Michigan Math. J. 73 (2023), 305–343
https://doi.org/10.1307/mmj/20205929

Automorphisms of the k-Curve Graph

Shuchi Agrawal, Tarik Aougab, Yassin Chandran,
Marissa Loving, J. Robert Oakley,
Roberta Shapiro, & Yang Xiao

Abstract. Given a natural number k and an orientable surface S of finite type, define the *k-curve graph* to be the graph with vertices corresponding to isotopy classes of essential simple closed curves on S and with edges corresponding to pairs of such curves admitting representatives that intersect at most k times. We prove that the automorphism group of the k-curve graph of a surface S is isomorphic to the extended mapping class group for all k satisfying $k \leq |\chi(S)| - 512$. We prove the same result for the so-called *systolic complex,* a variant of the curve graph with many complete subgraphs coming from interesting collections of systoles with respect to a hyperbolic metric. This resolves a conjecture of Schmutz Schaller.

1. Introduction

Let S be a connected orientable surface of genus g, possibly with finitely many punctures p, and let $\mathrm{Mod}^{\pm}(S)$ denote the extended mapping class group. The *curve complex* $\mathcal{C}(S)$ is a flag simplicial complex whose vertices correspond to isotopy classes of essential simple closed curves and whose edges represent pairs of such classes that can be realized disjointly on S. A celebrated theorem of Ivanov [18] identifies $\mathrm{Aut}(\mathcal{C}(S))$ with $\mathrm{Mod}^{\pm}(S)$ in all but finitely many cases. This result inspired a flurry of results in related contexts, where $\mathrm{Mod}(S)$ acts by simplicial automorphisms on some graph whose vertices represent homotopy classes of curves and/or arcs [6; 8; 15; 16; 19; 22; 29], or finite collections of curves and arcs [20; 7], or subsurfaces [6; 25].

In many of these papers the result is that the full automorphism group of the complex being considered is $\mathrm{Mod}^{\pm}(S)$, or at least virtually so in a finite number of sporadic cases, and the proofs all factor through Ivanov's original theorem by showing that any automorphism of a particular complex induces one of $\mathcal{C}(S)$. This led to Ivanov's metaconjecture, which we discuss further in Chapter 4: *Fifteen problems about the mapping class group* of [9, p. 77].

Received June 1, 2020. Revision received August 6, 2021.
Much of this paper is based on work supported by the National Science Foundation under Grant No. DMS-1439786 while the authors were in residence at the Institute for Computational and Experimental Research in Mathematics (ICERM) in Providence, RI, during the Summer@ICERM program. In addition, Aougab was partially supported by NSF grants DMS-1502623 and DMS-1807319, and Loving was supported by the NSF Graduate Research Fellowship under Grant No. DGE 1144245.

IVANOV'S METACONJECTURE. *Any "sufficiently rich" complex naturally associated with a surface should have* $\mathrm{Mod}^{\pm}(S)$ *as its group of automorphisms, and furthermore, there exists a proof of this, which factors through Ivanov's original theorem.*

The focus of this paper is to verify the metaconjecture for an infinite family of curve graphs whose edges represent bounded intersection. In particular, we will consider the following natural generalization of $\mathcal{C}(S)$. For any $k \in \mathbb{N}$, the *k-curve graph* is defined to be the graph whose vertices are those of $\mathcal{C}(S)$ and whose edges represent homotopy classes of curves with geometric intersection number at most k. Our main result characterizes $\mathrm{Aut}(\mathcal{C}_k(S))$ when $|\chi(S)|$ is sufficiently large relative to k.

THEOREM 1.1. *Suppose* $|\chi(S)| \geq k + 512$. *Then the natural map*

$$\mathrm{Mod}^{\pm}(S) \to \mathrm{Aut}(\mathcal{C}_k(S))$$

is an isomorphism.

When $k = 1$, Theorem 1.1 holds without the restriction on the Euler characteristic of S. We omit the proof for clarity since it is nearly identical to that of Theorem 1.3. Furthermore, the lower bound of $k + 512$ on $|\chi(S)|$ is not sharp. See Appendix for details on how this bound is derived.

Theorem 1.1 addresses Part (3) of Question 7.4 of Margalit's collection of open problems [7]. It represents a first step toward resolving Ivanov's metaconjecture in the cases where edges do not represent disjointness. In addition, to the authors' knowledge, it is only the third result in the literature that resolves Ivanov's conjecture for an infinite family of simplicial complexes. The first such result was the work of Brendle and Margalit [6] for *complexes of regions,* and the second is McLeay's extension [26] of their work from closed surfaces to punctured surfaces (including those of genus 0).

In addition to results concerning simplicial automorphisms mentioned above, there are a number of theorems characterizing simplicial injections [3; 5; 4; 14; 17], quasi-isometries [27], and other types of structure-preserving maps of $\mathcal{C}(S)$ and related complexes. For example, Raf and Schleimer [27] identify the group of quasi-isometries of $\mathcal{C}(S)$ with $\mathrm{Mod}^{\pm}(S)$. We remark that even though $\mathcal{C}_k(S)$ is quasi-isometric to $\mathcal{C}(S)$, this result does not imply Theorem 1.1. Indeed, a priori it is possible that an automorphism of $\mathcal{C}_k(S)$ moves every vertex a uniformly bounded distance and would therefore be equivalent to the identity as a quasi-isometry.

We also consider the following variant of the curve graph, which we denote $\mathcal{SC}(S)$. The vertices of this graph correspond to isotopy classes of essential curves, which are either nonseparating curves or separating curves that bound a twice punctured disk on one side. The edges represent pairs of such curves that intersect minimally, that is, at most once in the case that both vertices correspond to nonseparating curves, and at most twice when at least one of those vertices is a separating curve. The notation $\mathcal{SC}(S)$ is due to Schmutz Schaller [29] and stands for the *systolic complex,* as interesting sets of systoles on a hyperbolic surface correspond to complete subgraphs of $\mathcal{SC}(S)$. However, Anderson, Parlier, and Pettet [1] give examples of complete subgraphs of $\mathcal{SC}(S)$ that are not realizable as the set of systoles for any hyperbolic metric on S.

THEOREM 1.2. *If S is a closed surface with genus $g \geq 3$, then the natural map*

$$\mathrm{Mod}^{\pm}(S) \to \mathrm{Aut}(\mathcal{SC}(S))$$

is an isomorphism. If $g = 2$, then this map is surjective with kernel $\mathbb{Z}/2\mathbb{Z}$ corresponding to the hyperelliptic involution. If S is a surface of genus g with $p > 0$ punctures and $\chi(S) < 0$, then this map is an isomorphism for $(g, p) \neq (1, 2), (1, 3), (0, 5)$.

Theorem 1.2 represents an almost complete resolution to the conjecture of [29]; we remark that our techniques do not cover the cases $(g, p) = (1, 1), (1, 2), (1, 3), (0, 4), (0, 5)$. Following the outline of Ivanov's metaconjecture, our proof strategy relies on showing that any automorphism of $\mathcal{SC}(S)$ induces one of $\mathcal{C}(S)$. This fails when $(g, p) = (1, 2)$, since Luo proved in [22] that the curve complex of the twice punctured torus admits automorphisms that are not induced by homeomorphisms. In the cases $(g, p) = (1, 1)$ and $(0, 4)$ the systolic complex is isomorphic to the 1-skeleton of the Farey tessellation of the hyperbolic plane, whose automorphism group is $\mathrm{PGL}(2, \mathbb{Z})$, and so the theorem is known. Thus the only remaining cases are $(g, p) = (1, 2), (1, 3)$, and $(0, 5)$.

When $g \neq 0$, we can also consider the subgraph of $\mathcal{SC}(S)$ consisting only of nonseparating curves. Note that in the event that S is closed, this is the entirety of $\mathcal{SC}(S)$. We denote this graph by $\mathcal{N}_1(S)$ and give the following characterization of its automorphisms.

THEOREM 1.3. *Suppose that $g \geq 1$ and that $(g, p) \neq (1, 2)$. Then the natural map*

$$\mathrm{Mod}^{\pm}(S) \to \mathrm{Aut}(\mathcal{N}_1(S))$$

is an isomorphism for $(g, p) \neq (2, 0)$ and a surjection with kernel $\mathbb{Z}/2\mathbb{Z}$ otherwise.

1.1. Idea of Proofs

In both Theorems 1.2 and 1.1, we need to show that an automorphism of either $\mathcal{SC}(S)$ or of $\mathcal{C}_k(S)$ preserves edges that represent disjointness. In what follows, we let $\mathrm{link}(\cdot)$ denote the link of a vertex, the subgraph induced by the set of vertices adjacent to a given vertex. Given a pair of curves α, β connected by an edge, we study the subgraph $L(\alpha, \beta) = \mathrm{link}(\alpha) \cap \mathrm{link}(\beta)$; we refer to such a subgraph as the *link of an edge* or *edge link.*

In particular, we prove that the diameter of an edge link distinguishes between edges corresponding to disjoint curves and edges corresponding to curves intersecting once. For larger values of k, the diameters may not be sufficient to pick out the edges representing disjoint pairs, so we need a more careful analysis of the types of geodesics that exist in each edge link. We show that, under the additional hypothesis that the surface is sufficiently large, an edge link representing a pair of nondisjoint curves always has finite diameter. Furthermore, there must always exist a finite number of vertices, which we call *shortcut curves,* such that for any two vertices u, v in the link whose edge link distance is maximal, there is a geodesic from u to v that passes through a shortcut curve. This additional geometric property distinguishes edges representing disjoint curves from all other types of edges.

Throughout this paper, we employ both combinatorial and coarse-geometric techniques; for example, we use the technology of *subsurface projections* to compute exact diameters of the edge links. Given a pair of curves α and β intersecting k times, a standard surgery argument going back to Lickorish [21] yields a curve δ that intersects α at most once and β at most $k/2$ times. In the proof of Theorem 1.1, as opposed to such δ, we have need of a curve δ' that is disjoint from α and that intersects β at most $k/4$ times. For this, we use a variant of a proposition due to the second author, used to prove that curve graphs are uniformly hyperbolic [2].

1.2. Outline of Paper

Section 2 contains a brief introduction to curves on surfaces, several relevant graphs of curves associated with surfaces, the notion of subsurface projections, and some relevant coarse geometry. In Sections 3 and 4, we compute diameters of edge links in $\mathcal{N}_1(S)$ and use these to prove Theorem 1.3 at the end of Section 4. Section 5 provides a proof of Theorem 1.2, resolving the conjecture from [29]. Lastly, in Section 6, we prove our main result, Theorem 1.1. We also include an appendix, which contains a sketch of several known results needed in the proof of Theorem 1.1. Appendix A also provides an‘explanation of the restriction on the Euler characteristic of the surfaces required by Theorem 1.1.

