	References	1100

1. Introduction

For $g, p \in \mathbb{N}$, $g \geq 2$, $p \geq 0$, let $S_{g,p}$ denote the orientable surface of genus g with p punctures or marked points, let $\text{Mod}^*(S)$ denote the associated *extended mapping class group*, and let $\mathcal{J}(S)$ be the set of isotopy classes of essential simple closed curves on S . $\text{Mod}^*(S)$ acts naturally on $\mathcal{J}(S)$, and analyzing orbits of various finite subsets of $\mathcal{J}(S)$ has proven to be a fruitful way of probing the algebra and geometry of $\text{Mod}^*(S)$. For example, the $\text{Mod}^*(S)$ -orbits of pants decompositions have been used to estimate the Weil-Petersson diameter of the thick part of Moduli space [4].

The main focus of this paper is to explicitly construct many distinct $\text{Mod}^*(S)$ -orbits of collections of curves with intersection properties that are reminiscent of pants decompositions. We show:

Theorem 1. *On S_g , there exists at least $2^{g-3}/(g-1)$, and at most $(4g^2 + 2g)!$ distinct $\text{Mod}^*(S)$ orbits of collections of $2g + 1$ simple closed curves pairwise intersecting once.*

Malestein, Rivin, and Theran have shown that any collection of curves pairwise intersecting once has cardinality at most $2g + 1$. Thus we think of such collections as analogous to pants decompositions in the following way: pants decompositions are the maximal cliques of the *curve graph*, $\mathcal{C}(S)$, of S – the graph with vertex set $\mathcal{J}(S)$, and edges between two isotopy classes that can be realized disjointly on S . Similarly, the curve systems in Theorem 1 are the largest cliques of the *Schmutz or systole graph*, $\mathcal{SC}(S)$, of S , whose vertices correspond to the subset of \mathcal{J} of non-separating simple closed curves, and whose edges correspond to pairs of curves intersecting exactly once.

Remark 1.1. Unlike pants decompositions, there do exist non-maximum collections of curves pairwise intersecting once that are nonetheless maximal with respect to inclusion. See Figure 1 for a complete 1-system on a surface of genus one with two boundary components so that there is no triangle in the complement of the curves. Gluing the two boundary components together, we find a complete 1-system of three curves with no triangle in their complement. There is only one $\text{Mod}^*(S_2)$ -orbit of five curves pairwise intersecting once on S_2 (see [10, p. 231]), and this orbit has the property that every trio forms triangles, so the example from Figure 1 is not contained in a maximum complete 1-system in genus two.

We remark that the number of $\text{Mod}^*(S_g)$ -inequivalent pants decompositions grows at least factorially in g (cf. [2]), and we conjecture that the same is true for the types of curve systems considered here. However, though $\mathcal{SC}(S)$ and $\mathcal{C}(S)$ are $\text{Mod}^*(S)$ -equivariantly quasi-isometric, it is not necessarily the case that the corresponding growth rates of $\text{Mod}^*(S)$ -inequivalent maximal cliques are comparable. This question requires a more detailed understanding of the specific nature of $\mathcal{SC}(S)$.




Figure 1. A complete 1-system that is not contained in a maximum complete 1-system.

Schmutz has shown that the automorphism group of $\mathcal{SC}(S)$ is the (extended) mapping class group [14], and thus as a corollary to Theorem 1 we obtain:

Corollary. *The number of maximum cardinality cliques of $\mathcal{SC}(S_g)$, inequivalent under the action of $\text{Aut}(\mathcal{SC})$, grows at least exponentially in g .*

By a k -system, we mean any subset $\Gamma \subset \mathcal{J}(S)$ consisting of curves pairwise intersecting at most k times. A *complete k -system* is a k -system in which any two curves intersect *exactly* k times. Thus the main focus of this paper is the study of complete 1-systems of maximum possible size, or *maximum complete 1-systems*.

Malestein, Rivin, and Theran showed that such 1-systems are unique up to the action of $\text{Mod}^*(S_g)$, for $g = 1, 2$, and asked if this uniqueness persists for higher genera. The first author answered this question by subsequently constructing two distinct orbits of complete 1-systems of size $2g + 1$ on S_g , for all $g \geq 3$. Thus we view Theorem 1 as a further demonstration of the non-uniqueness of maximum complete 1-systems.

Note that ‘complete’-ness, for the 1-systems we consider, is a significant simplifying assumption. Though there has been substantial recent progress towards estimating the size of maximum 1-systems [11], even asymptotically precise counts are not currently available. While it would be interesting to examine the number of $\text{Mod}^*(S)$ -orbits of maximum 1-systems, the absence of any examples when $g \geq 3$ makes this seem difficult.¹

Our method of distinguishing $\text{Mod}^*(S)$ -orbits of a curve system Γ is to analyze the *dual cube complex* $C(\Gamma)$ to Γ , a complex built from cubes of various dimensions which encodes the combinatorics of the intersections between curves. This invariant is a useful way of organizing topological information about the $\text{Mod}^*(S)$ -orbit of Γ . Along the way in our analysis, we show:

¹ [10] calculate that there are two $\text{Mod}^*(S)$ -orbits of maximum 1-systems for $g = 2$.

Theorem 2. *Let Λ_1, Λ_2 be any two collections of curves which fill a closed surface S . Then Λ_1 and Λ_2 are equivalent under the action of the extended mapping class group if and only if there is an isomorphism of cube complexes $C(\Lambda_1) \cong C(\Lambda_2)$. The induced set map from Λ_1 to Λ_2 corresponds to the induced map between hyperplanes of $C(\Lambda_1)$ and hyperplanes of $C(\Lambda_2)$.*

Thus, the reader may view the main result as a construction of many non-isomorphic cube complexes, each dual to a maximum complete 1-system. In particular, we show:

Proposition 1.2. *For any even $k \in [\lfloor g/2 \rfloor, g]$, there exists a complete 1-system of size $2g + 1$ on S_g whose dual cube complex has dimension k .*

It is interesting to consider whether or not the minimal dimension of the cube complex dual to a maximum complete 1-system grows with the genus. At the moment this problem seems difficult. As a first step, one might determine whether there exists a maximum complete 1-system Γ whose dual cube complex is 2-dimensional:

Question 1. *For $g \geq 3$, does there exist a complete 1-system of size $2g + 1$ whose dual cube complex is 2-dimensional?*

In this case, the quotient of the dual cube complex $C(\Gamma)$ by the action of $\pi_1(S)$ produces a square-tiled copy of S , with at least 4 squares around each vertex. We conjecture that the answer to Question 1 is no.

In general, $C(\Gamma)$ can be a very complicated combinatorial object; indeed, one may interpret this as a consequence of Theorem 2, since $\text{Mod}^*(S)$ -orbits of curve systems can be difficult to distinguish (cf. [9]). In order to leverage $C(\Gamma)$ to useful information about $\text{Mod}^*(S)$, we make use of the following simplifying theorem for cube complexes dual to 1-systems:

Theorem 3. *Suppose $\Gamma = \{\gamma_1, \dots, \gamma_n\}$ is any 1-system on a closed surface S . Then the dimension of $C(\Gamma)$ is n if and only if the dimension of $C(\Gamma')$ is 3, for Γ' any triple of curves in Γ .*

Remark 1.3. We note that one direction of Theorem 3 is immediate: if the dimension of the entire cube complex is n , then any three curves must correspond to a 3-cube in the dual cube complex. However, as Figure 3 in §5 demonstrates, the converse is false if the assumption of being a 1-system is dropped.

Our main construction requires g to be odd, and we extend the conclusion of Theorem 1 and Proposition 1.2 to even g via the following simple process (see §9 for a slightly more careful description):

Beginning with a complete 1-system Γ of size $2g + 1$ on S_g for $g = 2k + 1$, excise a pair of small open disks which are locally on opposite sides of some $\gamma \in \Gamma$, and glue on an annulus A along the resulting boundary circles. Note that Γ is still a complete 1-system on S_{g+1} , and we extend Γ to a collection of $2g + 3$ curves by adding γ', γ'' , defined as follows: both γ', γ'' run parallel to γ in the complement of the new annulus A . Within A , both γ', γ'' run from one boundary component of A to the other, intersecting once in the interior of A .

Therefore γ', γ'' intersect each other exactly once, within A , and each intersects all of the original elements of Γ exactly once because γ does. If Γ on S_g is obtained from a complete 1-system Γ' on S_{g-1} as described above, we call Γ a *stabilization* of Γ' .

It is natural to ask whether or not *every* complete 1-system on S_g of size $2g + 1$ is obtained from one on S_{g-1} of size $2g - 1$ via this process:

Question 2. *Let Γ be a complete 1-system on S_g of size $2g + 1$. Is it always the case that Γ is a stabilization of some complete 1-system on S_{g-1} ?*

We observe that each complete 1-system we construct is indeed a stabilization of a complete 1-system on a lower genus surface. Furthermore, we note that a positive answer to Question 2 implies a negative answer to Question 1: Lemmas 5.2 and 9.1 imply that the dimension of the cube complex dual to any complete 1-system obtained via stabilization is at least three.

Organization of paper. In §2 and §3 we outline some preliminary notions regarding the mapping class group and Sageev's construction of dual cube complexes. In §4, we prove Theorem 2; in §5, we prove Theorem 3; in §6, we outline the main construction of our complete 1-systems; in §7, §8, and §9 we prove that these complete 1-systems are indeed inequivalent up to the action of the extended mapping class group, completing the proof of Theorem 1.

2. Background

Let $S = S_g$ be a compact oriented surface of genus g , and let $\Gamma = \{\gamma_1, \dots, \gamma_n\}$ be a collection of free homotopy classes of closed curves on S . Recall that a collection of curves form a *bigon* if there is an embedded disk in S whose boundary is the union of two arcs of the curves. A *minimal position realization* of Γ is a set $\lambda = \{\eta_1, \dots, \eta_n\}$ such that

- (i) each $\eta_i: S^1 \rightarrow S$ is a smooth immersion in the free homotopy class γ_i ;
- (ii) the union $\bigcup_i \eta_i(S^1)$ forms no bigons;
- (iii) the immersed submanifolds $\{\eta_1(S^1), \dots, \eta_n(S^1)\}$ intersect only at transverse double points.

We will refer to minimal position realizations simply as *realizations*. In everything that follows, we suppose that $\lambda = \{\eta_1, \dots, \eta_n\}$ is a realization of Γ . Condition (ii) above implies that λ minimizes the sum of the pairwise geometric intersection numbers of the curves in Γ [7, Proposition 1.7, p. 31]. See [7, Chapter 1] for background on curves on surfaces.

Let Λ indicate the set of the lifts of elements of λ to the universal cover \tilde{S} , so that the elements of Λ are curves in \tilde{S} . The union of the curves in Λ may be considered as an embedded graph $G \subset \tilde{S}$. Condition (iii) above guarantees that each vertex of this graph has valence four, and condition (ii) implies that every pair of curves in Λ intersect at most once (see [7, Lemma 1.8, p. 30]).

When the curves in Γ are disjoint, then the dual graph to the lifts Λ admits an isometric action of $\pi_1 S$, and the quotient graph is an invariant for Γ that can be used to distinguish mapping class group orbits. However, in general the dual graph to λ in S is not an invariant of Λ , as different realizations may yield non-isomorphic graphs. For example, the presence of a triangle in the complement of λ allows a Reidemeister type III move, creating a new realization but changing the isomorphism types of the dual graph, as shown in Figure 2.




Figure 2. Changing the realization λ by a Reidemeister type III move changes the isomorphism type of the dual graph.

While the dual graph depends essentially on the realization λ , we now describe the construction of a related cube complex ‘dual’ to Λ which is independent of realization. Originally due to Sageev, this produces an isometric action of $\pi_1 S$ on a finite-dimensional cube complex, which we denote by $\widetilde{C(\lambda)}$, with quotient $C(\Lambda) := \widetilde{C(\lambda)}/\pi_1 S$, a cube complex with finitely many maximal cubes. Though it will be unnecessary in this work, when Λ is filling $C(\Lambda)$ is non-positively curved, and this construction can be placed in a considerably more general context (see [12], [13], and [5]).

Recall that a *cube* of dimension n is $[-1, 1]^n$, endowed naturally with the structure of a cell complex with a Euclidean metric, and a *cube complex* is the quotient of disjoint cubes by a gluing map which is a Euclidean isometry on each cell. In such a complex, a *maximal cube* is a cube which is not contained in a higher dimensional cube. The *dimension* of a cube complex is the supremum of the dimensions of its cubes. A *square complex* is a two-dimensional cube complex. See [13] for an introduction to cube complexes.

The *local hyperplanes* of an n -cube are the intersections of $[-1, 1]^n$ with coordinate planes. We introduce an equivalence relation on local hyperplanes in a cube complex: two local hyperplanes are *hyperplane equivalent* if they intersect along the 1-skeleton of the cube complex, and a *hyperplane* of a cube complex is the hyperplane equivalence class of a local hyperplane. An isomorphism of cube complexes is a homeomorphism that is a cellular Euclidean isometry, which thus preserves hyperplanes. We will denote the hyperplanes of a cube complex C by \mathcal{H}_C . For example, 1-dimensional cube complexes are graphs whose hyperplanes (and local hyperplanes) are midpoints of edges.

3. Sageev's construction

Fix a choice of realization λ of Γ , with lifts Λ as before. We now describe the construction of a cube complex $\widetilde{C}(\lambda)$ that is independent of the choice of realization. Each element $\gamma \in \Lambda$ separates $\widetilde{S} \setminus \gamma$ into two connected components. Fix a choice of identification of these two half-spaces with $\{1, -1\}$ for each $\gamma \in \Lambda$. We will refer to a choice of one of these two as a *labeling* of γ , and a labeling of Λ is a labeling for each of the curves of Λ . Identifying 2 with $\{1, -1\}$, we say a labeling $v \in 2^\Lambda$ is *admissible* if the half-spaces $v(\alpha)$ and $v(\beta)$ intersect for every pair of curves $\alpha, \beta \in \Lambda$.

Let $\mathcal{V}_0 \subset 2^\Lambda$ denote the collection of admissible labelings. Note that for each connected region $U \subset \widetilde{S} \setminus \Lambda$, there is an admissible labeling $v_\lambda(U) \in \mathcal{V}_0$ defined as follows. For each curve in Λ , $v_\lambda(U)$ chooses the half-space containing U . Let \mathcal{V}'_1 be the graph whose vertex set is \mathcal{V}_0 , where two labelings are joined by an edge when they differ on exactly one element of Λ . Hence there is an element of Λ associated to each edge of \mathcal{V}'_1 .

Choose a connected region $U \subset \widetilde{S} \setminus \Lambda$, and let \mathcal{V}_1 denote the connected component of $v_\lambda(U)$ in \mathcal{V}'_1 . (The choice of connected region U is evidently not essential). When lifts $\gamma_1, \dots, \gamma_n \in \Lambda$ pairwise intersect, it is straightforward to check that \mathcal{V}_1 contains an embedded copy of the 1-skeleton of $[-1, 1]^n$. We obtain the cube complex $\widetilde{C}(\lambda)$ by adding in the interior of any cube whose 1-skeleton is contained in \mathcal{V}_1 ; the links of the resulting cube complex are all flag simplicial complexes (see [5] for a more detailed exposition).

The action of $\pi_1 S$ on Λ induces a permutation action of $\pi_1 S$ on 2^Λ . For $g \in \pi_1 S$ and admissible labeling $v \in \mathcal{V}_0$, the labeling $g \cdot v \in \mathcal{V}_0$ is given by $(g \cdot v)(\gamma) = g \cdot v(g^{-1} \cdot \gamma)$ for each $\gamma \in \Lambda$. Since the action on \widetilde{S} is by deck transformations, it is straightforward to check that the action of $\pi_1 S$ preserves the admissible set \mathcal{V}_0 and induces an action by graph automorphisms on the 1-skeleton \mathcal{V}_1 . This action extends to $\widetilde{C}(\lambda)$ by definition.

The complex $\widetilde{C}(\lambda)$ with an action of $\pi_1 S$ is independent of the realization λ : a choice of half space for a lift $\gamma \in \Lambda$ determines, and is determined by, a choice of complement of $\partial\gamma \subset S^1 = \partial\pi_1 S$. By [7, Proposition 1.7, p. 31], choices of half-spaces for two lifts $\gamma, \gamma' \in \Lambda$ intersect if and only if the endpoints of γ link with those of γ' on S^1 . The latter is independent of the choice of realization. We thus denote $\widetilde{C}(\lambda)$ by $\widetilde{C}(\Gamma)$ when convenient.

Moreover, by Gromov's link condition [3] it is immediate that $\widetilde{C}(\Gamma)$ is non-positively curved. A direct argument shows that any loop in $\widetilde{C}(\Gamma)$ must backtrack, which implies that the complex is also simply-connected, and thus CAT(0). We collect the relevant information about $\widetilde{C}(\Gamma)$:

Theorem 3.1. [13] *Suppose that Γ is a filling collection of curves. Then the cube complex $\widetilde{C}(\Gamma)$ is CAT(0), and the action of $\pi_1 S$ is free, properly discontinuous, and cocompact. Given a realization λ , there is a $\pi_1 S$ -equivariant incidence-preserving identification of $\mathcal{H}_{\widetilde{C}(\Gamma)}$ with the lifts Λ .*

We will invoke the correspondence in Theorem 3.1 often, which we may refer to as the *curves to hyperplanes correspondence*. Briefly, this correspondence proceeds as follows. Each local hyperplane corresponds to a single coordinate of a cube, i.e. a single lift in Λ , and hyperplane equivalent local hyperplanes correspond to the same lift. Conversely, a single lift γ in Λ determines a subcomplex of $\widetilde{C}(\Gamma)$ spanned by admissible labelings whose chosen halfspaces, in each coordinate, intersect γ . In each maximal cube of this subcomplex, there is a local hyperplane which corresponds to switching the label at γ , and one can check that these local hyperplanes are all hyperplane equivalent. In what follows, we do not return to the combinatorial definition of $C(\Gamma)$.

4. Mapping class group orbits of collections of curves

We now characterize a filling curve system from the isomorphism type of its dual cube complex. Recall the *extended mapping class group* $\text{Mod}^*(S) = \text{Diff}(S)/\text{Diff}_0(S)$. We recall Theorem 2:

Theorem 2. *Two filling curve systems Λ_1 and Λ_2 are equivalent under the action of $\text{Mod}^*(S)$ if and only if there is an isomorphism of cube complexes $C(\Lambda_1) \cong C(\Lambda_2)$. The induced set map from Λ_1 to Λ_2 corresponds to the induced map between hyperplanes of $C(\Lambda_1)$ and hyperplanes of $C(\Lambda_2)$.*

Proof. One direction is straightforward. Suppose $\phi \cdot \Lambda_1 = \Lambda_2$ for $\phi \in \text{Mod}^*(S)$. Choose a realization λ_1 for Λ_1 , and a homeomorphism ϕ' realizing ϕ . In this case ϕ' induces a $\pi_1 S$ -equivariant isomorphism $\widetilde{\phi'}: \widetilde{C(\lambda_1)} \cong \widetilde{C(\phi \cdot \lambda_1)}$, where the induced map of hyperplanes corresponds to the set map induced by ϕ' from λ_1 to $\phi' \cdot \lambda_1$. Since $\phi' \cdot \lambda_1$ is a realization of Λ_2 , we have $C(\Lambda_1) \cong C(\lambda_1)$ and $C(\Lambda_2) \cong C(\phi \cdot \lambda_1)$. The result follows.

On the other hand, suppose $\Phi: C(\Lambda_1) \cong C(\Lambda_2)$ is an isomorphism of cube complexes. Choose realizations λ_1 and λ_2 for Λ_1 and Λ_2 . By Theorem 3.1, we have that $\widetilde{C(\Lambda_i)}$ is simply-connected and the action of $\pi_1 S$ is free and properly discontinuous, for each $i = 1, 2$ – note that this step requires the assumption that Λ_i is filling. Thus $\widetilde{C(\Lambda_i)}$ is the universal cover of $C(\Lambda_i)$ and there are isomorphisms $\pi_1 S \cong \pi_1 C(\Lambda_i)$ for each i . Composing these isomorphisms with Φ_* , we find an automorphism $\phi_*: \pi_1 S \rightarrow \pi_1 S$. By construction, we may lift Φ to a ϕ_* -equivariant isomorphism of cube complexes $\widetilde{\Phi}: \widetilde{C(\Lambda_1)} \rightarrow \widetilde{C(\Lambda_2)}$, in the sense that

$$\widetilde{\Phi}(g \cdot x) = \phi_*(g) \cdot \widetilde{\Phi}(x)$$

for each $x \in \widetilde{C(\Lambda_1)}$ and $g \in \pi_1 S$. This induces a corresponding equivariant map on hyperplanes, which in turn induces a correspondence of the collections of conjugacy classes of stabilizers of hyperplanes of $\widetilde{C(\Lambda_1)}$ with those of $\widetilde{C(\Lambda_2)}$.

The curves to hyperplanes correspondence guarantees that there is an identification of the collection of conjugacy classes of hyperplane stabilizers of $\pi_1 S$ acting on $\widetilde{C(\Lambda_i)}$ with the collection of conjugacy classes of stabilizers of curves in $\widetilde{\lambda_i}$, and so we arrive at a correspondence of the collection of conjugacy classes of stabilizers of curves in $\widetilde{\lambda_1}$ with those of $\widetilde{\lambda_2}$. The conjugacy classes of the stabilizers of curves in $\widetilde{\lambda_i}$ are naturally identified with the collection of conjugacy classes of Λ_i in $\pi_1 S$, from which it follows that the automorphism ϕ_* of $\pi_1 S$ takes the conjugacy classes determined by Λ_1 to those of Λ_2 . By the Dehn–Nielsen–Baer Theorem [7, Chapter 8, Theorem 8.1], there is $\phi \in \text{Mod}^*(S)$ inducing ϕ_* , so that $\phi \cdot \Lambda_1 = \Lambda_2$. □

5. Recognizing n -cubes dual to a 1-system

There remains the problem of recognizing quantitative information about $C(\Lambda)$ from a given set of curves Λ . In this section, we prove Theorem 3, giving a criterion for recognizing the dimensions of cubes in the complex dual to a 1-system.

A realization of a collection of curves forms a *triangle* if there is an embedded disk on S whose boundary components are three arcs of the curves, and so that these arcs intersect pairwise exactly once on the boundary of the disk. A collection of homotopy classes of curves forms a triangle if there is a realization of the curves that forms a triangle.