2. Preliminaries

Throughout the paper, unless otherwise noted, S is an orientable surface of finite type, possibly with punctures and/or boundary components.

2.1. Curves and Arcs

A *simple closed curve* on S is a homotopy class of maps $S^1 \to S$ admitting a representative that is an embedding. We will often abuse notation and identify a simple closed curve with an embedded representative, and further identify this embedded representative with its image in S. A simple closed curve is *essential* if it is not homotopically trivial and if it is not homotopic to a map whose image bounds a once-punctured disk on one side.

Let $f, g : (0, 1) \to S$ be two embeddings such that $\lim_{t\to 0,1} f(t)$ are either punctures or points on boundary components, and similarly for g. Then f and g are *boundary-slide homotopic* if there is a homotopy $H : [0, 1] \times (0, 1) \to S$ from

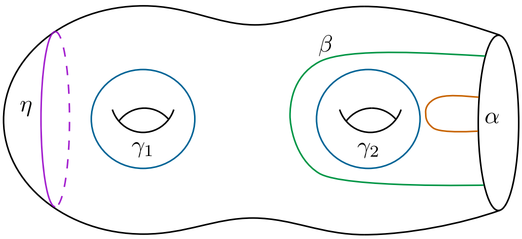

Figure 1 The boundary slide homotopy class of β is an essential arc, whereas α is not since it cobounds a disk with an arc of ∂S. The homotopy class of each γ_i is essential, and that of η is not since η bounds a disk on the surface. The collection $\{\gamma_1, \gamma_2\}$ is a multicurve.

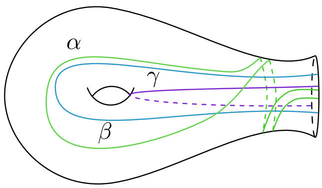

Figure 2 The green arc α and the blue arc β are boundary-slide homotopic. Neither is boundary-slide homotopic to γ. The boundary-slide homotopy class of α and β, denoted $[\alpha]$, is an essential arc, as is the class of $[\gamma]$. Since $\gamma \cap \beta = \emptyset$ and $\beta \in [\alpha]$, $\{[\alpha], [\gamma]\}$ is a multiarc.

f to g such that $\lim_{s\to 0} H(t,s)$ is either a puncture or on a fixed boundary component for all t, and similarly as $s \to 1$ (see Figures 1 and 2). Then an *essential arc* will be a nontrivial boundary-slide homotopy class of such maps.

Given a pair of essential simple closed curves or arcs α, β, their *geometric intersection number* $i(\alpha, \beta)$ is the minimum of $|\alpha' \cap \beta'|$ taken over all representative images $\alpha' \subset S$ of α and $\beta' \subset S$ of β. Note that here α can be a curve, and β can be an arc. If α', β' realize the geometric intersection number for their respective homotopy classes, then they are said to be in *minimal position.*

A *multicurve* (resp., *multiarc*) is a collection of pairwise distinct essential simple closed curves (resp., arcs) whose pairwise geometric intersection numbers are all 0. A collection of pairwise disjoint curves and arcs will, by convention, be referred to as a multicurve. As is well known, for $(g, p) \neq (1, 0)$, any multicurve on S consisting of curves contains at most $3g + p - 3$ connected components, and this bound is realizable.

Lastly, we introduce the notion of a *weighted multiarc,* which will be used in the proof of Proposition 6.7 and in Appendix A. A *weighted multiarc* is a multiarc with positive integer weights assigned to each arc. We use $|\alpha|$ to denote the number of arcs in a multiarc $\alpha = \{a_1, a_2, \ldots, a_n\}$, and $w(\alpha) = \sum_{i=1}^{n} w_i$ to denote the total weight, where w_i is the weight assigned to arc $a_i \in \alpha$.

2.2. Relevant Graphs and Their Automorphisms

In this section, we introduce various graphs whose vertices will represent curves or arcs in S and with edges corresponding to various constraints on the geometric intersection number. We call an edge connecting vertices v and w an *n-edge* if the curves corresponding to v and w minimally intersect n times. An edge is a *nonzero* edge if $n \neq 0$.

Let $\mathcal{AC}(S)$, the *arc and curve graph,* be the graph whose vertices correspond to essential simple closed curves and arcs on S, and whose edges correspond to disjoint pairs, so all edges are 0-edges. When S is an annulus, $\mathcal{AC}(S)$ consists only of arcs connecting the two boundary components. We note that $\mathcal{AC}(S)$ is connected when $3g+n+b \geq 5$: Theorem 4.3 in [10] states that $\mathcal{C}(S)$ is connected, and any arc is disjoint from at least one curve.

Define $\mathcal{N}(S)$, the *nonseparating curve graph,* to be the graph whose vertices are nonseparating simple closed curves with edges between classes that admit disjoint representatives. All edges in $\mathcal{N}(S)$ are 0-edges. It is known that $\mathcal{N}(S)$ is connected for genus $g \geq 2$ (see Theorem 4.4 in [10]). In addition, $\mathcal{N}(S)$ is of infinite diameter. The argument of Masur and Minsky [23] for the infinite diameter of $\mathcal{C}(S)$ apply fairly directly to $\mathcal{N}(S)$. Irmak [15] showed that for surfaces with $g > 1$, $\mathrm{Aut}(\mathcal{N}(S)) \cong \mathrm{Mod}^{\pm}(S)$.

The *systolic complex,* denoted $\mathcal{SC}(S)$, is defined differently for closed and punctured surfaces. When S is closed, $\mathcal{SC}(S)$ has vertices corresponding to isotopy classes of nonseparating curves and whose edges represent pairs of such curves with geometric intersection number at most 1. In this case, $\mathcal{SC}(S)$ contains both 0-edges and 1-edges. If S is not closed, then $\mathcal{SC}(S)$ has an additional vertex for each separating curve that bounds a twice punctured disk on one side; such vertices are connected to others by an edge when there are at most two intersections. We observe that for closed surfaces, $\mathcal{SC}(S)$ is equal to $\mathcal{N}_1(S)$.

We define $\mathcal{N}_1(S)$, the *nonseparating 1-curve graph,* as the subgraph of $\mathcal{SC}(S)$ consisting only of nonseparating curves. To the best of the authors' knowledge, this graph has not been otherwise named or studied extensively in the literature. Note that $\mathcal{N}(S)$ is a subgraph of $\mathcal{N}_1(S)$ with the same vertex set, and therefore $\mathcal{N}_1(S)$ is connected when $g \geq 2$ because $\mathcal{N}(S)$ is connected. When $g \geq 2$, $\mathcal{N}_1(S)$ is quasi-isometric to $\mathcal{N}(S)$, and therefore it is of infinite diameter. See Remark 2.4.

Along with $\mathcal{SC}(S)$, Schmutz Schaller defined the graph $\mathcal{G}(S)$. When $g \geq 1$, $\mathcal{G}(S)$ has as its vertex set the collection of all nonseparating curves and has edges corresponding to pairs of curves intersecting exactly once. This means that for closed surfaces of positive genus, $\mathcal{G}(S)$ is a subgraph of $\mathcal{SC}(S)$. When $g = 0$, $\mathcal{G}(S)$ has as its vertex set the collection of all curves bounding a twice punctured disk on one side and whose edges correspond to pairs of curves intersecting exactly twice. Again, this is a subgraph of $\mathcal{SC}(S)$. See Figure 3 for an illustration of how thse graphs differ. It is a result of Schaller [29] that $\mathrm{Aut}(\mathcal{G}(S)) \cong \mathrm{Mod}^{\pm}(S)$. In the same paper, Schaller conjectured the following, which we resolve in Theorem 1.2 for all but $(g, p) = (1,3), (0,5)$.

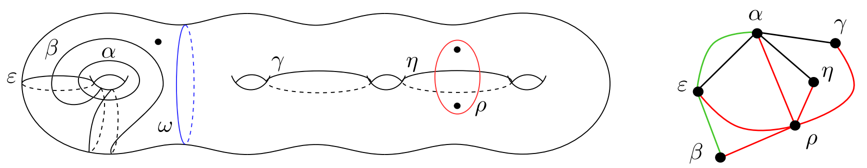

Figure 3 The black curves α, β, ε, η, and γ are all vertices of $\mathcal{N}_1(S)$. The black edges in the graph record adjacencies in $\mathcal{N}_1(S)$. The curve ρ is not in $\mathcal{N}_1(S)$, but it is in $\mathcal{SC}(S)$. The red edges represent adjacencies in $\mathcal{SC}(S)$. Note that the black edges are also in $\mathcal{SC}(S)$. The green edge represents adjacency in $\mathcal{G}(S)$. The curve ω is not a vertex in $\mathcal{N}_1(S)$, $\mathcal{SC}(S)$, or $\mathcal{G}(S)$. Note that not all the edges in this subgraph are shown to avoid cluttering the diagram.

Conjecture 2.1 (Schmutz Schaller, [29]). *The automorphism group of $\mathcal{SC}(S)$ is isomorphic to* $\mathrm{Mod}^{\pm}(S)$.

For any one of the above-mentioned graphs and for the curve complex $\mathcal{C}(S)$, we obtain a metric on the vertex set by identifying each edge with the unit interval and defining the distance between two vertices to be the minimum number of edges contained in any edge path between them. Given one of these graphs G, the distance function will be denoted by $d_G(\cdot, \cdot)$. All graphs mentioned above are of infinite diameter; for all but finitely many surfaces, this follows directly from the work of Masur and Minsky [23] on the infinite diameter of $\mathcal{C}(S)$.

The following lemma establishes a quasi-isometry between $\mathcal{AC}_1(S)$ and $\mathcal{AC}(S)$, where $\mathcal{AC}_1(S)$ has the same vertex set as the standard arc and curve graph but with edges when there is at most one intersection.

Proposition 2.2. *If S is a surface with punctures or with nonempty boundary and so that $(g, p) \neq (0, 3)$, then*

$$\mathcal{AC}_1(S) \cong_{\mathrm{QI}} \mathcal{AC}(S),$$

where $\mathcal{AC}_1(S)$ is the arc and curve graph of S with both 0-edges and 1-edges. Moreover, both are of infinite diameter.

Proof. Let $\phi : \mathcal{AC}_1(S) \longrightarrow \mathcal{AC}(S)$ be the identity map on the vertices.

For any vertices u, v in $\mathcal{AC}_1(S)$), suppose $d_{\mathcal{AC}_1}(u, v) = l$ for some $l > 0$. If $l = 1$ and $i(u, v) = 1$, then either the genus g is nonzero, or at least one of u, v is an arc. If $g > 1$ or if the number of punctures $p > 1$, then there is an essential curve in $S \setminus (u \cup v)$, and thus $d_{\mathcal{AC}}(\phi(u), \phi(v)) = 2$.