Lemma 5.1. *If a collection of homotopy classes of closed curves forms a triangle then every realization of the curves forms a triangle.*

Proof. By [6, Theorem 1] any two realizations are homotopic through isotopies and finitely many Reidemeister type III moves. Neither of these changes the existence of a triangle in the complement. \square

Lemma 5.2. *If $\Gamma = \{\gamma_1, \gamma_2, \gamma_3\}$ is a realization of simple closed curves on S , then $\dim C(\Gamma) = 3$ if and only if $S \setminus \Gamma$ has a connected component that is a triangle.*

Proof. If $\dim C(\Gamma) = 3$ and Γ is in minimal position, then the curves to hyperplanes correspondence guarantees that there are three mutually intersecting lifts of Γ in \tilde{S} . As Γ is a realization, the intersections of these lifts are not concurrent, and there is a triangle T in their complement. Since T is compact, the intersection $T \cap \tilde{\Gamma}$ consists of finitely many arcs. As each such arc doesn't form a bigon with the boundary of the triangle, the complement of the arc has a triangular component. Thus there is a triangular component $T' \subset T \setminus \tilde{\Gamma}$ (cf. with *innermost* bigon, [7, p. 31]).

Let $\pi: \tilde{S} \rightarrow S$ be the covering map. If $\text{int}(T')$ does not embed under π , then there is an element $g \in \pi_1 S$ so that $\text{int}(T') \cap g \cdot \text{int}(T') \neq \emptyset$. In this case, by the Jordan Curve Theorem there is an intersection $p \in \partial T' \cap g \cdot \partial T'$. Since $\partial T'$ consists of arcs from lifts of Γ , this violates $T' \subset T \setminus \tilde{\Gamma}$. Thus $\text{int}(T')$ embeds in S , and we have found a triangle in the complement of Γ .

Conversely, suppose Γ forms a triangle T . Lift this topological disk to \tilde{S} , and observe that the arcs of the boundary are contained in curves that pairwise intersect. By the curves to hyperplanes correspondence, there is a 3-cube corresponding to this collection. Using the correspondence again, a 4-cube would yield four lifts of curves in Γ that pairwise intersect. Since $|\Gamma| = 3$, two of these lifts are in the same $\pi_1 S$ -orbit. As each of the γ_i are simple, this is impossible. \square

Recall that a *1-system* is a collection of homotopy classes of simple closed curves whose pairwise geometric intersection number is at most one. It is interesting to note that our proof of Theorem 3 below makes essential use of the orientation of S .

Theorem 3. *Suppose $\Gamma = \{\gamma_1, \dots, \gamma_n\}$ is a 1-system on S . Then the dimension of $C(\Gamma)$ is n if and only if the dimension of $C(\Gamma')$ is three, for every triple $\Gamma' \subset \Gamma$.*

One direction is straightforward. If $\dim C(\Gamma) = n$, then the correspondence of curves to hyperplanes guarantees that there is a set of n lifts of the curves of Γ which mutually intersect. In this case, every trio of these lifts mutually intersect and form a 3-cube in the complex dual to that trio. Towards the other direction, using Lemma 5.2 we may assume that each trio forms a triangle.

We define an *almost-realization* of a curve system to be a minimal position realization (see §2) with one slight change: the words ‘only at transverse double points’ in condition (iii) should be replaced by ‘transversally.’ (The terms *pairwise minimal position realization* would be more descriptive, but less economical). Note that [7, Proposition 1.7, p. 31] still guarantees that almost-realizations minimize pairwise geometric intersection numbers, as with realizations.

Proposition 5.3. *If $\{\gamma_1, \dots, \gamma_n\}$ is a 1-system so that every trio $\{\gamma_i, \gamma_j, \gamma_k\}$ forms a triangle, then there is an almost-realization $\{\alpha_1, \dots, \alpha_n\}$ of $\{\gamma_1, \dots, \gamma_n\}$ so that the curves $\alpha_1, \dots, \alpha_n$ have a single common intersection point.*

Note that it is essential that the curves form a 1-system, as Figure 3 exhibits four curves so that every trio forms a triangle, but the conclusion of Proposition 5.3 fails. Note as well that, if the curves form a 1-system and every trio forms triangles, then each pair intersects exactly once, which we assume below. We first prove Theorem 3, assuming Proposition 5.3.

Proof of Theorem 3. By Proposition 5.3, the curves $\gamma_1, \dots, \gamma_n$ have an almost-realization as $\alpha_1, \dots, \alpha_n$, so that the α_i have a single common intersection point. Choose a lift of this point to \tilde{S} , and consider the lifts of the α_i passing through this point. Since the α_i form an almost-realization, there are no bigons formed by the curves, and the chosen lifts pairwise intersect in exactly one point. This implies that in any realization of the curves $\gamma_1, \dots, \gamma_n$ these lifts will pairwise intersect. By the curves to hyperplanes correspondence, there is an n -cube dual to these curves. \square

In the proof of Proposition 5.3, we will need a straightforward adaptation of [7, Proposition 1.7, p. 31]:

Lemma 5.4. *If $\{\alpha_1, \dots, \alpha_n\}$ is an almost-realization of a collection of simple closed curves, then for any other simple closed curve β , there is a curve β' , homotopic to β , so that $\{\alpha_1, \dots, \alpha_n, \beta'\}$ is an almost-realization.*




Figure 3. Four curves for which every trio forms a triangle, but so that there is no almost-realization with a single fourfold intersection point.

Proof. By [7, Proposition 1.7, p. 31], two curves are in minimal position if and only if they form no bigons. One half of this proof shows that if a bigon is formed [7, p. 32], one may homotope one of these curves across the bigon while the other remains fixed. Choose any realization β_0 of β . If β_0 is in minimal position with α_1 , then let $\beta_1 = \beta_0$. If β_0 is not in minimal position with α_1 , find a bigon that β_0 forms with α_1 . Use this bigon to find a curve β'_0 , homotopic to β_0 , so that $|\beta'_0 \cap \alpha_1| < |\beta_0 \cap \alpha_1|$.

It is crucial to observe that this step may be done so that $|\beta'_0 \cap \alpha_j| \leq |\beta_0 \cap \alpha_j|$, for each j . The new curve $\beta'_0 = \beta_+ \cup \beta_-$ is composed of two arcs, where β_+ follows β_0 and β_- follows α_1 . If an arc of α_j intersects β'_0 in the arc β_+ , then there is a corresponding point of intersection of α_j with β_0 . If an arc of α_j intersects β'_0 in the arc β_- , then this arc of α_j enters the bigon formed by β_0 and α_1 along the side contained in α_1 . In this case, α_j must exit the bigon through the side contained in β_0 , since α_1 and α_j form no bigons by hypothesis (see Figures 4 and 5). Thus there is again a point of intersection of β_0 with this arc of α_j that corresponds to the intersection of α_j with β'_0 . We conclude that $|\beta'_0 \cap \alpha_j| \leq |\beta_0 \cap \alpha_j|$, for $j \neq 1$.

We apply these finitely many homotopies across bigons to β_0 , one for each of the bigons formed by β_0 and α_1 . The result is a curve β_1 , homotopic to β , so that β_1 and α_1 are in minimal position, and so that $|\beta_1 \cap \alpha_j| \leq |\beta_0 \cap \alpha_j|$, for $j \neq 1$. Do this one-by-one for each α_k , and the result is a curve β' , homotopic to β , so that β is in minimal position with α_j for each $j = 1, \dots, n$. Thus $\{\alpha_1, \dots, \alpha_n, \beta'\}$ is an almost-realization. \square




Figure 4. A point in $\alpha_j \cap \beta'_0$ corresponds to a point in $\alpha_j \cap \beta_0$.




Figure 5. ‘New’ intersections of β'_0 with α_j , violating the assumption that α_1 and α_j are in minimal position.

Proof of Proposition 5.3. We proceed by induction on n . For $n = 3$, choose a realization of the three curves that forms a triangle. Note that since the curves form a 1-system, there is one boundary arc of the triangle contained in each of the curves. Fixing the curves outside of a disk containing the triangle, homotope one of the curves across the triangle so that it transversally crosses the other two curves at their intersection point. This produces an almost-realization of the curves, since no bigons have been created.

Assume now that we have $n + 1$ curves $\{\gamma_1, \dots, \gamma_{n+1}\}$ so that every trio forms a triangle. By the inductive hypothesis, there is an almost-realization $\{\alpha_1, \dots, \alpha_n\}$ of $\{\gamma_1, \dots, \gamma_n\}$ so that the α_i all have a unique common intersection point. Suppose we are given a curve in the homotopy class of γ_{n+1} that is in minimal position with the α_i . Our strategy of proof will use the intersection properties of the curves – namely the fact that each pair intersects exactly once, and that every trio forms triangles – to ‘weave’ this curve around the other curves. Note that in the process of achieving the desired arrangement, the curves α_i remain unchanged.

Lemma 5.4 implies that we may find a curve β , homotopic to γ_{n+1} , so that $\{\alpha_1, \dots, \alpha_n, \beta\}$ is an almost-realization. After possibly renaming the curves $\alpha_1, \dots, \alpha_n$, we may assume that these curves are arranged around their unique common intersection in the counter-clockwise order $\alpha_1, \dots, \alpha_n$. Let p_0 indicate the common intersection point of the α_i , and choose a disk neighborhood C of p_0 , embedded on S , so that $\alpha_i \cap C$ is connected for each $i = 1, \dots, n$.

Let $p_i := \beta \cap \alpha_i$. Replacing β with a homotopic curve if necessary, we may assume that all of the p_i are contained in the arcs $\alpha_i \cap C$, and that $\beta \cap C$ consists only of arcs that intersect an arc $\alpha_i \cap C$. We now have a picture where β weaves in and out of C , intersecting each of the arcs $\alpha_i \cap C$ precisely once. Outside of C the curves $\{\alpha_1, \dots, \alpha_n, \beta\}$ are disjoint. The argument that follows applies various homotopies to β , maintaining the fact that β and α_i are in minimal position, for each $i = 1, \dots, n$. Note that as we apply such homotopies, the p_i move as well.

In fact, for each p_i there is a homotopy of the curve β that slides the intersection p_i along α_i , out of C , until p_i “reappears” on the other side of p_0 on $\alpha_i \cap C$. Since the curves α_i do not intersect outside of C , this homotopy leaves the collection $\{\alpha_1, \dots, \alpha_n, \beta\}$, pairwise, in minimal position. This homotopy of β will be exploited in the following argument, where we will refer to it as the ‘slide move applied to p_i .’

Our goal is to apply a homotopy to β to achieve a special formation, in which intersections of β with α_i have been ‘localized’ to C . We will say a curve is a *localized realization* of β if it is homotopic to β , $\beta \cap C$ is connected, and $p_i \in C$, for each i . See Figure 6 for an illustrated example. The bulk of the proof of the inductive step concerns existence of a localized realization of β . We will proceed by analyzing the triangles formed by β with α_i and α_{i+1} , as i goes from 1 to $n-1$. At each step, we either find a localized realization of β , or we apply a homotopy to β so that the number of connected components of $\beta \cap C$ is at most $n-i$.




Figure 6. A ‘localized realization’ of β : $\beta \cap C$ is connected, and $p_i \in C$.

Given vertices a_1, a_2, a_3 of the triangle T , we will say that T ‘realizes (a_1, a_2, a_3) ’ if the counter-clockwise orientation of ∂T induces the cyclic order (a_1, a_2, a_3) . We will refer to a sufficiently small neighborhood of a_i as a ‘corner’ of T , or ‘the corner at a_i ’ when the points are distinct, where the sufficiency is fulfilled when the intersection of T with this neighborhood is connected and disjoint from the side opposite the vertex.

Suppose that two curves η and δ in minimal position, together with a third curve, form a triangle T , so that T has a corner at the intersection point $p \in \eta \cap \delta$. In this case, the complement of $\eta \cup \delta$ in a small enough neighborhood of p consists of four components. Note that these four components correspond to the four possibilities for the placement of a corner of T at p . (When η and δ are simple curves it is straightforward to check that these possibilities are mutually exclusive). This is exploited repeatedly below, for the placement of corners of triangles at p_0 .

By hypothesis, and by Lemma 5.1, there is a triangular component T_1 of $S \setminus \{\alpha_1, \alpha_2, \beta\}$. Because the curves form a 1-system, the vertices of T_1 are necessarily given by p_0, p_1 , and p_2 . Of the four ways for there to be a corner

formed by α_1 and α_2 at p_0 , two of them – namely, the choice of corner so that T realizes (p_0, p_2, p_1) (see Figure 7) – would achieve a localized realization. The composition, if necessary, of slide moves applied to p_1 and p_2 with a homotopy of β across the triangle T_2 would form a localized realization. We thus assume that T realizes (p_0, p_1, p_2) , in which case, applying slide moves to p_1 and p_2 if necessary, we have the situation pictured in Figure 8. Applying further slide moves when necessary, we may now assume that the p_i , for $i > 2$, all lie on the same side of $\alpha_1 \cap C$ inside C , as in Figure 9.

We now consider a triangular component T_2 of $S \setminus \{\alpha_2, \alpha_3, \beta\}$, with vertices p_0, p_2 , and p_3 . As before, there are four possible placements of the corner of T_2 at p_0 . Two of them correspond to a situation in which T_2 realizes (p_0, p_3, p_2) , again allowing a homotopy of β that forms a localized realization. Assuming then that T_2 realizes (p_0, p_2, p_3) , we now describe why there is only one possible placement of the corner of T_2 at p_0 .

First suppose the corner of T_2 at p_0 is placed as pictured in Figure 10. That is, suppose T_2 does not share the side between p_0 and p_2 , contained in the arc $\alpha_2 \cap C$, with T_1 . Because T_2 realizes (p_0, p_2, p_3) , the counter-clockwise orientation of ∂T_2 turns left from α_2 to β at p_2 . There are two ways to make a left turn from α_2 to β at p_2 . Of them, the one so that T_2 does not share the side between p_0 and p_2 with T_1 contains the arc of the curve β that lies between the vertices p_2 and p_3 and containing p_1 . Consequently, one side of the arc $\alpha_1 \cap C$ leaves C inside T_2 , while the other leaves C outside of T_2 . Since α_1 may only intersect the sides of T_2 inside C , this is impossible.

This leaves only one possibility for the placement of the corner of T_2 at p_0 in which T_2 realizes (p_0, p_2, p_3) (see Figure 11). This case allows us to apply a homotopy to β supported in the triangle T_2 , leaving the intersections p_i inside C , and ensuring that the number of connected components of $\beta \cap C$ is at most $n - 2$ (see Figure 12).

Similarly, at the k th step, the triangular component T_k of $S \setminus \{\alpha_k, \alpha_{k+1}, \beta\}$ either provides a homotopy of β that achieves a localized realization, or provides a homotopy of β , supported inside C , that ensures that the number of connected components of $\beta \cap C$ is at most $n - k$. When $k = n - 1$, we have ensured the existence of a localized realization of β .

Finally, after replacing β with its localized realization, we assume that $\beta \cap C$ is connected, and that $p_i \in C$ for each i . Note that this implies that $\beta \cap \partial C$ consists of exactly two points, which we denote p_- and p_+ . A straightforward application of the Jordan Curve Theorem ensures that p_- and p_+ are in diametrically opposed components of $\partial C \setminus \{\alpha_1, \dots, \alpha_n\}$, and we may apply a homotopy to β making it into a diameter passing through p_0 , completing the inductive step. \square




Figure 7. The triangle T_1 realizes (p_0, p_2, p_1) , forming a localized realization of β .




Figure 8. The triangle T_1 realizes (p_0, p_1, p_2) , where no localized realization is yet assured.




Figure 9. Given that the triangle T_1 realizes (p_0, p_1, p_2) , the arcs $\beta \cap C$ may appear as pictured.




Figure 10. Given triangle T_1 realized as (p_0, p_1, p_2) , the darkly shaded oriented triangle T_2 cannot have a corner at p_0 as pictured.




Figure 11. The remaining case where triangle T_2 realizes (p_0, p_2, p_3) .




Figure 12. Applying a homotopy to β using T_2 , getting ‘closer’ to a localized realization.

A word of caution. It is not generally true that if $\Gamma' \subset \Gamma$ then the complex $C(\Gamma')$ is a subcomplex of $C(\Gamma)$ (see Figure 13). The following corollary is a weaker version of such a statement that will suffice for our application. Given a collection of curves Γ , we will say that a subset of n curves $\Gamma' \subset \Gamma$ form an n -cube in $C(\Gamma)$ if there are n hyperplanes corresponding to the curves of Γ' intersecting in an n -cube of $C(\Gamma)$.




Figure 13. The complex $C(\{\alpha, \gamma\})$ is not a subcomplex of $C(\{\alpha, \beta, \gamma\})$.

Corollary 5.5. *If Γ is a 1-system of curves, and $\{\gamma_1, \dots, \gamma_n\} = \Gamma' \subset \Gamma$, then the curves of Γ' form an n -cube in $C(\Gamma)$ if and only if every triple of curves from Γ' form a 3-cube.*

Proof. Using the curves to hyperplanes correspondence, the curves of Γ' form an n -cube in $C(\Gamma)$ if and only if there is a choice of lifts $\{\tilde{\gamma}_1, \dots, \tilde{\gamma}_n\}$ so that these lifts pairwise intersect, which in turn occurs if and only if $\dim C(\Gamma') = n$, at which point we apply Theorem 3. □

6. A family of maximum complete 1-systems

We now construct many maximum complete 1-systems on a surface S of any odd genus $g = 2n + 1$. Consider the genus g surface as the unit sphere in \mathbb{R}^3 with g handles attached at evenly spaced disks centered on an equator, making the order g homeomorphism σ that cyclically permutes the handles apparent. Let x be a point fixed by this homeomorphism. Consider the presentation

$$\pi_1(S, x) = \left\langle r_i, s_i \mid \prod_{i=1}^g [r_i, s_i] \right\rangle,$$

where the generators r_1 and s_1 are as pictured in 14, and $r_i = \sigma(r_{i-1})$ and $s_i = \sigma(s_{i-1})$ for $i = 2, \dots, g$.




Figure 14. The generators of $\pi_1(S, x)$ in one handle.

Let $\alpha_1, \beta_1,$ and δ be given by

- (1) $\alpha_1 = [r_{n+2}r_{n+3} \dots r_{2n+1}s_1]$,
- (2) $\beta_1 = [r_{n+2}r_{n+3} \dots r_{2n+1}s_1^{-1} \prod_{i=1}^{n+1} [s_i, r_i]]$, and
- (3) $\delta = [r_1r_2 \dots r_{2n+1}]$.

The orbit of α_1 under σ gives g curves which we denote by $\{\alpha_1, \alpha_3, \dots, \alpha_{2g-1}\}$. Similarly, we denote the σ -orbit of β_1 by $\{\beta_1, \beta_3, \dots, \beta_{2g-1}\}$. We complete these collections to sequences $\{\alpha_i\}_{i=1}^{2g}$ and $\{\beta_i\}_{i=1}^{2g}$ by defining $\alpha_{2i} = \tau_i(\alpha_{2i-1})$ and $\beta_{2i} = \tau_i^{-1}(\beta_{2i-1})$ for $i = 1, \dots, g$, and where τ_i is the right Dehn twist around r_i . See Figures 15, 16, 17, and 18 for illustrative examples.

Grouping these curves together, we will refer to $A = \{\alpha_1, \dots, \alpha_{2g}\}$ as the set of ‘up’ curves and $B = \{\beta_1, \dots, \beta_{2g}\}$ as the set of ‘down’ curves. We will refer to the pair of up curves (resp. down curves) α_{2i-1} and α_{2i} (resp. β_{2i-1} and β_{2i}) as ‘partners’, for $i = 1, \dots, g$.

It is immediate that $\alpha_i \cap \alpha_j = \beta_i \cap \beta_j = 1$ for $i \neq j$. When α_i and α_j are not partners, then $\alpha_i \cap \beta_j = 1$. This calculation makes it clear that we may form many maximum complete 1-systems. For each $i = 1, \dots, g$, choose one of the two pairs of partners from $\{\alpha_{2i-1}, \alpha_{2i}\}$ and $\{\beta_{2i-1}, \beta_{2i}\}$. Together with the δ curve, this forms $2g + 1$ curves that pairwise intersect once. This is maximal by [10, Theorem 1.4].

More precisely, for $\epsilon \in \{1, -1\}^g$, let $A(\epsilon) = \{\alpha_{2i-1}, \alpha_{2i} \mid \epsilon_i = 1\}$ and $B(\epsilon) = \{\beta_{2i-1}, \beta_{2i} \mid \epsilon_i = -1\}$, and let $\Gamma(\epsilon) = \{\delta\} \cup A(\epsilon) \cup B(\epsilon)$. Several examples in genus 5 and 7 are shown in Figures 19, 20, 21, and 22.




Figure 15. The curves $\alpha_1, \alpha_2 \in A$




Figure 16. The curves $\beta_1, \beta_2 \in B$




Figure 17. The curves $\alpha_3, \alpha_4 \in A$ and $\beta_5, \beta_6 \in B$

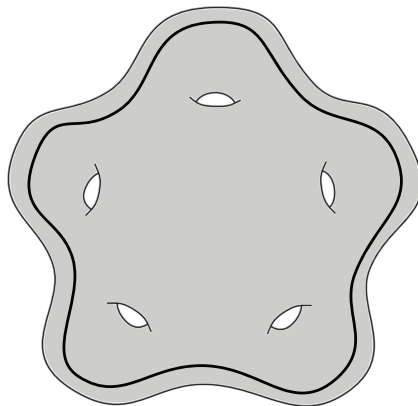


Figure 18. The curve δ




Figure 19. The system of curves $\Gamma(1, 1, 1, 1, 1)$




Figure 20. The system of curves $\Gamma(1, -1, 1, 1, -1)$




Figure 21. The system of curves $\Gamma(1, 1, 1, 1, 1, 1, 1)$




Figure 22. The system of curves $\Gamma(1, -1, 1, -1, -1, 1, -1)$

Lemma 6.1. *For each $\epsilon \in \{-1, 1\}^g$, the collection of curves $\Gamma(\epsilon)$ forms a maximum complete 1-system.*

There is an action of the dihedral group D_g on $\{1, -1\}^g$ given by letting the generators act by a g -cycle and a reversal of the list, respectively. Letting $\mathbb{Z}/2\mathbb{Z}$ act by taking ϵ to $-\epsilon$, we obtain an action of $\mathbb{Z}/2\mathbb{Z} \oplus D_g$ on $\{1, -1\}^g$, and there is naturally an action of $\text{Mod}^*(S)$ on systems of conjugacy classes of curves on S . The following proposition, whose proof occupies the bulk of our analysis in §7 and §8, implies that the maximum complete 1-systems from Lemma 6.1 represent many distinct orbits.