If $(g, p) = (1, 1)$ and u, v are curves, then without loss of generality one is the $1/0$ and the other is the $0/1$ curve, and then there are disjoint arcs λ_u, λ_v such that $i(\lambda_u, u) = i(\lambda_v, v) = 0$, and so $d_{\mathcal{AC}}(\phi(u), \phi(v)) \leq 3$.

If one of u, v is a curve and the other is an arc, then there is an arc disjoint from both. These are shown in Figure 4.

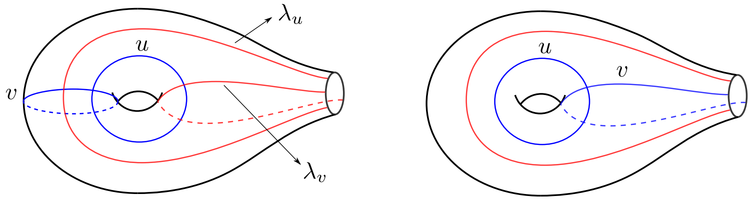

Figure 4 Two possible configurations when $(g, p) = (1, 1)$: on the left when u and v are curves, and on the right when u is a curve and v is an arc.

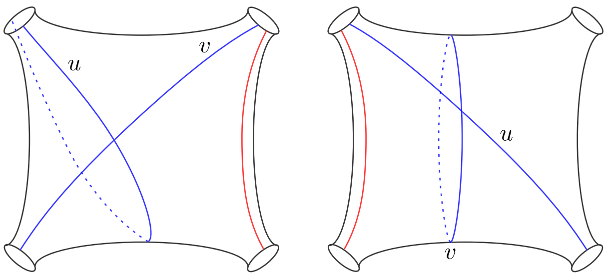

Figure 5 Two possible configurations when $(g, p) = (0, 4)$. On the left when u and v are both arcs and u is separating. On the right when u is an arc and v is a curve.

If both are arcs, then Hatcher's original surgery argument for the contractibility of the arc complex (see the Main Theorem of [12] or Theorem 5.5 of [10]) implies that the distance from u to v in the arc complex of S is at most 2.

Finally, if $g = 0$, then we will argue in such a way that the proof for $p = n$ implies a proof for $p > n$, and thus we can assume that $p = 4$.

If both u and v are arcs, then at most one can be separating. In this case, it is easy to find an arc in the complement of u and v: without loss of generality, u separates two punctures from another, and v only witnesses one of the two punctures on one side of u. Thus there is an arc λ connecting the two punctures on one side of u, disjoint from v. If neither u nor v separate, then cutting along one produces a 3-holed sphere, in which the other arc becomes two arcs. Given any two disjoint arcs in a 3-holed sphere, there is always a third essential arc disjoint from both, as shown in Figure 5.

Lastly, if u is an arc and v is a curve, then u cannot be separating, and so up to homeomorphism there is a unique configuration: v separates two punctures from the other two, and u connects one puncture on one side of v to one on the other. It is then straightforward to find an arc disjoint from both.

If $l > 1$ and if the shortest path between u and v contains no 1-edges, then $d_{\mathcal{AC}}(\phi(u), \phi(v)) = l$ as well. If the shortest path contains a 1-edge (ρ, η), then as above, we choose a path of length at most 3 through 0-edges from ρ to η. Therefore $d_{\mathcal{AC}}(\phi(u), \phi(v)) \leq 3l$.

When $(g, p) \neq (1, 1)$, that $\mathcal{AC}(S)$ has infinite diameter follows from the fact that $\mathcal{C}(S)$ has infinite diameter [23] and from Theorem 1.3 of [19], which states that $\mathcal{AC}(S) \cong_{\mathrm{QI}} \mathcal{C}(S)$. When $(g, p) = (1, 1)$, the graph $\mathcal{AC}_1(S)$ is quasi-isometric to the Farey graph, which is of infinite diameter. $\square$

REMARK 2.3. Masur and Minsky [23] not only show that $\mathcal{C}(S)$ is infinite diameter, but that the orbit of any pseudo-Anosov mapping class is as well. Since the quasi-isometries between $\mathcal{AC}(S)$, $\mathcal{C}(S)$, and $\mathcal{AC}_1(S)$ are all $\mathrm{Mod}(S)$-equivariant, the same is true for $\mathcal{AC}(S)$ and $\mathcal{AC}_1(S)$.

REMARK 2.4. When $g \geq 2$, a simplified version of the above argument demonstrates that $\mathcal{N}(S)$ and $\mathcal{N}_1(S)$ are quasi-isometric, via the identity map on the vertices.

We conclude this subsection with an argument due originally to Hempel [13] (see also [28]), which implies the connectedness of the curve complex. We reference the argument explicitly since it will be called upon in Section 6. Hempel's argument proves that given simple closed curves α, β on a surface S representing vertices of $\mathcal{C}(S)$,

$$d_{\mathcal{C}(S)}(\alpha, \beta) \leq 2\log_2(i(\alpha, \beta)) + 2.$$

Hempel considers a simple surgery that replaces α with a curve α' intersecting β at most half the number of times that α intersects β. The curve α' is obtained by following along α until the arrival at a chosen intersection with β, and then following along β until the arrival at a subsequent intersection, at which point α' follows along α again until closing up.

2.3. *Subsurface Projections*

An *essential subsurface* of S is a pair (Y, i_Y) where Y is a surface (potentially with boundary), and $i_Y : Y \hookrightarrow S$ is a π_1-injective map and an embedding on the interior of Y, so that each component of ∂Y maps to either an essential simple closed curve on S or a component of ∂S. We will often identify an essential subsurface with its image in S. When Y is an annulus, we say that it is an *annular* subsurface, and otherwise that Y is *nonannular*. An essential simple closed curve or arc is said to be in *minimal position* with an essential subsurface Y when it is in minimal position with all components of ∂Y (see Figure 6).

Given a nonannular essential subsurface $Y \subset S$, the *subsurface projection*

$$\pi_Y : \mathcal{C}(S) \longrightarrow \mathcal{P}(\mathcal{AC}(Y))$$

takes a vertex $\alpha \in \mathcal{C}(S)$ to the multicurve in Y obtained by taking all distinct homotopy classes occurring in the intersection of α with Y, after Y and α are put in minimal position (see [24] for more detail).

When Y is an annulus, we first consider the cover S_Y of S corresponding to $\pi_1 Y$. This cover compactifies to an annulus, and we let $\pi_Y(\alpha)$ be any lift of α to this annulus that connects its two boundary components (see Figure 7).

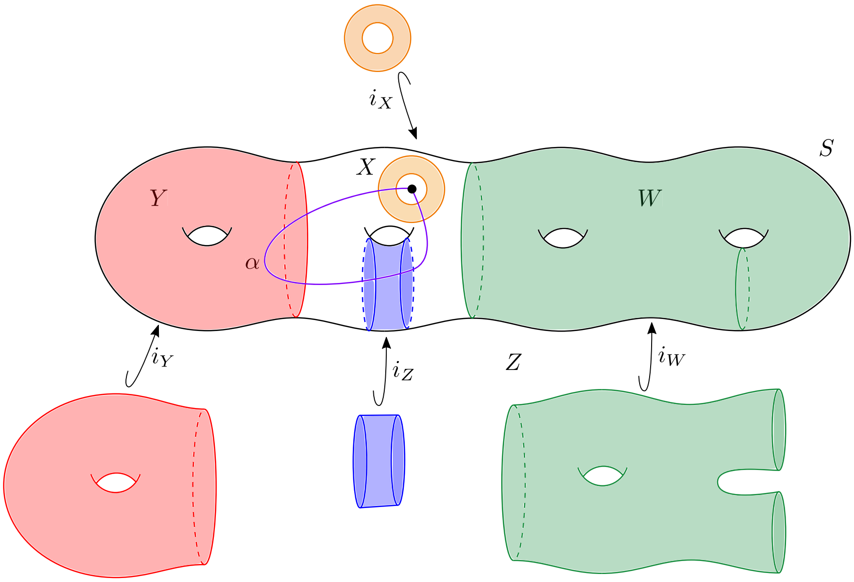

Figure 6 The pairs (Y, i_Y) and (W, i_W) are examples of essential nonannular subsurfaces, and (Z, i_Z) is an essential annular subsurface. The arc α is in minimal position with Z and W but not with Y.

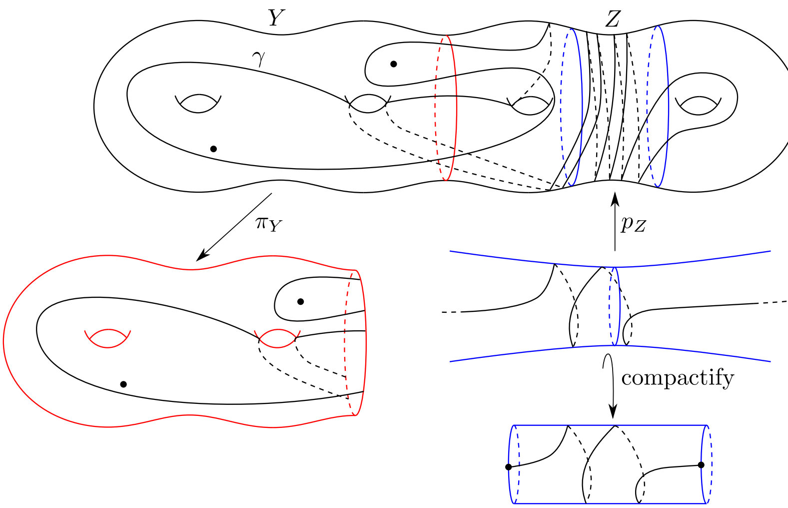

Figure 7 The subsurface projection of γ to Y consists of three pairwise disjoint essential arcs. The projection of γ to Z is obtained first by lifting γ, via the covering map p_Z, to the annular cover associated with Z, and then compactifying to obtain an essential arc.

We define the *Y-subsurface distance* as $d_Y(\alpha, \beta) := \mathrm{diam}_{\mathcal{AC}(Y)}(\pi_Y(\alpha) \cup \pi_Y(\beta))$.

3. Curves That Intersect Exactly Once

Let G be a graph, and let v be a vertex of G. From now on, we will use $V(G)$ to refer to the vertex set of G. We define the *link of a vertex* v, $\mathrm{link}(v)$, to be the induced subgraph of G containing the set of all vertices adjacent to v. Note that $v \notin \mathrm{link}(v)$.