Proposition 6.2. *If $\Gamma(\epsilon)$ and $\Gamma(\epsilon')$ are in the same $\text{Mod}^*(S)$ -orbit, then ϵ and ϵ' are in the same $(\mathbb{Z}/2\mathbb{Z} \oplus D_g)$ -orbit in $\{1, -1\}^g$.*

We will also require the simple observation:

Lemma 6.3. *A maximum complete 1-system is filling.*

Proof. Suppose not. Then there is a simple closed curve α disjoint from the curves in our 1-system. Cut open S along α , and cap off the two resulting boundary components created with disks. Note that the resulting surface may be disconnected. In any case, the set of $2g + 1$ curves obtained forms a maximum complete 1-system of curves on a surface of genus $g' \leq g - 1$, contradicting [10, Theorem 1.4]. \square

In fact, one can show that any $2g$ curves from a maximum complete 1-system are filling, but we will not require this stronger statement.

Let $N(g)$ indicate the number of $\text{Mod}^*(S)$ -orbits among maximum complete 1-systems. A simple argument involving the square complex dual to a realization of curves on S provides an upper bound for $N(g)$ below. For g odd, Proposition 6.2 allows using the $\Gamma(\epsilon)$ to provide a lower bound for $N(g)$. For g even, we will use the stabilization procedure described in detail in §9 to obtain lower bounds for $N(g)$. Given a choice of realization for a maximum complete 1-system λ , and an arc α intersecting each of the curves in λ once, stabilization produces a maximum complete 1-system on a surface of genus $g + 1$.

When necessary below, we identify $\Gamma(\epsilon)$ with fixed choices of realization for each such collection. In §9 we prove:

Proposition 6.4. *There is a choice of arc α so that the stabilizations of $\Gamma(\epsilon)$ and $\Gamma(\epsilon')$ along α are in the same $\text{Mod}^*(S_{g+1})$ -orbit if and only if $\Gamma(\epsilon)$ and $\Gamma(\epsilon')$ are in the same $\text{Mod}^*(S_g)$ -orbit.*

We restate and prove Theorem 1:

Theorem 1. *We have the bounds*

$$(4g^2 + 2g)! \geq N(g) \geq \frac{2^{g-1}}{4(g-1)}.$$

Proof. The lower bound follows from Proposition 6.2 when g is odd, and from Proposition 6.4 when g is even.

Towards the upper bound, consider the set of isomorphism classes of square complexes \mathcal{S}_λ that are dual to realizations λ of maximum complete 1-systems. For filling systems of curves, the dual square complex is isomorphic to the surface S , and the hyperplanes of the square complex are in the homotopy classes of the curves one started with. Lemma 6.3 now guarantees that an isomorphism of square complexes $\mathcal{S}_\lambda \cong \mathcal{S}_{\lambda'}$ yields a homeomorphism of S taking λ to λ' . Thus there is a well-defined map from the set of isomorphism classes of square complexes dual to maximum complete 1-systems to the set of $\text{Mod}^*(S)$ -orbits of maximum complete 1-systems. This map is evidently surjective, so that an upper bound for the number of possible square complexes dual to a realization of a maximum complete 1-system produces an upper bound for $N(g)$.

For each realization λ of a maximum complete 1-system, each of the curves in λ passes through exactly $2g$ squares of \mathcal{S}_λ . We may thus view \mathcal{S}_λ as the quotient of the disjoint union of $2g + 1$ annuli, each of which is built from $2g$ squares, where the quotient map identifies squares in pairs. There are at most $\binom{2g(2g+1)}{2, \dots, 2}$ pairings, and each pair of matched squares has two possible identifications, giving at most

$$2^{g(2g+1)} \binom{2g(2g+1)}{2, \dots, 2} = 2^{g(2g+1)} \cdot \frac{(2g(2g+1))!}{2^{g(2g+1)}} = (2g(2g+1))!$$

square complexes \mathcal{S}_λ . □

7. Restricting mapping class group orbits via $C(\Gamma(\epsilon))$

This section is the first step towards the proof of Proposition 6.2.

Proposition 7.1. *If $\phi \in \text{Mod}^*(S)$ satisfies $\phi \cdot \Gamma(\epsilon) = \Gamma(\epsilon')$, then*

- (1) $\phi \cdot \delta = \delta$, and
- (2) $\phi \cdot \{A(\epsilon), B(\epsilon)\} = \{A(\epsilon'), B(\epsilon')\}$.

Moreover, the map ϕ sends partner curves to partner curves.

In other words, either ϕ preserves the sets of up and down curves, or it exchanges them. The proof of this proposition will follow from a coarse picture of the cube complex $C(\Gamma(\epsilon))$.

Lemma 7.2. *In the complex $C(A \cup B \cup \{\delta\})$, the triples that form 3-cubes are the following:*

- (1) $\{\alpha_i, \alpha_j, \alpha_k\}$ or $\{\beta_i, \beta_j, \beta_k\}$, for distinct $i, j, k \in \{1, \dots, 2g\}$;
- (2) $\{\alpha_{2i-1}, \alpha_{2i}, \delta\}$ or $\{\beta_{2i-1}, \beta_{2i}, \delta\}$, for $i \in \{1, \dots, g\}$;
- (3) $\{\alpha_{2j-1}, \alpha_{2j}, \beta_i\}$ or $\{\beta_{2j-1}, \beta_{2j}, \alpha_i\}$, for $i \in \{1, \dots, 2g\}$ and $j \in \{1, \dots, g\}$;
- (4) $\{\alpha_i, \beta_j, \delta\}$ for $i, j \in \{1, \dots, 2g\}$.

Proof. Using Lemmas 5.1 and 5.2, we determine whether a triple of curves forms a 3-cube by choosing a realization of the curves, and observing whether there is a triangular component of the complement. For each of the curves, we fix choices of realizations as in Figures 15, 16, and 18.

If δ is one of the three curves, we arrange the possible ways to choose the other two curves according to whether the curves are chosen as ‘up’ or ‘down’ (i.e. from A or B). If both of the other curves are up, then there is a triangle in the complement of the trio if and only if the other two curves were partners. If one of the curves is up and one is down, there is such a triangle. The other cases are similar.

On the other hand, if δ is not one of the three curves, if all of the curves are up, there is such a triangle. If two of the curves are up and one is down, there is a triangle in their complement if and only if the two up curves are partners. The other cases are similar. \square

We proceed with an examination of hyperplanes of maximal cubes in the cases for $\epsilon \in \{1, -1\}^g$ where $|A(\epsilon)|, |B(\epsilon)| > 2$.

Lemma 7.3. *When $|A(\epsilon)|, |B(\epsilon)| > 2$, the sets of hyperplanes of maximal cubes of $C(\Gamma(\epsilon))$ correspond to one of the following lists of curves:*

- (1) the $2|\epsilon^{-1}(1)|$ curves $A(\epsilon)$;
- (2) the $2|\epsilon^{-1}(-1)|$ curves $B(\epsilon)$;
- (3) the 5 curves $\{\alpha_{2i-1}, \alpha_{2i}, \beta_{2j-1}, \beta_{2j}, \delta\}$, for $i, j \in \{1, \dots, g\}$ such that $\alpha_{2i} \in A(\epsilon)$ and $\beta_{2j} \in B(\epsilon)$.

Proof. Using Corollary 5.5, in order to check whether a subset of curves from $\Gamma(\epsilon)$ forms an n -cube, it is enough to check whether every triple forms a 3-cube. By Lemma 7.2 we have a complete list of such 3-cubes.

The curves $A(\epsilon)$ and $B(\epsilon)$ form cubes of dimensions $2|\epsilon^{-1}(1)|$ and $2|\epsilon^{-1}(-1)|$, respectively, by Lemma 7.2 and Corollary 5.5. If one adds a down curve β_i to $A(\epsilon)$, then a pair of up curves that are not partners will not form a 3-cube with this down curve β_i , by Lemma 7.2. If one adds δ to $A(\epsilon)$, then again a pair of up curves that are not partners will not form a 3-cube with δ . The analogous statements hold for $B(\epsilon)$. By Corollary 5.5, the cubes of dimension $2|\epsilon^{-1}(1)|$ and $2|\epsilon^{-1}(-1)|$

containing these sets of hyperplanes, respectively, must be maximal. The same analysis shows that a maximal cube containing δ must contain a pair of partner up curves and a pair of partner down curves. \square

The cases in which either of $|A(\epsilon)|$ or $|B(\epsilon)|$ are less than or equal to 2 are quite similar, so we list the relevant information without proof.

Lemma 7.4. *When $|B(\epsilon)| = 0$ (resp. $|A(\epsilon)| = 0$), the sets of hyperplanes of maximal cubes of $C(\Gamma(\epsilon))$ correspond to one of the following lists of curves:*

- (1) *the $2g$ curves $A(\epsilon)$ (resp. $B(\epsilon)$);*
- (2) *the 3 curves $\{\alpha_{2i-1}, \alpha_{2i}, \delta\}$ (resp. $\{\beta_{2i-1}, \beta_{2i}, \delta\}$), for $i \in \{1, \dots, g\}$.*

When $|B(\epsilon)| = 2$ (resp. $|A(\epsilon)| = 2$), the sets of hyperplanes of maximal cubes of $C(\Gamma(\epsilon))$ correspond to one of the following lists of curves:

- (1) *the $2g - 2$ curves $A(\epsilon)$ (resp. $B(\epsilon)$);*
- (2) *the 5 curves $\{\alpha_{2i-1}, \alpha_{2i}, \beta_{2j-1}, \beta_{2j}, \delta\}$, for $i, j \in \{1, \dots, g\}$ such that $\alpha_{2i} \in A(\epsilon)$ and $\beta_{2j} \in B(\epsilon)$.*

Proof of Proposition 7.1. The simple observation we exploit is that cube complex isomorphisms must send maximal cubes to maximal cubes.

Suppose $|A(\epsilon)|, |B(\epsilon)| > 2$. By Lemma 7.3 and Lemma 7.4, the maximal cubes that the hyperplane corresponding to δ passes through are all 5-dimensional, while the maximal cubes that the hyperplane corresponding to α_i (resp. β_i) passes through, for any $i \in \{1, \dots, 2g\}$, include one of the two even-dimensional maximal cubes corresponding to $A(\epsilon)$ and $B(\epsilon)$. Thus any isomorphism of cube complexes $\Phi: C(\Gamma(\epsilon)) \cong C(\Gamma(\epsilon'))$ must take the hyperplane corresponding to δ to itself. By Theorem 2, the corresponding mapping class ϕ fixes δ . Similarly, the even-dimensional maximal cubes whose hyperplanes correspond to $A(\epsilon)$ and $B(\epsilon)$ must be sent to the pair of even-dimensional maximal cubes whose hyperplanes correspond to $A(\epsilon')$ and $B(\epsilon')$. The remaining cases are similar.

Finally, a pair of partner curves are simultaneously up or down. By Lemma 7.2, they form a 3-cube with δ while a pair of non-partner curves that are simultaneously up or down do not. It follows that ϕ sends partner curves to partner curves. \square

While we may conclude that the pair of numbers $|A(\epsilon)|$ and $|B(\epsilon)|$ is equal to $|A(\epsilon')|$ and $|B(\epsilon')|$ if $\Gamma(\epsilon)$ and $\Gamma(\epsilon')$ are $\text{Mod}^*(S)$ -equivalent, we are not yet able to prove Proposition 6.2. We turn to finer invariants of $\Gamma(\epsilon)$.

8. The labeled polygon $P(\epsilon)$ associated to $\Gamma(\epsilon)$

Towards the proof of Proposition 6.2, we introduce a more detailed invariant. The information inherent to the invariant we produce is easily packaged as a polygon. The essential tool to building this invariant is the ordering induced on intersection points of a realization of an oriented curve.