Let u, v be adjacent vertices in a graph. We will denote the edge between u and v by (u, v). Define the *link of an edge* (u, v) to be

$$L(u, v) := \mathrm{link}(u) \cap \mathrm{link}(v).$$

In this section, we focus on the graph $\mathcal{N}_1(S)$, which, when S is closed, agrees with $\mathcal{SC}(S)$. There are two types of edges in $\mathcal{N}_1(S)$: we call edges that connect vertices admitting disjoint representatives *0-edges* and those that minimally intersect once *1-edges*. To prove that an automorphism of $\mathcal{N}_1(S)$ induces an automorphism of $\mathcal{C}(S)$, we will give a graph-theoretic criterion to distinguish between 0-edges and 1-edges.

The arguments in this section will be used to prove Theorem 1.2 and will also be useful for the proof of Theorem 1.3. In particular, the first step of the proof of Theorem 1.3 is showing the following:

Proposition 3.1. *The diameter of a link of an edge in $\mathcal{N}_1(S)$ equals* 4 *if and only if e is a* 1*-edge.*

This section is devoted to proving one direction of this statement, which we record as Lemma 3.2. We will prove the other direction in Section 4.

Lemma 3.2. *The diameter of the link of a* 1*-edge in $\mathcal{N}_1(S)$ is* 4.

The proof of Lemma 3.2 will break into two pieces, Lemmas 3.3 and 3.4, which we will prove in Subsections 3.1 and 3.2, respectively. We begin by establishing some helpful language and notation.

Let $u, v \in V(\mathcal{N}_1(S))$ be such that (u, v) is a 1-edge. Let $S_1 = N(u \cup v)$ be the regular neighborhood of $u \cup v \subset S$, which is homeomorphic to a torus with one boundary component.

Recall that the set of isotopy classes of essential simple closed curves on a torus with one boundary component is in bijection with the set $\mathbb{Q} \cup \{\frac{1}{0}\}$, which we will call *slopes*. In particular, the meridian curve is associated with $0/1$, and the longitude is associated with $1/0$. Moreover, the boundary-slide isotopy classes of essential simple arcs on a torus with one boundary component are in one-to-one correspondence with the isotopy classes of essential simple closed curves on the torus. We thus may refer to curves or arcs on S_1 by their associated slopes, which are of the form p/q, where p, q are coprime integers. Note that an essential

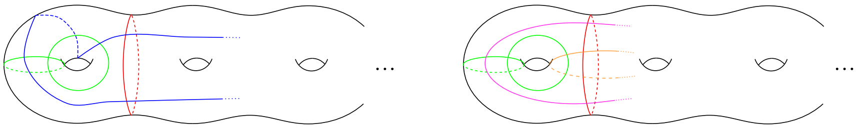

Figure 8 On the left: a 1-curve α in S_1. On the right: a 2-curve α in S_1.

simple closed curve or simple arc with slope p/q intersects the meridian $|p|$ times and the longitude $|q|$ times in minimal position.

Up to a change of coordinates if necessary, we may assume that u and v are the $0/1$ and $1/0$ curves, respectively, as shown by the green curves in Figure 8. We will also denote the $1/1$ and $-1/1$ curves in S_1 by γ^+ and γ^-. Observe that if $\alpha \subset S$ is a simple closed curve with $i(\alpha, u), i(\alpha, v) \leq 1$, then $\alpha \cap S_1$ has at most two nontrivial components: each component of $\alpha \cap S_1$ is an essential arc in S_1, and any essential arc in S_1 must intersect the $0/1$ curve or the $1/0$ curve at least once.

By a *1-curve* we will mean any curve α such that $i(\alpha, u) = i(\alpha, v) = 1$ and such that $\alpha \cap S_1$ has a single component. If α is a 1-curve, then $i(\alpha, u) = i(\alpha, v) = 1$, and $\alpha \cap S_1$ must be a $1/1$ arc or $-1/1$ arc. One of these two possible configurations is shown on the left of Figure 8. Similarly, a *2-curve* is any curve intersecting each of u and v exactly once, and so that its intersection with S_1 has two components. An illustration of a 2-curve is shown on the right of Figure 8.

3.1. *The Link of a* 1*-Edge Has Diameter at Most* 4

We are now prepared to show that the link of a 1-edge has diameter at most 4.

LEMMA 3.3. *If* (u, v) *is a* 1*-edge in* $\mathcal{N}_1(S)$, *then* $\operatorname{diam}(L(u, v)) \leq 4$.

Proof. By assumption, $i(u, v) = 1$. Let $S_2 = \overline{S \setminus S_1}$. Up to a change of coordinates, we assume that u and v are the $0/1$ and $1/0$ curves in S_1, respectively.

Let $\alpha \in L(u, v)$. Up to a homeomorphism of S_1 exchanging u and v, the possible configurations of α relative to u and v are as follows:

(1) $i(u, \alpha) = i(v, \alpha) = 1$,
(2) $i(u, \alpha) = i(v, \alpha) = 0$,
(3) $i(u, \alpha) = 0$ and $i(v, \alpha) = 1$.

We will refer to a curve $\alpha \in L(u, v)$ as being of either type (1), (2), or (3) depending on which of the above three holds. Recall that γ^+ and γ^- are the $1/1$ and $-1/1$ curves on S_1, respectively.

CASE 1. If $\alpha \in L(u, v)$ is of type (1), then $d_L(\alpha, \{\gamma^+, \gamma^-\}) \leq 2$.

Proof. If $\alpha \subset S_1$, then it must be either γ^+ or γ^-, and we are done.

If $\alpha \not\subset S_1$, then α intersects S_1 in either one or two connected components, that is, α is either a 1- or 2-curve. If $\alpha \not\subset S_1$ is a 1-curve, then $\alpha \cap S_1$ must be either a $1/1$ or $-1/1$ arc, and therefore we have either $d_L(\alpha, \gamma^+) = 1$ or $d_L(\alpha, \gamma^-) = 1$.

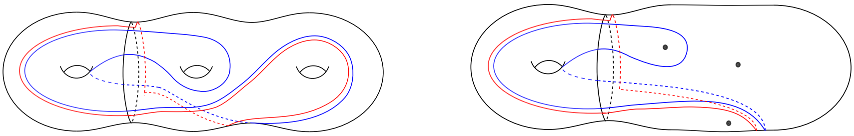

Figure 9 The constructed curve β is shown in red.

Otherwise, $\alpha \not\subset S_1$ is a 2-curve, and hence α has two connected components in S_1 and they can be arranged as in Figure 8. Note that this implies $\alpha \cap S_2$ also has two disjoint connected components.

Observe that the endpoints of the two components of $\alpha \cap S_2$ are unlinked on ∂S_2, that is, if we read off the cyclic order of these endpoints on ∂S_2 clockwise or counterclockwise, then the two endpoints of each component are adjacent to each other (see the two blue arcs in S_2 in Figure 9).

We can now construct a curve β as follows: starting at a point of intersection $\alpha \cap \partial S_1$, follow along one component of $\alpha \cap S_1$. Upon arriving back at ∂S_1, continue following along α through one component of $\alpha \cap S_2$. When arriving back at ∂S_2, the combinatorics of the points $\alpha \cap \partial S_2$ described in the previous paragraph implies that there is a choice of arc along ∂S_2 that ends where we began and so that the resulting closed curve is simple and intersects α exactly once. In particular, β is nonseparating; it also intersects each of γ^+ and γ^- exactly once, and thus $d_L(\alpha, \gamma^+), d_L(\alpha, \gamma^-) \leq 2$ (see Figure 9). □

Case 2. If $\alpha \in L(u, v)$ is of type (2), then $d(\alpha, \gamma^+) = d(\alpha, \gamma^-) = 1$.

Proof. In this case, $\alpha \subset S_2$. Thus $d(\alpha, \gamma^+) = d(\alpha, \gamma^-) = 1$, and we are done. □

Case 3. If $\alpha \in L(u, v)$ is of type (3), then $d_L(\alpha, \gamma^+) = d_L(\alpha, \gamma^-) = 1$.

Proof. In this case, α intersects both γ^+ and γ^- exactly once. Hence we have that $d_L(\alpha, \gamma^+) = d_L(\alpha, \gamma^-) = 1$, as desired. □

Now, let $\alpha, \beta \in L(u, v)$. Suppose both α and β are two type (1) curves. If both α and β are 2-curves, then by Case 1 they are distance at most 2 from both γ^+ and γ^-, and thus $d_L(\alpha, \beta) \leq 4$. If α is a 1- curve and β is a 2-curve, then again we have that $d_L(\alpha, \beta) \leq 4$ because $d_L(\beta, \gamma^+), d_L(\beta, \gamma^-) \leq 2$, and thus $d_L(\alpha, \{\gamma^+, \gamma^-\}) = 1$. If both α and β are 1-curves, then both are at a distance at most 1 from either γ^+ or γ^- and therefore are at a distance of at most 4 from one another, since $d_L(\gamma^+, \gamma^-) = 2$.

If α is a 1-curve and $\beta \subset S_1$, then $d_L(\alpha, \beta) \leq 3$, since $\beta \in \{\gamma^+, \gamma^-\}$ and either $d_L(\alpha, \gamma^+) = 1$ or $d_L(\alpha, \gamma^-) = 1$. Finally, if α is a 2-curve and $\beta \subset S_1$, then $d_L(\alpha, \beta) \leq 2$, since $\omega \in \{\gamma^+, \gamma^-\}$ and $d_L(\alpha, \gamma^+), d_L(\alpha, \gamma^-) \leq 2$.

It remains to consider curves of types (2) and (3). A type (2) curve intersects neither u nor v and therefore is disjoint from S_1. Thus it is at distance 1 from both γ^+ and γ^-, and therefore at distance at most 3 from any type (1) curve and at a distance of at most 2 from any other type (2) curve. A type (3) curve must also be at distance 1 from both γ^+ and γ^- and so is also distance at most 3 from any type (1) curve and at distance at most 2 from any type (2) or (3) curve. □

3.2. The Link of a 1-Edge Has Diameter at Least 4

We end this section by showing that the diameter of the link of a 1-edge is at least 4.

LEMMA 3.4. *Suppose* $(g, p) \neq (1, 2)$. *If* (u, v) *is a* 1*-edge in* $\mathcal{N}_1(S)$, *then* $\operatorname{diam}(L(u, v)) \geq 4$.

Proof. Let u, v, S_1, and S_2 be as in Lemma 3.3. Again, let γ^+ be the 1/1 curve, and let γ^- be the $-1/1$ curve in S_1. It suffices to show that there exist two 2-curves α and β whose shortest connecting path passes through γ^+ or γ^-.