However, while an oriented curve in a realization of a curve system determines an ordering of its intersection points, this ordering may not be an invariant of the collection of homotopy classes; the presence of 3-cubes implies the existence of a Reidemeister move that will make it possible to switch the ordering of intersection points. We state the following only for complete 1-systems for ease in exposition, but with slightly more detail a more general statement could be made.

Lemma 8.1. *Let $\{\gamma_1, \dots, \gamma_k, \gamma, \gamma'\}$ be a complete 1-system of curves, and let $\vec{\gamma}$ be a choice of orientation of γ . Suppose that γ does not form 3-cubes with any pair γ_i and γ_j . Then the cyclic ordering of $\gamma_1, \dots, \gamma_k$ induced by $\vec{\gamma}$ is invariant of the choice of realization, and the choice of a ‘first curve’ γ_j induces a well-defined ordering of $\{\gamma_i: i \neq j\}$. Moreover, if γ and γ' form 3-cubes with γ_i , for all i , then the cyclic orderings of $\gamma_1, \dots, \gamma_k$ induced by the two orientations of γ' coincide with those of the two orientations of γ .*

Proof. Choose a realization of the curve system $\{\gamma_1, \dots, \gamma_k, \gamma, \gamma'\}$. One obtains a cyclic ordering of $\{\gamma_1, \dots, \gamma_k\}$ induced by $\vec{\gamma}$. This ordering is invariant of the chosen realization, since any realization can be obtained from any other realization by applying a sequence of Reidemeister moves, and births or deaths of monogons or bigons [8, Lemma 5.6].

Finally, suppose γ and γ' intersect at p and form 3-cubes with each γ_i , with triangles indicated by $\Delta(\gamma, \gamma', \gamma_i)$. In this case, an argument similar to the inductive step in the proof of Lemma 5.3 guarantees that the triangles $\Delta(\gamma, \gamma', \gamma_i)$ must determine corners that are either identical or opposite at p (see Figure 23). Applying Reidemeister moves if necessary, we may arrange for all the corners to be identical at p , in which case the fact that $\gamma, \gamma_i, \gamma_j$ don't form triangles guarantees the conclusion. \square




Figure 23. The corner of the darkly shaded triangle $\Delta(\gamma, \gamma', \gamma_j)$ at p cannot be as pictured with respect to the corner of $\Delta(\gamma, \gamma', \gamma_i)$ at p .

We refer to maximal sets of non-partner up (resp. down) curves as *full*, and we fix choices of full sets of up and down curves U and D , by choosing one curve in each pair of partner curves. In our setting, by Lemma 7.2 the curve α_i does not form 3-cubes with any pair of non-partner down curves. Thus, choosing an orientation $\vec{\alpha}_i$, we may apply Lemma 8.1 to $\vec{\alpha}_i$ with D . Note that this cyclic ordering of D is unchanged by replacing α_i with its partner curve by Lemma 8.1, since a pair of partner up curves form 3-cubes with any down curve by Lemma 7.2. We conclude that a choice of orientation $\vec{\alpha}_i$ induces a cyclic order to any full set of down curves.

Moreover, as long as $|A(\epsilon)| > 2$, the cyclic order of D induced by $\vec{\alpha}_i$ can be upgraded, canonically, into a bona fide ordering. If $|A(\epsilon)| > 2$, there is an up curve α_j which is not a partner of α_i . By Lemma 7.2, the curve α_i does not form a 3-cube with α_j and β_l , for any l , so that Lemma 8.1 applies to $\vec{\alpha}_i$ and the union of D with α_j . Moreover, it is evident that this order does not depend on the choice of α_j among up curves that are not partners of α_i . We conclude that a choice of orientation $\vec{\alpha}_i$ induces an ordering of D , which we refer to as the $\vec{\alpha}_i$ -ordering of D .

We choose an almost realization of $\Gamma(\epsilon)$ such that the curves in U intersect at a single point p , and curves in D intersect at a single point p' . Let c be a small circle centered at p , and c' a small circle centered at p' . An orientation of c (resp. c') induces a cyclic ordering on the finite set of points $U \cap c$ (resp. $D \cap c'$). Assuming that both U and D are non-empty, using the orientation of the surface S , we equip c with a clockwise orientation, and c' with a counter-clockwise orientation. This induces a cyclic ordering on the finite collection of points $U \cap c$ and on $D \cap c'$.

Remark 8.2. We note the pair given by the cyclic ordering on $U \cap c$ and its reverse (resp. for $D \cap c'$) is an invariant of the $\text{Mod}^*(S)$ -orbit of $\Gamma(\epsilon)$. This follows from the fact that the cyclic order on $U \cap c$ corresponds to the cyclic order on the set of endpoints of lifts to \mathbb{H}^2 of geodesics in U induced by an orientation of $\partial\mathbb{H}^2$ (with respect to any hyperbolic structure on S), and therefore this ordering is detected by the dual cube complex to $\Gamma(\epsilon)$. In particular, the definition of this cyclic ordering on $U \cap c$ does not depend on our particular choice of almost realization for $\Gamma(\epsilon)$.

Given a choice of orientation of a curve $\gamma \in U$ (resp. D), the two points of $\gamma \cap c$ (resp. $\gamma \cap c'$) are each equipped with an arrow that either points away from, or towards p (resp. p'). We refer to arrows pointing towards p (resp. p') as *inward pointing* and the others as *outward pointing*. Given a set of choices of orientations for each of the curves in U , the set of inward and outward pointing arrows partitions the set of points $U \cap c$ into two sets. Moreover, there is an involution ι of $U \cap c$ that exchanges these two sets, for any set of choices of orientations for the curves in U : given $v \in U \cap c$, $\iota(v)$ is the only other point of $U \cap c$ on the same curve of $\Gamma(\epsilon)$ as v .

Lemma 8.3. *There exist choices of orientations $\{\vec{\alpha}_i\}$ for the curves in U such that the cyclic orderings of D induced by $\vec{\alpha}_i$ and $\vec{\alpha}_j$ are cyclically equivalent, for all i and j .*

Proof. Choose orientations for α_i and α_j as in Figure 24. The down curves D are partitioned by this choice into four sets D_1, D_2, D_3, D_4 , and with the chosen orientations on α_i and α_j , the $\vec{\alpha}_i$ - and the $\vec{\alpha}_j$ -orderings of D are cyclically equivalent. \square




Figure 24. The oriented up curve u_1 induces the ordering (D_1, D_2, D_3, D_4) of the down curves, while u_2 induces (D_4, D_1, D_2, D_3) .

A set of choices of orientations of the curves in U is *coherent* if it satisfies the conclusion of Lemma 8.3. The orientations chosen in Lemma 8.3 (see Figure 24) are a convenient choice, and we refer to these choices of orientations of up curves as the *standard orientations* for curves in U . It will be useful to have a chosen orientation of the down curves as well. Note the orientation-reversing involution of the surface that exchanges each up curve α_i with the down curve β_i . An orientation of a down curve is *standard* if the image under this involution is a standardly oriented up curve.

Remark 8.4. We note that there are exactly two coherent orientations on U ; this follows from the fact that a coherent orientation is completely determined by choosing the orientation on one curve of U . We will refer to the coherent set of choices of orientations for U that is not the standard one as *non-standard*.

Definition 8.5. The *labeled polygon* $P(\epsilon)$ associated to $\Gamma(\epsilon)$ is a $2|U|$ -gon, satisfying the following conditions.

- (1) The vertices are labeled from elements of $U \cap c$, in the cyclic ordering induced by the orientation of c ; thus $P(\epsilon)$ comes equipped with a preferred orientation of its boundary. By placing an outward pointing arrow at each point of $U \cap c$, the vertices of $P(\epsilon)$ determine orderings of D . By Remark 8.4, these orderings partition the vertices into two sets P_1 and P_2 , exchanged by ι , that correspond to the two cyclic equivalence classes of orderings of D determined by possible coherent choices of orientations for U . For each $l = 1, 2$, we let $R_l(\epsilon)$ denote the polygon formed by the cyclically ordered vertices in P_l . Note that the edges of $R_l(\epsilon)$ are diagonals of $P(\epsilon)$;
- (2) Each edge of $R_l(\epsilon)$ is decorated with a pair of integers $\mathcal{M}_l(e)$ and $\mathcal{N}_l(e)$, defined as follows. Let $e = (\tilde{\alpha}_i, \tilde{\alpha}_j)$ be an edge of $R_l(\epsilon)$ with initial and terminal vertices corresponding to the oriented curves $\tilde{\alpha}_i$ and $\tilde{\alpha}_j$ in U , respectively. We set $\mathcal{M}_l(e)$ equal to the number of vertices of $P(\epsilon)$ between $\tilde{\alpha}_i$ and $\tilde{\alpha}_j$. If $\beta_{l_1}, \dots, \beta_{l_m}$ denotes the $\tilde{\alpha}_i$ -ordering of D , then by construction there exists some r such that the $\tilde{\alpha}_j$ -ordering on D is

$$\beta_{l_m-r+1}, \dots, \beta_{l_m}, \beta_{l_1}, \dots, \beta_{l_m-r}.$$

We define $\mathcal{N}_l(e) = r$.

Lemma 8.6. $P(\epsilon)$ is a well-defined invariant of $\Gamma(\epsilon)$.

Proof. By Remark 8.4, the standard and non-standard pair of choices of orientations for the curves in U are the only coherent such choices. Thus the partition of the vertices of $P(\epsilon)$ into the two sets P_1 and P_2 is well-defined. The labels $\mathcal{M}_l(e)$ and $\mathcal{N}_l(e)$ are evidently invariant under a different set of choices of full sets U and D . □

Lemma 8.7. For each of the labeled polygons $P(\epsilon)$ there exists a choice of l_1 and l_2 so that $\{l_1, l_2\} = \{1, 2\}$ and

$$\sum_e \mathcal{N}_{l_1}(e) = |D|$$

and

$$\sum_e \mathcal{N}_{l_2}(e) = |D| \cdot (|U| - 1).$$

Proof. One of the polygons has vertices that determine the standard orientations of the curves in U , while the other has vertices that determine the non-standard ones. By examining Figure 24, one sees that the sum for the polygon with standard orientations determined at its vertices is the first sum above, while the sum for the other polygon is the second sum above. \square

After relabeling, we assume from now on that $R_1(\epsilon)$ refers to the polygon whose vertices are points of $U \cap c$ with outward pointing arrows in the standard orientations of the curves in U . See Figure 25 for examples corresponding to $\epsilon = (1, 1, 1, -1, -1, 1, -1)$ and $\epsilon = (-1, 1, 1, -1, -1, 1, 1)$.

Definition 8.8. An *isomorphism of labeled polygons* from $P(\epsilon)$ to $P(\epsilon')$ is a permutation $\psi \in S_{2|U|}$, where S_n denotes the symmetric group on n symbols, that induces a label-preserving isomorphism between the labeled 1-skeletons of $P(\epsilon) \cup R_1(\epsilon) \cup R_2(\epsilon)$ and $P(\epsilon') \cup R_1(\epsilon') \cup R_2(\epsilon')$.


By Proposition 7.1, a homeomorphism of the surface that takes $\Gamma(\epsilon)$ to $\Gamma(\epsilon')$ takes $A(\epsilon)$ to either $A(\epsilon')$ or $B(\epsilon')$. In the first case, by construction, the homeomorphism induces an isomorphism of labeled polygons $P(\epsilon) \cong P(\epsilon')$. Note that it is a consequence of Lemma 8.7 that we may conclude that, if $\psi: P(\epsilon) \cong P(\epsilon')$ is an isomorphism of labeled polygons, then ψ takes $R_1(\epsilon)$ to $R_1(\epsilon')$ and $R_2(\epsilon)$ to $R_2(\epsilon')$.

Remark 8.9. As a consequence, any homeomorphism taking $\Gamma(\epsilon)$ to $\Gamma(\epsilon')$ that takes $A(\epsilon)$ to $A(\epsilon')$ must preserve the standard orientations of curves.


Proposition 8.10. *If $\psi: P(\epsilon) \cong P(\epsilon')$ is an isomorphism of labeled polygons, then ϵ and ϵ' are in the same $(\mathbb{Z}/2\mathbb{Z} \oplus D_g)$ -orbit of $\{-1, 1\}^g$.*

Proof. Our strategy below is to make various choices for labels and orientations of curves that make the structure of $P(\epsilon)$ more transparent – by Lemma 8.6 these choices don't affect the isomorphism type of $P(\epsilon)$. With the structure of $P(\epsilon)$ clear, the result will follow easily. Roughly speaking, the constructions below are careful treatments of ‘how the curves in $\Gamma(\epsilon)$ look,’ when they are drawn conveniently on the surface. In what follows, we fix the almost realization described at the beginning of this section, and fix as well the standard orientations of each of the curves.

We first establish a cyclic ordering on all the curves in $U \cup D$. Consider the collection of curves $\omega_1, \dots, \omega_g$ in Figure 26, realized in minimal position with $\Gamma(\epsilon)$. Note that for each i , the curve ω_i intersects exactly one curve in $U \cup D$. Given $\gamma \in U \cup D$, the *transversal* to γ is the unique curve in $\{\omega_1, \dots, \omega_g\}$ intersecting γ . Thus there is a cyclic ordering on $U \cup D$ associated to the cyclic ordering $[1, \dots, g]$ on the indices of the ω_i curves.



$$\epsilon = (1, 1, 1, -1, -1, 1, -1)$$



$$\epsilon = (-1, 1, 1, -1, -1, 1, 1)$$

Figure 25. Full sets of curves from the curve system $\Gamma(\epsilon)$, the ‘small circle’ c , the polygon $P(\epsilon)$, and the polygon $R_1(\epsilon)$. Edge labels are written as the ordered pair $(\mathcal{M}_1(e), \mathcal{N}_1(e))$.




Figure 26. The transversals $\omega_1, \dots, \omega_g$.

We will now describe a cyclic order on $(U \cap c) \sqcup (D \cap c')$. For $\gamma \in U \cup D$, suppose $\gamma \in U$, and let ω denote the transversal to γ . Then c separates γ into two arcs, both bounded by the two points $s_1(\gamma), s_2(\gamma)$ of $\gamma \cap c \subset U \cap c$. Let $\Lambda(\gamma)$ denote the arc whose interior is not contained in the disk bounded by c . Then $\Lambda(\gamma)$ is subdivided further into two sub-arcs, which we denote $S_1(\gamma)$ and $S_2(\gamma)$; for each $i = 1, 2$, the arc $S_i(\gamma)$ is the sub-arc of $\Lambda(\gamma)$ bounded by $s_i(\gamma)$ and by $\gamma \cap \omega$ (see Figure 27). We remark that the arcs $S_1(\gamma)$ and $S_2(\gamma)$ are distinguished by the choice that $S_1(\gamma)$ contains none of the intersection points with D in the chosen realization of $\Gamma(\epsilon)$. Similarly, if $\gamma \in D$ then $\Lambda(\gamma)$ is the arc of γ not contained within c' , and $\Lambda(\gamma)$ is subdivided into $S_1(\gamma)$ and $S_2(\gamma)$, where $S_1(\gamma)$ is the sub-arc of $\Lambda(\gamma)$ containing no intersections with U .

Starting at some point v of $U \cap c$, let $\gamma \in U \cup D$ denote the up curve associated to v , and let ω_i denote the transversal to γ . The point of $(U \cap c) \sqcup (D \cap c')$ immediately following v is obtained as follows. We assume that $v = s_1(\gamma)$. Let γ' be the curve obtained from γ by moving $(g+1)/2$ around the cyclic order on $U \cup D$. Thus the transversal to γ' is $\omega_{i+(g+1)/2}$, where the addition is interpreted modulo g . Then the point of $(U \cap c) \sqcup (D \cap c')$ immediately following v is $s_2(\gamma')$.

The next point is obtained in a similar fashion. Starting at $s_2(\gamma')$ we move to the curve whose index in the cyclic order on $U \cup D$ is $(g+1)/2$ from that of γ' , in the clockwise direction. Letting γ'' denote this curve, the next point in the cyclic order on $(U \cap c) \sqcup (D \cap c')$ is $s_1(\gamma'')$. The ordering is defined by iterating this procedure: add $(g+1)/2$ to the cyclic index, and alternate between s_1 and s_2 . See Figure 28.




Figure 27. The intersections $s_1(\gamma)$ and $s_2(\gamma)$ of $\gamma \in U$ with the small circle c , and the sub-arc $S_1(\gamma)$ of γ . (The arc $S_2(\gamma)$ is not pictured).




Figure 28. In the cyclic order on $(U \cap c) \sqcup (D \cap c')$, $s_2(\gamma')$ immediately follows $s_1(\gamma)$, and $s_1(\gamma'')$ immediately follows $s_2(\gamma')$.

Thus far, we have established a cyclic ordering on $U \cup D$, and an associated cyclic ordering on $(U \cap c) \sqcup (D \cap c)$; we also recall that the curves in U and D are equipped with standard orientations. Note that these choices are compatible in the following sense. Recall that the standard orientations of the curves induce inward and outward pointing arrows on the points in $(U \cap c) \sqcup (D \cap c)$. In the established cyclic ordering of $(U \cap c) \sqcup (D \cap c)$, each inward pointing arrow is followed by an outward pointing arrow, and likewise each outward pointing arrow is followed by an inward pointing arrow.

Lemma 8.11. *For an edge $e = e(\vec{\alpha}_i, \vec{\alpha}_j)$ of $R_1(\epsilon)$, the integer $\mathcal{M}_1(e)$ (resp. $\mathcal{N}_1(e)$) is equal to the number of outwardly pointing arrows of U (resp. D) which follow $\iota(\vec{\alpha}_i)$ and which precede $\iota(\vec{\alpha}_j)$.*

Proof. With the given choices made, the integer $\mathcal{M}_1(e)$ is equal to the number of inward pointing arrows that follow $\vec{\alpha}_i$ and precede $\vec{\alpha}_j$. Apply the involution ι and the claim follows for $\mathcal{M}_1(e)$. For $\mathcal{N}_1(e)$, the edge e partitions the down curves into four pieces D_1, D_2, D_3, D_4 , as in Figure 24. The label $\mathcal{N}_1(e)$ of the edge e is equal to $|D_4|$, which is also the number of outwardly pointing arrows between $\iota(\vec{\alpha}_i)$ and $\iota(\vec{\alpha}_j)$. \triangle

Henceforth, by an *interval* of $(U \cap c) \sqcup (D \cap c')$, we mean the subset which follows a particular point in $(U \cap c) \sqcup (D \cap c')$ and which precedes some other point, as in the statement of Lemma 8.11. Evidently, given an interval I , $\iota(I)$ is an interval bounded by the image of the end points of I under ι .

To finish the proof of Proposition 8.10, it suffices to show that the orbit of ϵ can be constructed from information about $P(\epsilon)$ which is preserved under isomorphism of labeled polygons. The number of 1's in ϵ is equal to $|A(\epsilon)|/2$, which is equal to half of the number of edges of $P(\epsilon)$. Thus all that remains is to determine how to interleave the $|B(\epsilon)|/2$ necessary -1 's to obtain ϵ . The following lemma completes the proof of Proposition 8.10. \square

Lemma 8.12. *If $e = (\vec{\alpha}_i, \vec{\alpha}_j)$ is an edge of $R_1(\epsilon)$, then the number of -1 's between ϵ_i and ϵ_j is equal to $\mathcal{M}_1(e) + \mathcal{N}_1(e) - 1$.*

Proof. Using the standard orientations for the curves in $U \cup D$ again, the consecutive elements of $(U \cap c) \sqcup (D \cap c')$ are equipped with opposite pointing arrows. Each -1 between the 1's associated to $\vec{\alpha}_i$ and $\vec{\alpha}_j$ corresponds to a down curve whose outwardly-pointing arrow lies between the outwardly-pointing arrows for $\vec{\alpha}_i$ and $\vec{\alpha}_j$. Since e is an edge of $R_1(\epsilon)$, there are no outwardly pointing arrows for up curves between the outwardly pointing arrows of $\vec{\alpha}_i$ and $\vec{\alpha}_j$.

By construction, the inward and outward pointing arrows alternate in $(U \cap c) \sqcup (D \cap c')$. Thus the number of outward pointing arrows that are between the outward pointing arrows at $\vec{\alpha}_i$ and $\vec{\alpha}_j$ is equal to one less than the number

of inward pointing arrows in this same interval I (see Figure 29). Each inward pointing arrow in this interval corresponds to an outward pointing arrow in $\iota(I)$.

In turn, each outward pointing arrow in $\iota(I)$ is associated to either an up curve or a down curve, and is therefore accounted for by either $\mathcal{M}_1(e)$ or $\mathcal{N}_1(e)$ by Lemma 8.11. This completes the proof of Lemma 8.12. \square

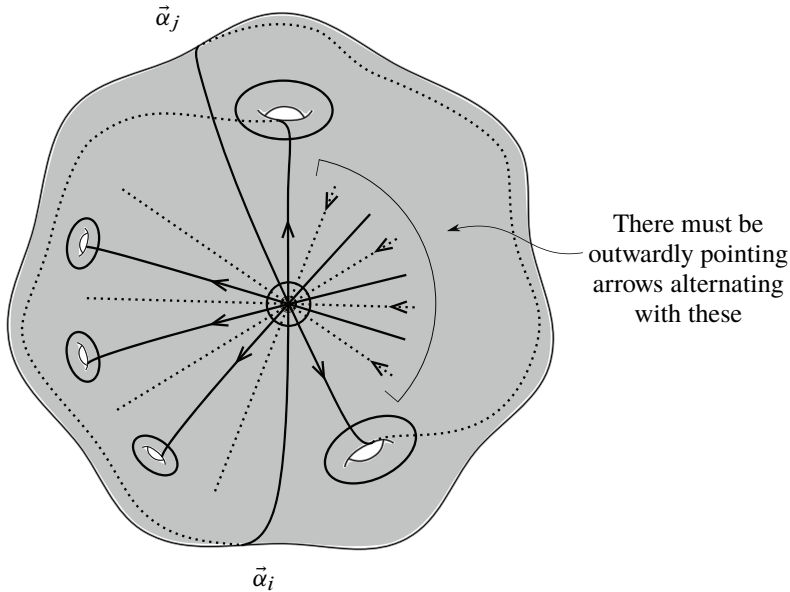


Figure 29. Given standard orientations of the curves in U and D , the arrows at the points of $(U \cap c) \sqcup (D \cap c')$ alternate between inward and outward pointing.