Since S_2 is not a 3-holed sphere, the diameter of $\mathcal{AC}(S_2)$ is infinite. Hence there exist arcs $\delta, \eta \in \mathcal{AC}(S_2)$ such that $d_{\mathcal{AC}(S_2)}(\delta, \eta) \geq 23$. There also exists arcs $\delta', \eta' \in \mathcal{AC}(S_2)$ such that $d_{\mathcal{AC}(S_2)}(\delta, \delta') = 1 = d_{\mathcal{AC}(S_2)}(\eta, \eta')$, and so that the endpoints of δ and δ' (respectively, η and η') do not link along ∂S_2. Note that $d_{\mathcal{AC}(S_2)}(\delta', \eta') \geq 21$.

Now construct a 2-curve α from δ and δ' as follows: choose an endpoint from each of δ and δ' that are not consecutive on ∂S_2. Connect these through S_1 via the $1/0$ arc and connect the remaining two endpoints with the $0/1$ arc. Construct β in a similar manner from η and η'. Note that $\pi_{S_2}(\alpha) = \{\delta, \delta'\}$ and $\pi_{S_2}(\beta) = \{\eta, \eta'\}$.

Let $\rho = \{\alpha, v_0, \ldots, v_n, \beta\}$ be any shortest path in $L(u, v)$ connecting α to β. If each v_i has a nontrivial projection to S_2, by choosing one vertex from each of $\pi_{S_2}(v_i)$ we obtain a path of length $m \leq n + 2$ in $\mathcal{AC}_1(S_2)$ from δ' to η'. Note that m may be strictly less than $n + 2$ if there exist v_i and v_j with the same projection to S_2. This, in turn, yields a path of length at most $3m$ in $\mathcal{AC}(S_2)$ between δ' and η'. The factor of 3 comes from the quasi-isometry established in Proposition 2.2. Since $d_{\mathcal{AC}(S_2)}(\delta', \eta') \geq 21$, $3(n + 2) \geq 3m \geq 21$, and so the length of ρ is greater than or equal to 5, a contradiction to Lemma 3.3.

This implies there exists some v_i that projects trivially to S_2, and therefore $v_i \in \{\gamma^+, \gamma^-\}$. Therefore we have that the length of ρ is at least 4. □

4. Disjoint Curves

In this section, we will complete the proof of Proposition 3.1 and use it to prove Theorem 1.3 by proving the following lemma.

LEMMA 4.1. *If* (u, v) *is a* 0*-edge in* $\mathcal{N}_1(S)$, *then* $\operatorname{diam} L(u, v) \neq 4$.

This lemma is proved by considering two cases: when the two vertices u, v of the edge (u, v) are jointly separating, as considered in Lemma 4.2, and when they are jointly nonseparating, as considered in Lemma 4.5.

4.1. The Jointly Separating Case

In this section, we will prove Lemma 4.1 for a 0-edge (u, v) when u and v are jointly separating.

LEMMA 4.2. *Suppose $u, v \in \mathcal{N}_1(S)$ are disjoint curves that are jointly separating. If both components of $S \setminus u \cup v$ contain nonseparating curves of S, then $L(u, v)$ has diameter at most 3. Otherwise, $L(u, v)$ has infinite diameter.*

Proof. Denote the two connected components of $S \setminus (u \cup v)$ by S_1 and S_2. Then any curve in $L(u, v)$ is contained in either S_1 or S_2 or has nontrivial intersection with both subsurfaces. Let $\alpha, \beta \in L(u, v)$. Assume first that both S_1 and S_2 contain nonseparating curves of S.

If α and β are contained in the same component of $S \setminus (u \cup v)$, say S_1, then $d_L(\alpha, \beta) \leq 2$ since by assumption there is a nonseparating curve contained entirely in S_2. On the other hand, if α and β are contained in different components of $S \setminus (u \cup v)$, then they are disjoint, and thus $d_L(\alpha, \beta) = 1$.

Without loss of generality, suppose α is contained in S_1 and β has nontrivial intersection with both S_1 and S_2. Then $\beta \cap S_2$ is a nonseparating arc on S_2 with endpoints on distinct boundary components. Letting ω denote a nonseparating curve of S contained in S_2 (which exists by assumption), by the classification of surfaces there is a homeomorphism f of S_2 fixing u and v that sends $\beta \cap S_2$ to an arc crossing ω at most once. If ω does not separate u from v, then we can choose f so that $f(\beta \cap S_2)$ is disjoint from ω. Otherwise, we can choose f so that $f(\beta \cap S_2)$ crosses ω exactly once. Then $f^{-1}(\omega)$ is a nonseparating curve on S that is at an $L(u, v)$-distance of 1 from both α and β.

Finally, if both α and β have nontrivial intersections with both components of $S \setminus (u \cap v)$, then each of them intersects S_1 and S_2 in a nonseparating arc. Then by the same argument used in the previous paragraph there are nonseparating curves $\omega_1 \subset S_1$, $\omega_2 \subset S_2$ such that $d_L(\alpha, \omega_1) = d_L(\beta, \omega_2) = 1$. Thus $d_L(\alpha, \beta) \leq 3$, since $i(\omega_1, \omega_2) = 0$.

It remains to consider the case where either S_1 or S_2 contains no nonseparating curve of S. Since S is not a twice punctured torus, at least one component of $S \setminus (u \cup v)$, say S_1, is not a 3-holed sphere. Now for $n \in \mathbb{N}$, $n > 2$, by Proposition 2.2 choose arcs ω_1, ω_n with the same endpoints, each with one endpoint on u and the other on v so that

$$d_{\mathcal{AC}_1(S_1)}(\omega_1, \omega_n) \geq n.$$

Then choose an arc $\lambda \subset S_2$ with the same endpoints as ω_1, ω_n so that the concatenation of λ with ω_1 and with ω_n yields nonseparating curves η_1, η_n (these will be nonseparating because they both cross each of u and v exactly once). As S_2 contains no nonseparating curves, any curve in $L(u, v)$ must project nontrivially to S_1, and so a path in $L(u, v)$ from η_1 to η_n gives rise to a path in $\mathcal{AC}_1(S_1)$ of length on the order of n. It follows that $d_L(\eta_1, \eta_n)$ is at least on the order of n. As n was arbitrary, the diameter of $L(u, v)$ must be infinite. □

We record the following remark as it will be useful in the proof of Conjecture 2.1.

REMARK 4.3. If u is a separating curve on S that bounds a 3-holed sphere on one side and v is another curve representing a vertex of $\mathcal{SC}(S)$ so that $i(u, v) = 0$, then the proof of Lemma 4.2 implies that the diameter of $L(u, v)$ in $\mathcal{SC}(S)$ is infinite. Indeed, in this case, S_1 consists of either a single 3-holed sphere or a disjoint union of two 3-holed spheres. Thus there is no curve in $L(u, v)$ that does not project to $S \setminus S_1$, which is the only assumption used in the last two paragraphs of the above proof.

REMARK 4.4. If the genus of S is 1, then if u, v are disjoint nonseparating curves, then they must be jointly separating. Indeed, cutting along u (or v) produces a planar surface. Therefore, in the next subsection, it will suffice to assume that the genus of S is at least 2.

4.2. *The Jointly Non-separating Case*

Let u, v be disjoint curves in S such that u and v are jointly nonseparating. In this case, we show that the diameter of $L(u, v)$ in $\mathcal{N}_1(S)$ is infinite. This concludes the proof of Lemma 4.1, which in turn completes the proof of Proposition 3.1.

LEMMA 4.5. *Let $g \geq 2$. If $u, v \in \mathcal{N}_1(S)$ are disjoint curves that are jointly non-separating, then $L(u, v)$ has infinite diameter.*

Proof. Consider $S' = S \setminus (u \cup v)$. Since the genus of S is at least 2, S' is not a 3-holed sphere, and so by Proposition 2.2, $\mathcal{AC}_1(S')$ has infinite diameter. Let λ_u, λ_v denote simple closed curves on S so that λ_u (resp., λ_v) crosses u (resp., v) exactly once.

Choose a pseudo-Anosov mapping class $\phi \in \mathrm{Mod}(S')$. By Remark 2.3, given $n \in \mathbb{N}$, there exists $N \geq 1$ such that

$$d_{\mathcal{AC}_1(S')}(\pi_{S'}\lambda_u, \phi^N(\pi_{S'}\lambda_v)) > n.$$

Let $\lambda_v^{\phi^N}$ denote the simple closed curve on S obtained by turning $\phi^N(\pi_{S'}\lambda_v)$ into a simple closed curve by including its intersection point with u. Then, since any essential simple closed curve on S projects nontrivially to S', an $L(u, v)$-path from λ_u to $\lambda_v^{\phi^N}$ gives rise to a path in $\mathcal{AC}_1(S')$ of comparable length, between their projections. Thus we have produced vertices in $L(u, v)$ that are arbitrarily far apart in $L(u, v)$, and so $L(u, v)$ has infinite diameter, as desired. $\square$

4.3. *Proof of Theorem 1.3*

We are now in a position to prove Theorem 1.3.

THEOREM 1.3. *Suppose that $g \geq 1$ and that $(g, p) \neq (1, 2)$. Then the natural map*

$$\mathrm{Mod}^{\pm}(S) \to \mathrm{Aut}(\mathcal{N}_1(S))$$

is an isomorphism for $g \neq (2, 0)$ and a surjection with kernel $\mathbb{Z}/2\mathbb{Z}$ otherwise.

Proof. Let $f \in \mathrm{Aut}(\mathcal{N}_1(S))$. By Proposition 3.1 we can conclude that graph automorphisms of $\mathcal{N}_1(S)$ preserve edge types. Thus f induces a graph automorphism of $\mathcal{N}(S)$ and of $\mathcal{G}(S)$ by restriction. Since the vertex sets of $\mathcal{N}(S)$, $\mathcal{N}_1(S)$, and $\mathcal{G}(S)$ are the same, $\mathrm{Aut}(\mathcal{N}_1(S))$ injects into $\mathrm{Aut}(\mathcal{G}(S))$. Hence $\mathrm{Aut}(\mathcal{N}_1(S)) \leq \mathrm{Aut}(\mathcal{G}(S))$.

Then by Theorem A of [29] f is induced by a mapping class of S. Conversely, when $(g, p) \neq (2, 0)$, every mapping class gives rise to a distinct automorphism of $\mathcal{N}_1(S)$. When $(g, p) = (2, 0)$, mapping classes give rise to distinct graph automorphisms exactly when they reside in distinct cosets of the centralizer of $\mathrm{Mod}^{\pm}(S)$. □

Remark 4.6. When $g \geq 2$, Theorem 1.3 can also be proved by appealing to Theorems 1.1 and 1.2 of [15] together with Lemmas 3.3, 3.4, and 4.2 and Proposition 4.5.

5. Automorphisms of the Systolic Complex

The purpose of this section is to prove Theorem 1.2, which we do in the final subsection. Recall that when a surface S has multiple punctures, the graph $\mathcal{SC}(S)$ includes vertices representing separating curves that bound a 3-holed sphere on one side. Such a vertex is connected to another vertex v by an edge whenever the corresponding curves are disjoint or intersect exactly twice.

Two main tools in the proof of Theorem 1.2 are the following propositions. The first proposition characterizes the diameter of the link of a 2-edge (u, v) when both u and v are separating.

Proposition 5.1. *Suppose $g \geq 1$ and $(g, p) \neq (0, 5), (1, 3)$. If (u, v) is a 2-edge in $\mathcal{SC}(S)$ with u and v both separating, then* $\mathrm{diam}(L(u, v)) = 4$.

The second proposition characterizes the diameter of the link of a 2-edges (u, v) when exactly one of u and v is separating.

Proposition 5.2. *Suppose $g \geq 1$ and $(g, p) \neq (0, 5), (1, 3)$. If (u, v) is a 2-edge in $\mathcal{SC}(S)$ with u nonseparating and v separating, then* $\mathrm{diam}(L(u, v)) = 3$. *Furthermore, whenever $d_L(\alpha, \beta) = 3$, there exists a path of length 3 in $L(u, v)$ from α to β that passes through $\{b_1, b_2\}$.*

The proof of Proposition 5.1 will break into two pieces, Lemmas 5.3 and 5.5. Likewise, the proof of Proposition 5.2 will break into Lemmas 5.4 and 5.6.

5.1. Upper Bounds on the Diameter of the Link of a 2-Edge in $\mathcal{SC}(S)$

Let (u, v) be an edge in $\mathcal{SC}(S)$ with $i(u, v) = 2$; we refer to such an edge as a 2-*edge*. In this case, at least one of u, v is separating, and the subsurface obtained by thickening the union of u and v is necessarily a 4-holed sphere. As in the case of the punctured torus, the simple closed curves on a 4-holed sphere are also naturally parameterized by slopes in $\mathbb{Q} \cup \infty$, and without loss of generality, we identify u, v with the $1/0$ and $0/1$ curves.

LEMMA 5.3. *Suppose* $(g, p) \neq (0, 5)$. *If* (u, v) *is a* 2*-edge in* $\mathcal{SC}(S)$ *with both* u *and* v *separating, then* $\operatorname{diam}(L(u, v)) \leq 4$.

Proof. Let S_1 be the 4-holed sphere that forms the regular neighborhood of $u \cup v$. Denote the $1/1$ and $-1/1$ curves in S_1 by γ^+ and γ^-, respectively. Note that $\gamma^\pm$ both intersect u and v twice.

Since both u and v are separating, three of the four boundary components of S_1 necessarily correspond to punctures of S, which implies that both γ^+ and γ^- are separating and bound a 3-holed sphere on one side. Therefore both are elements of $L(u, v)$.

Now consider $\alpha \in L(u, v)$. Note that $|\alpha \cap S_1| \leq 2$. This is because every component of $\alpha \cap S_1$ has its endpoints on a single boundary component (the one that is not a puncture of S) and α intersects both u and v either 0 or 2 times, since u and v are both separating.

If $|\alpha \cap S_1| = 1$, then α will be at an $L(u, v)$-distance of 1 from $\gamma^\pm$.

If $|\alpha \cap S_1| = 2$, then $\alpha \cap S_1$ must be (up to homeomorphism) as pictured in Figure 10.

First, suppose that α is separating. Take one of the arcs of α contained in S_2. This can be concatenated with the blue arc shown in Figure 10 to obtain a simple closed curve η, which is disjoint from u and intersects each of v, α, and (without loss of generality) γ^+ twice. Based on the topology of S_2 and the arcs of $\alpha \cap S_2$, η can be chosen to be either nonseparating or a separating curve that bounds a 3-holed sphere on one side.

Thus $\eta \in L(u, v)$, $d_L(\alpha, \eta) = 1$, and $d_L(\eta, \gamma^+) = 1$. The possible configurations for η are shown in Figure 11.

Now suppose α is nonseparating. Then we we claim that it is possible to build a simple closed curve η contained entirely in S_2, which intersects α exactly once and is thus nonseparating (and therefore in $L(u, v)$). Indeed, take an essential arc λ in S_2 disjoint from α with endpoints on b_1 (the existence of λ is guaranteed because α is nonseparating). We can then concatenate λ with a subarc of b_1 that intersects α exactly once. This concatenation gives us our desired curve η, illustrated in Figure 12. It follows that α is at an $L(u, v)$-distance of at most 2 from both γ^+ and γ^-. □

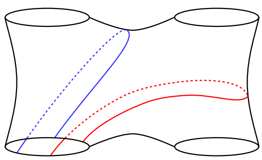

Figure 10 The intersection pattern of $\alpha \in L(u, v)$ with S_1 when $|\alpha \cap S_1| = 2$ and both u and v are separating.

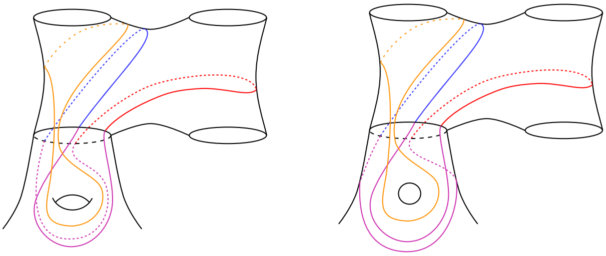

Figure 11 The curve η (drawn in orange) is either nonseparating (shown on the left) or separating (shown on the right) depending on the configuration of $\alpha \cap S_2$.

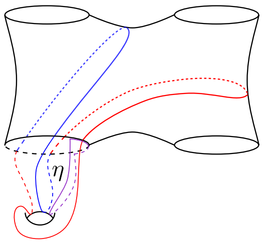

Figure 12 Constructing η when α is nonseparating.

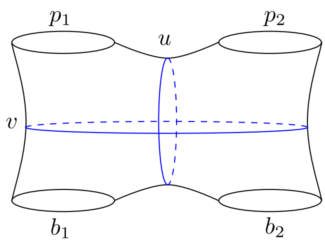

Figure 13 The regular neighborhood S_1 of u and v when u is nonseparating and v is separating.

LEMMA 5.4. *Suppose* $(g, p) \neq (0, 5)$. *If* (u, v) *is a* 2*-edge in* $\mathcal{SC}(S)$ *with* u *nonseparating and* v *separating*, *then* $\mathrm{diam}(L(u, v)) \leq 3$.

Proof. Let $\alpha, \beta \in L(u, v)$ with u nonseparating and v separating as shown in Figure 13. The boundary curves b_1 and b_2 of S_1 will play a crucial role in each case of this proof. Note that b_1 and b_2 are necessarily nonseparating. If they were both separating, then u would also be separating and if only one of them was separating, then the other would be a boundary component, and again u would be separating.

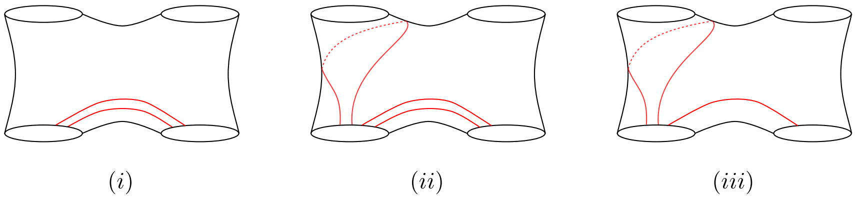

Figure 14 Possible intersection patterns of a curve in $L(u, v)$ with S_1 when exactly one of u and v is separating. Note that it is also possible for α to intersect S_1 in a submultiarc of one of the multiarcs pictured above. The arguments outlined below will apply to this situation.

We will make use of the following possible intersection patterns shown in Figure 14. In each case, we will show that α has an $L(u, v)$-distance of 1 from at least one of $\{b_1, b_2\}$. It is also possible for α to intersect S_1 in a submultiarc of one of the multiarcs pictured below, but this possibility is subsumed by the three arguments outlined below.

Case 1: Suppose α and β are both separating. Since they are separating, they are both in either configuration (i) or (ii). This implies that there exists $i \in \{1, 2\}$ such that α intersects b_i twice. The same is true for β, although not necessarily with the same value of i.

Assume first that $d_L(\alpha, b_1) = d_L(\beta, b_1) = 1$. Then $d_L(\alpha, \beta) \leq 2$.

Otherwise, $d_L(\alpha, b_1) = 1$ and $d_L(\beta, b_2) = 1$. Since b_1 and b_2 are disjoint, $d_L(b_1, b_2) = 1$. So $d_L(\alpha, \beta) \leq 3$.

Case 2: Suppose α is separating and β is nonseparating. Since β is nonseparating, it must be in configuration (iii). So there exists $i \in \{1, 2\}$ such that $d_L(\beta, b_i) = 1$. Since α is separating, it must be either in configuration (i) or (ii). By the same reasoning as in Case 1, $d_L(\alpha, b_i) = 1$ for some $i \in \{1, 2\}$. The argument then resolves in the same way as in Case 1.

Case 3: Suppose α and β are both nonseparating. Then α and β must be in configuration (iii). So there exists $i \in \{1, 2\}$ such that $d_L(\beta, b_i) = 1$. The same is true for β, although not necessarily with the same value of i. Once again, the argument resolves as in Case 1. □

5.2. *Lower Bounds on the Diameter of the Link of a 2-Edge in $\mathcal{SC}(S)$*

In this section, we finish the proofs of Propositions 5.1 and 5.2 by proving the following two lemmas.

LEMMA 5.5. *Suppose $g \geq 1$ and $(g, p) \neq (1, 3)$. If (u, v) is a 2-edge in $\mathcal{SC}(S)$ with both u and v separating, then* $\operatorname{diam}(L(u, v)) \geq 4$.

Proof. Let S_1 denote the regular neighborhood of $u \cup v$. As mentioned in the proof of Lemma 5.3, S_1 must be a sphere with four boundary components, exactly three of which are boundary components of S.