We are now ready to prove Proposition 6.2:

Proof of Proposition 6.2. Let $\phi \in \text{Mod}^*(S)$ have $\phi \cdot \Gamma(\epsilon) = \Gamma(\epsilon')$. By Proposition 7.1, we have either $\phi \cdot A(\epsilon) = A(\epsilon')$ or $\phi \cdot A(\epsilon) = B(\epsilon')$. After composing if necessary with the reflection of the surface that exchanges all up curves A with all down curves B (preserving the $(\mathbb{Z}/2\mathbb{Z} \oplus D_g)$ -orbit of ϵ'), we assume that $\phi \cdot A(\epsilon) = A(\epsilon')$. In this case, ϕ induces an isomorphism of the labeled polygons $P(\epsilon) \cong P(\epsilon')$. By Proposition 8.10, this implies that ϵ and ϵ' are in the same $(\mathbb{Z}/2\mathbb{Z} \oplus D_g)$ -orbit. \square

9. Stabilizing 1-systems

We return to stabilization. Suppose λ is an almost realization of a maximal complete 1-system on the surface S_g of genus g . Given any arc α on S_g , we can

delete tiny disks at its endpoints (in the complement of any of the curves of λ), and glue in an annulus. The result is a surface S_{g+1} of genus $g + 1$, and a complete 1-system naturally in correspondence with Γ . We refer to this system of curves by Γ as well, the distinction being clear from context. Note that Γ is not maximum in S_{g+1} , as it consists of two too few curves.

When α intersects each of the curves of λ once (on S_g), there are two readily available curves to add. Concatenate α with an arc that crosses the annulus to form a new simple closed curve α' . By construction, α' intersects each of the curves of Γ once. Moreover, the Dehn twist of α' around the core curve of the annulus, which we denote α'' , intersects each of the curves of $\Gamma \cup \{\alpha'\}$ once. Thus, given an arc α that intersects each of the curves of a realization of a maximum complete 1-system λ once, we produce the *stabilization* of λ along the arc α , denoted $\text{stab}(\lambda, \alpha) := \Gamma \cup \{\alpha', \alpha''\}$, a maximum complete 1-system on S_{g+1} . The following is immediate.

Lemma 9.1. *For each $\gamma \in \Gamma$, the trio $\{\alpha', \alpha'', \gamma\}$ in $\text{stab}(\lambda, \alpha)$ forms a triangle. If three curves don't form a triangle in Γ , then the corresponding curves don't form a triangle in $\text{stab}(\lambda, \alpha)$.*

There is a special case of this stabilization procedure. Choose a curve $\gamma \subset \lambda$, and fix an identification of γ with $S^1 = [0, 1]/0 \sim 1$, and the image of 0 with $p \in \gamma$. Identify an ϵ -neighborhood of γ with the annulus $S^1 \times (0, 1)$, with coordinates chosen, with ϵ small enough, such that the intersections of the other curves in $\lambda \setminus \{\gamma\}$ with the annulus each consist of a single vertical arc. In these coordinates, let α be the arc

$$\left\{ \left(t, \frac{1}{2} - t \right) : t \in [0, 1] \right\}.$$

We will refer to the stabilization $\text{stab}(\lambda, \alpha)$ by $\text{stab}(\lambda, p, \gamma)$, which we identify with $\Gamma \cup \{\gamma', \gamma''\}$. Note that the stabilization obtained may depend on the realization λ and on the point $p \in \gamma$ chosen, in the sense that it is possible that $\text{stab}(\lambda, p, \gamma)$ and $\text{stab}(\eta, q, \gamma)$ are $\text{Mod}^*(S_{g+1})$ -inequivalent, either for $p \neq q$ or for non-isotopic realizations $\lambda \not\sim \eta$.

In the case of this special version of stabilization, we note that there is a one-holed torus subsurface $\Sigma \subset S_{g+1}$ so that the curves in $\Gamma \setminus \{\gamma\}$ are disjoint from Σ , and so that the curves γ, γ' , and γ'' are homotopic as properly embedded arcs in the surface with boundary $S_{g+1} \setminus \Sigma$.

Lemma 9.2. *For each $\beta \in \Gamma \setminus \{\gamma\}$, the trios $\{\beta, \gamma, \gamma'\}, \{\beta, \gamma, \gamma''\} \subset \text{stab}(\lambda, p, \gamma)$ form triangles in S_{g+1} .*

Note that Lemmas 5.2 and 9.1 imply that any maximum complete 1-system obtained via stabilization has dual cube complex of dimension at least three, and Lemma 9.2, with Theorem 3 then implies that any obtained via the special case above has dimension at least four.

Remark 9.3. All of the maximum complete 1-systems in our construction can be obtained by a sequence of the more specialized stabilizing procedure, applied successively to the canonical example in genus two. In particular, they all have dual cube complexes of dimension at least three. Question 2 in §1 presents itself.

We fix choices of almost realizations for $\Gamma(\epsilon)$ as in §8, which we denote as well by $\Gamma(\epsilon)$, the distinction being clear from context.

Lemma 9.4. *For any choice of $p \in \delta$, any trio from $A(\epsilon) \subset \text{stab}(\Gamma(\epsilon), p, \delta)$ (resp. $B(\epsilon) \subset \text{stab}(\Gamma(\epsilon), p, \delta)$) forms a triangle in S_{g+1} .*

Proof. In the chosen realizations, the triangles formed among the curves of $A(\epsilon)$ are all disjoint from the stabilizing arcs (which track δ very closely), and from the disks chosen in the stabilization process. Thus the triangles persist. \square

Proposition 9.5. *If $\phi \in \text{Mod}^*(S_{g+1})$ satisfies*

$$\phi \cdot \text{stab}(\Gamma(\epsilon), p, \delta) = \text{stab}(\Gamma(\epsilon'), p, \delta),$$

then

- (1) $\phi \cdot \{\delta, \delta', \delta''\} = \{\delta, \delta', \delta''\}$, and
- (2) $\phi \cdot \{A(\epsilon), B(\epsilon)\} = \{A(\epsilon'), B(\epsilon')\}$.

Proof. By Lemma 9.4, the hyperplanes corresponding to curves of $A(\epsilon)$ (resp. $B(\epsilon)$) pass through maximal cubes of dimension $|A(\epsilon)|$ (resp. $|B(\epsilon)|$), which is even. The hyperplanes corresponding to δ, δ' , and δ'' pass through maximal cubes of odd dimension (e.g. the cube corresponding to the curves $\delta, \delta', \delta'', \alpha_{2i-1}, \alpha_{2i}, \beta_{2j-1}$, and β_{2j} , when $|A(\epsilon)|, |B(\epsilon)| \geq 1$). The conclusion follows as in the proof of Proposition 7.1. \square

We are now able to prove Proposition 6.4.

Proof of Proposition 6.4. Suppose $\phi \in \text{Mod}^*(S_{g+1})$ sends $\text{stab}(\Gamma(\epsilon), p, \delta)$ to $\text{stab}(\Gamma(\epsilon'), p, \delta)$, so that by Proposition 9.5 we have that ϕ preserves the set $\{\delta, \delta', \delta''\}$.

Consider the one-holed torus subsurface $\Sigma \subset S_{g+1}$ that arises in the course of the stabilizations $\text{stab}(\Gamma(\epsilon), p, \delta)$. Consider the mapping class $\check{\phi}$ induced on S_g by deleting Σ and identifying the resulting boundary component to a point. Since the curves of $\Gamma(\epsilon) \setminus \{\delta\}$ (and likewise $\Gamma(\epsilon') \setminus \{\delta\}$) can be made disjoint from Σ ,

the induced map $\tilde{\phi}$ takes the induced curves $\Gamma(\epsilon) \setminus \{\delta\}$ to $\Gamma(\epsilon') \setminus \{\delta\}$. Moreover, in the resulting surface δ, δ' , and δ'' go to the same homotopy class of curve, namely $\delta \subset S_g$. Since $\phi \cdot \{\delta, \delta', \delta''\} = \{\delta, \delta', \delta''\}$, we find that $\tilde{\phi} \cdot \delta = \delta$, so that $\Gamma(\epsilon)$ and $\Gamma(\epsilon')$ are equivalent under the action of $\text{Mod}^*(S_g)$. \square

References

- [1] T. Aougab, Large collections of curves pairwise intersecting exactly once. *Geom. Dedicata* **172** (2014), 293–302. [Zbl 1303.57018](#) [MR 3253783](#)
- [2] B. Bollobás, The asymptotic number of unlabelled regular graphs. *J. London Math. Soc. (2)* **26** (1982), no. 2, 201–206. [Zbl 0504.05051](#) [MR 0675164](#)
- [3] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren der Mathematischen Wissenschaften, 319. Springer Verlag, Berlin, 1999. [Zbl 0988.53001](#) [MR 1744486](#)
- [4] W. Cavendish and H. Parlier, Growth of the Weil–Petersson diameter of moduli space. *Duke Math. J.* **161** (2012), no. 1, 139–171. [Zbl 1244.32008](#) [MR 2872556](#)
- [5] I. Chatterji and G. Niblo, From wall spaces to CAT(0) cube complexes. *Internat. J. Algebra Comput.* **15** (2005), no. 5–6, 875–885. [Zbl 1107.20027](#) [MR 2197811](#)
- [6] M. de Graaf and A. Schrijver, Making curves minimally crossing by Reidemeister moves. *J. Combin. Theory Ser. B* **70** (1997), no. 1, 134–156. [Zbl 0888.57001](#) [MR 1441263](#)
- [7] B. Farb and D. Margalit, *A primer on mapping class groups*. Princeton Mathematical Series, 49. Princeton University Press, Princeton, N.J., 2012. [Zbl 1245.57002](#) [MR 2850125](#)
- [8] W. M. Goldman, Invariant functions on Lie groups and Hamiltonian flows of surface group representations. *Invent. Math.* **85** (1986), no. 2, 263–302. [Zbl 0619.58021](#) [MR 0846929](#)
- [9] G. Levitt and K. Vogtmann, A Whitehead algorithm for surface groups. *Topology* **39** (2000), no. 6, 1239–1251. [Zbl 0973.20024](#) [MR 1783856](#)
- [10] J. Malestein, I. Rivin, and L. Theran, Topological designs. *Geom. Dedicata* **168** (2014), 221–233. [Zbl 1284.57020](#) [MR 3158040](#)
- [11] P. Przytycki, Arcs intersecting at most once. *Geom. Funct. Anal.* **25** (2015), no. 2, 658–670. [Zbl 1319.57016](#) [MR 3334237](#)
- [12] M. Sageev, Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)* **71** (1995), no. 3, 585–617. [Zbl 0861.20041](#) [MR 1347406](#)
- [13] M. Sageev, CAT(0) cube complexes and groups. In M. Bestvina, M. Sageev, and K. Vogtmann (eds.), *Geometric group theory*. (Park City, UT, 2012) IAS/Park City Mathematics Series, 21. American Mathematical Society, Providence, R.I., and Institute for Advanced Study (IAS), Princeton, N.J., 2014, 7–54. [MR 3329724](#)

- [14] P. Schmutz Schaller, Mapping class groups of hyperbolic surfaces and automorphism groups of graphs. *Compositio Math.* **122** (2000), no. 3, 243–260. [Zbl 0981.57004](#)
[MR 1781329](#)

Received November 10, 2015

Tarik Aougab, Mathematics Department, Brown University, Box 1917, 151 Thayer Street,
Providence, RI 02912, USA

e-mail: tarik_aougab@brown.edu

Jonah Gaster, Department of Mathematics, Boston College, 541 Maloney Hall,
Chestnut Hill, MA 02467, USA

e-mail: gaster@bc.edu