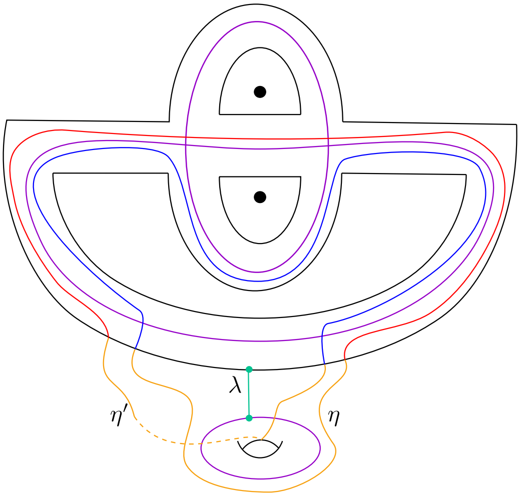

Figure 15 A diagram of S_1 and the construction of α in the event that S_2 has positive genus.

Since $(g, p) \neq (0, 5)$ or $(1, 3)$, the complementary subsurface $S_2 = S \setminus S_1$ is not a 3-holed sphere, and thus the diameter of $\mathcal{AC}_1(S_2)$ is infinite. As in the proof of Lemma 3.4, we will construct a pair of 2-curves α and β whose shortest path in $L(u, v)$ passes through $\{\gamma^+, \gamma^-\}$, both of which are in $L(u, v)$. Using the same notation as before, we will first specify the construction for $\alpha \cap S_2 = \{\eta, \eta'\}$.

The assumptions in the statement of the proposition guarantee that S_2 contains a nonseparating curve c of S, and let b be the essential boundary component of S_1 (i.e., the boundary component of S_1 not corresponding to a boundary component of S). Let λ be an embedded arc connecting b to c. We consider an arc $\eta \subset S_2$ with both endpoints on b, obtained by traveling along λ, then around c, and then back to b along the inverse of λ.

Note that S_2 has positive genus since we assumed that S has genus at least 1. This situation is illustrated in Figure 15. So on S_2 there must exist a second essential arc η' disjoint from η and whose endpoints link with those of η along b. Furthermore, we can choose η' so that it intersects c exactly once. As shown in Figure 16, the existence of η' is guaranteed so long as S_2 is not planar.

Now we specify the construction of $\alpha \cap S_1 = \{r_1, r_2\}$ as shown in Figure 10 (r_1 is the blue arc, and r_2 is the red arc). Choose r_1, r_2 so that the endpoints of r_1 coincide with one endpoint of η and with one of η', and similarly for r_2. Then the concatenation $\eta \cdot r_1 \cdot \eta' \cdot r_2$ yields a simple closed curve α, which is necessarily nonseparating; indeed, by construction it intersects c exactly once. Moreover, α intersects both γ^+ and γ^- more than twice.

By applying to α a high power of a mapping class fixing S_1 and acting as a pseudo-Anosov on S_2, we obtain a second nonseparating curve β whose projection to S_2 is arbitrarily far away from the projection of α to S_2 in $\mathcal{AC}_1(S_2)$. The

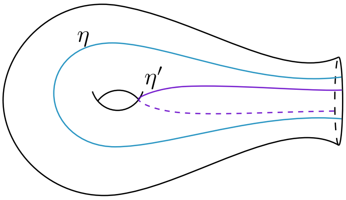

Figure 16 In any surface with at least one boundary component and genus $g \geq 1$, given an arc η constructed from a nonseparating curve as described above, there exists an arc η' disjoint from η and whose endpoints link with those of η.

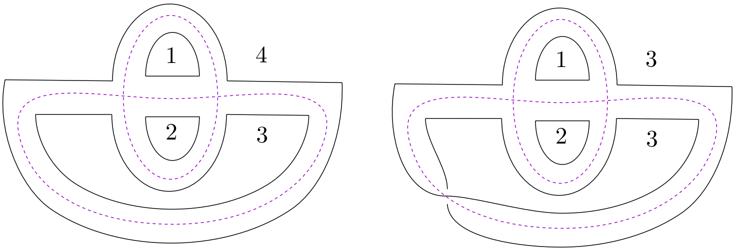

Figure 17 The possible configurations for u, v, and S_1 with u and v shown in dotted lines. On the left, S_1 is a sphere with 4 boundary components. On the right, S_1 is a torus with 3 boundary components, which can be ruled out by Euler characteristic considerations.

lemma now follows because both α and β are at least distance 2 from $\{\gamma^+, \gamma^-\}$, the only two curves in $L(u, v)$ that project trivially to S_2. □

LEMMA 5.6. *Suppose $g \geq 1$ and $(g, p) \neq (1, 3)$. If (u, v) is a 2-edge in $\mathcal{SC}(S)$ with u nonseparating and v separating, then* $\operatorname{diam}(L(u, v)) \geq 3$. *Furthermore, whenever $d_L(\alpha, \beta) = 3$, there exists a path of length 3 in $L(u, v)$ from α to β that passes through $\{b_1, b_2\}$.*

Proof. We note that the last sentence is implied by the proof of Lemma 5.4: in each of the cases outlined there, there exists a path of length 3 connecting α to β that passes through b_1 or b_2.

Without loss of generality, we may assume that v is a separating curve bounding two punctures of S. Let S_1 denote the regular neighborhood of $u \cup v$. See Figure 17. We claim that S_1 is topologically a sphere with four boundary components. Indeed, S_1 has at least three boundary components, two of which correspond to the pair of punctures on one side of v. There are three boundary components if and only if S_1 has positive genus, but in that case, its Euler characteristic is $2 - 2 - 3 = -3$. Each complementary region of $u \cup v$ on S_1 is either a disk or a punctured disk and as a graph on S_1 with vertices corresponding to intersection points, $u \cup v$ is 4-valent and therefore has twice as many edges as vertices. It follows that

$$\chi(S_1) = -3 = i(u, v) - 2i(u, v) + D,$$

where $D \geq 0$ is unknown. Hence $3 \leq i(u, v)$, a contradiction. So S_1 is a 4-holed sphere with two boundary components corresponding to boundary components of S and two boundary components b_1 and b_2 corresponding to boundaries of $S_2 = S \setminus S_1$. The curves b_1 and b_2 are necessarily nonseparating as shown in the proof of Lemma 5.4.

We will construct a pair of curves α and β so that any path from one to the other such that each vertex projects nontrivially to S_2 must be longer than 3. On the other hand, b_1 and Eb_2 are the only two curves in $L(u, v)$ that do not project to S_2. Note that $\gamma^{\pm}$ are *not* in $L(u, v)$ since both intersect u twice and if either one of them is separating, it can not bound a 3-holed sphere on either of its sides by the topological assumptions made on S in the statement of the lemma.

Let α be a 2-curve in configuration (iii) and let $\alpha \cap S_2 = \{\eta, \eta'\}$. Letting ρ denote a homeomorphism S_2 taking b_1 to b_2, let $\beta = \rho(\alpha)$. Therefore both β and α are nonseparating; α (resp., β) intersects b_1 (resp., b_2) three times; and α (resp., β) intersects b_2 (resp., b_1) exactly once.

We claim that there is $K > 0$ such that if

$$d_{\mathcal{AC}(S_2)}(\pi_{S_2}(\alpha), \pi_{S_2}(\beta)) > K,$$

then if $\Gamma = \{\alpha, v_1, \ldots, v_n, \beta\}$ is any path in $L(u, v)$ so that v_i projects to S_2 nontrivially, then $n > 2$. Indeed, $i(v_j, v_{j+1}) \leq 2$, and thus Γ may not be a path in $\mathcal{AC}(S_2)$, but it constitutes a sequence that makes uniformly bounded jumps. The existence of Γ therefore implies an upper bound on the distance in $\mathcal{AC}(S_2)$ between the projections of α and β that depends only on n.

Now assume that the projections of α, β to S_2 are at an $\mathcal{AC}(S_2)$ distance of at least K, where K is as above. It follows that a shortest path from α to β must pass through at least one of b_1, b_2 since the $L(u, v)$ distance between them is no more than 3. Since $d_L(\alpha, b_1)$ and $d_L(\beta, b_2)$ are both at least 2, a shortest path from α to β cannot have length 2 and therefore must be of length 3. □

In the context of the previous lemma, the curves $\{b_1, b_2\}$ play the role of a *shortcut set,* a notion we will introduce formally in Section 6. Furthermore, the proof of Lemma 6.7 applies in our setting here exactly as written, and shows the following:

LEMMA 5.7. *If (u, v) is a 0-edge such that both u and v are non-separating, then the diameter of $L(u, v)$ is either infinite, or it is 3. In the latter case, there can not exist a pair of curves b_1, b_2 in $L(u, v)$ so that any shortest path of length 3 in $L(u, v)$ passes through either b_1 or b_2.*

5.3. Proof of Schaller's Conjecture

We are now ready prove Conjecture 2.1.

THEOREM 1.2. *The automorphism group of $\mathcal{SC}(S)$ is isomorphic to* $\mathrm{Mod}^{\pm}(S)$.

Proof. We begin by showing that automorphisms of $\mathcal{SC}(S)$ preserve 0-edges. Let (u, v) be a 0-edge, and assume first that S has genus at least 1.

If u and v are both nonseparating, then either $\operatorname{diam}(L(u, v)) = \infty$ or $\operatorname{diam}(L(u, v)) = 3$, and in this case, $L(u, v)$ does not possess a shortcut set—curves for which any shortest path of length 3 must pass through at least one of them. So (u, v) cannot be sent to a 1-edge, since links of 1-edges have diameter 4, nor can it be sent to a 2-edge (since links of 2-edges either have diameter 4 or diameter 3 and possess the curves b_1, b_2).

If one or both of u and v are separating, then $\operatorname{diam}(L(u, v)) = \infty$. So again, (u, v) cannot be sent to a 1- or 2-edge.

If S has genus 0, then all curves are separating, and we only need to distinguish between 0-edges and 2-edges. Note that links of 0-edges have infinite diameter and links of 2-edges (with both curves separating) have diameter at most 4. Thus automorphisms of $\mathcal{SC}(S)$ preserve 0-edges, as desired.

To complete the proof, it remains to show that any automorphism ψ of the subgraph of $\mathcal{SC}(S)$ consisting of all vertices but only 0-edges is induced by a mapping class. If $g = 0$, then this amounts to saying that the graph with vertices corresponding to separating curves that bound a 3-holed sphere on one side and with edges corresponding to disjointness, has automorphism group isomorphic to the extended mapping class group. This follows readily from McLeay's extension [26] of Brendle–Margalit's work [6] on *complexes of regions,* and we outline the idea as follows. A *complex of regions* is a simplicial complex whose vertices are essential subsurfaces chosen from some specific subset of mapping class group orbits of all subsurfaces, and edges are determined by disjointness. When $g = 0$, we can interpret the subgraph of $\mathcal{SC}(S)$ corresponding only to 0-edges as a complex of regions, where each vertex represents a 3-holed sphere. Theorem 2 of [26] states that the automorphism group of a complex of regions is isomorphic to the extended mapping class group when every *minimal* vertex (roughly speaking, a subsurface that does not properly contain another subsurface representing a vertex of that complex of regions) is *small,* a technical assumption that amounts to saying that the topology of the entire surface is sufficiently complicated relative to the topology of the subsurface. In our context, every vertex will be minimal, and every vertex will be small so long as $p \geq 7$. The only remaining genus 0 case is $p = 6$; since the argument involves ideas from [26] that are not relevant to the rest of the remaining cases of Theorem 1.2, we cover this in Appendix B.

If $g \neq 0$, then ψ induces an automorphism of $\mathcal{N}(S)$. Indeed, any nonseparating curve is involved in a 0-edge of diameter 3, but a 0-edge involving a separating curve has infinite diameter. Thus there is a mapping class f such that ψ coincides with f when restricted to the nonseparating vertices. Consider $\phi \circ f^{-1}$. This is an automorphism of $\mathcal{SC}(S)$ that pointwise fixes each vertex corresponding to a nonseparating curve.

Let u be a vertex of $\mathcal{SC}(S)$ corresponding to a separating curve, let $v = \phi \circ f^{-1}(u)$, and suppose $v \neq u$.

Claim. *There exists a nonseparating curve γ that is disjoint from u and intersects v at least three times.*

Proof. Suppose that u and v are disjoint. Note that since S has genus at least 1, the complement of u containing v has genus at least 1. So there exists a nonseparating curve γ disjoint from u that intersects v arbitrarily many times.

Suppose u and v intersect. Once again, the component of the complement of u which is not a 3-holed sphere, call it S', has genus at least 1. Note that $v \cap S'$ is an essential multiarc, call it ν. We can then pick any component of ν and find a nonseparating curve γ in S' that intersects it arbitrarily many times. □

By the claim there exists a nonseparating curve γ that is at distance 1 from u but at distance > 1 from v. This is a contradiction since $\phi \circ f^{-1}(\gamma) = \gamma$. Thus $u = v$, and therefore $\phi \circ f^{-1}$ is the identity on the full vertex set of $\mathcal{SC}(S)$, as desired. □

6. The k-Curve Graph

We are now ready to show that the automorphism group of the k-curve graph is the extended mapping class group for $|\chi(S)|$ sufficiently large with respect to k. Throughout this section, we will assume that S is a connected oriented surface with negative Euler characteristic. As before, we call edges in $\mathcal{C}_k(S)$ that connect vertices admitting disjoint representatives *0-edges.* We call all other edges *nonzero edges.* Distinguishing between 0-edges and nonzero edges in $\mathcal{C}_k(S)$ is a more delicate process than distinguishing between 0 and 1-edges in $\mathcal{N}_1(S)$ and $\mathcal{SC}(S)$. In addition to the diameter, we will also record two other properties of the edge links. First, we will consider the cardinality of the edge links, namely whether the edge link contains a finite or infinite number of vertices. Second, we will define a finite collection of curves associated with an edge link called a *shortcut set,* whose existence (or nonexistence) will be our final tool for distinguishing between edge links (see Figure 18). Throughout this section, we assume that S satisfies $|\chi(S)| \geq k + 512$ unless specified otherwise. See Appendix (particularly, Remark A.3) for an explanation of the relevance of this inequality.

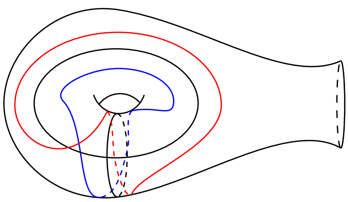

Figure 18 A shortcut set is a generalization of $\gamma^{\pm}$ in the case $k = 1$. Indeed, when $i(u, v) = 1$, $\text{diam}(L(u, v)) = 4$, and as demonstrated in Lemma 3.3, there is a path of length at most 4 between them that passes through $\{\gamma^{+}, \gamma^{-}\}$.

Table 1 Strategy for distinguishing between 0-edges and nonzero edges. N/A indicates that the existence of a shortcut set was not checked for these edge links. Note that we do not include the case of a nonzero edge with u, v filling since if $i(u, v) \leq k$ and $|\chi(S)| \geq k + 512$, then they necessarily do not fill.

Edge-type of (u, v)	$\lvert L(u, v)\rvert$	$\mathrm{diam}(L(u, v))$	Shortcut Set
0-edge with u, v jointly nonseparating	∞	∞	N/A
0-edge with u, v jointly separating			
i) Neither component of $S \setminus (u \cup v)$ is a 3-holed sphere	∞	3	Does not exist
ii) At least one component of $S \setminus (u \cup v)$ is a 3-holed sphere	∞	∞	N/A
nonzero edge with u, v not filling	∞	1, 2, or 4	N/A
nonzero edge with u, v not filling	∞	3	Exists

The goal of this section is to prove the partition of edge types shown in Table 1 based on the three characteristics that we have just outlined.

6.1. Diameter and Cardinality of Links

We will now compute the diameters of various types of edges.

Let (u, v) be an edge in $\mathcal{C}_k(S)$. We will begin by considering the case where u and v are a filling pair.

LEMMA 6.1. *If (u, v) is an edge in $\mathcal{C}_k(S)$ such that $u \cup v$ fills S, then there are finitely many vertices in $L(u, v)$.*

Proof. Since u and v fill, their union gives rise to the 1-skeleton of a polygonal decomposition of S where some polygons may be once-punctured. Any other essential curve γ can be isotoped to be in minimal position with respect to $u \cup v$, and so it defines an equivalence class of cyclically ordered sequences each of length $i(\gamma, u \cup v)$; we simply read off the edges of the polygonal decomposition in accordance with the order in which γ meets them. However, this does not yield a uniquely defined cyclic sequence because γ can be homotoped over a vertex of one polygon and into another. We will consider any two sequences related in this way to be equivalent.

There are at most finitely many sequences of edges in the polygonal decomposition of length at most k, and therefore there are at most finitely many (equivalence classes of) cyclic sequences of length at most k. This implies that there are at most finitely many curves that intersect both u and v at most k times. □

We will next consider the links of nonzero edges (u, v) when u and v are not a filling pair.

Lemma 6.2. *If (u, v) is a nonzero edge in $\mathcal{C}_k(S)$ such that u and v do not fill S, then*

$$\operatorname{diam}(L(u, v)) \le 4.$$

Proof. Let $\alpha \in L(u, v)$, and let $F(u, v)$ be the subsurface of S filled by u and v (potentially with some complementary disks glued back in so as to make F essential). Note that $u \cup v$ can be thought of as a 4-valent graph Γ with vertices in $u \cap v$ and edges given by arcs of either u or v running between intersection points. Note that Γ has exactly twice as many edges as vertices. Since $F(u, v)$ is a thickening of Γ, we get

$$\chi(F(u, v)) = -i(u, v).$$

It follows that $|\chi(F(u, v))| \le k$, and hence

$$|\chi(S \setminus F(u, v))| \ge 512. \tag{1}$$

There exists $\gamma \in L(u, v)$ such that $\gamma \subset F(u, v)$. This can be seen, for example, by surgering along intersections of u and v as in Hempel's argument (see [13, Lemma 2.1]). If $\alpha \subset S \setminus F(u, v)$, then $d_L(\alpha, \gamma) = 1$. Otherwise, if $\alpha \subset F(u, v)$, there exists a simple closed curve $\beta \subset S \setminus F(u, v)$ so that $\beta \in L(u, v)$, since u and v do not fill S. This yields a path between α and γ in $L(u, v)$ of length 2. Hence $d_L(\alpha, \gamma) \le 2$.

Lastly, if α nontrivially intersects both $F(u, v)$ and its complement, then we consider the multiarc formed by $\alpha \cap (S \setminus F(u, v))$. Abusing notation slightly, we will denote this multiarc by α. Lemma A.1 implies that there exists some essential simple closed curve $\eta \subset S \setminus F(u, v)$ such that $i(\alpha, \eta) \le k$. Thus $d_L(\alpha, \eta) = 1$, and so $d_L(\alpha, \gamma) \le 2$, since $d_L(\eta, \gamma) = 1$.

It follows from the above cases that the diameter of $L(u, v)$ is at most 4. □

Next we consider the diameter of a 0-edge (u, v) when u and v are jointly nonseparating.

Lemma 6.3. *Let $u, v \in \mathcal{C}_k(S)$. If (u, v) is a 0-edge such that u and v are jointly nonseparating, then $L(u, v)$ has infinite diameter.*

To prove Lemma 6.3 we first need to establish the following quasi-isometry between $\mathcal{AC}(S)$ and $\mathcal{AC}_k(S)$.

Proposition 6.4. *Let S be a surface of genus at least 2. Consider $\mathcal{AC}_k(S)$, the graph with the same vertex set as $\mathcal{AC}(S)$ and with edges connecting arcs and curves that intersect essentially at most k times. Then $\mathcal{AC}(S) \cong_{\mathrm{QI}} \mathcal{AC}_k(S)$.*

Proof. Let α and β be two vertices in $\mathcal{AC}_k(S)$. We use $d_{\mathcal{AC}_k}(\cdot, \cdot)$ to denote the distance in $\mathcal{AC}_k(S)$ and $d_{\mathcal{AC}}(\cdot, \cdot)$ the distance in $\mathcal{AC}(S)$. Let $\phi : \mathcal{AC}_k(S) \to \mathcal{AC}(S)$ be the identity map on the vertices. Since every edge in $\mathcal{AC}(S)$ is also present in $\mathcal{AC}_k(S)$, we have that

$$d_{\mathcal{AC}_k}(\alpha, \beta) \le d_{\mathcal{AC}}(\phi(\alpha), \phi(\beta)).$$

On the other hand, for any α, β connected by a nonzero edge in $\mathcal{AC}_k(S)$, by surgering along the intersections of α and β as in Hempel's argument (see [13] Lemma 2.1), there is a path between α and β consisting of only 0-edges of length at most $2\log_2(k)+2$ in $\mathcal{AC}_k(S)$. This path is mapped bijectively into $\mathcal{AC}(S)$ by ϕ. Thus

$$d_{\mathcal{AC}}(\phi(\alpha), \phi(\beta)) \leq (2\log_2(k) + 2) \cdot d_{\mathcal{AC}_k}(\alpha, \beta). \qquad \square$$

We can now prove Lemma 6.3.

Proof of Lemma 6.3. Let u and v be jointly nonseparating disjoint curves on S, and consider $S' = S \setminus (u \cup v)$. By Propositions 2.2 and 6.4, we know that $\mathcal{AC}_k(S') \cong_{\mathrm{QI}} \mathcal{AC}(S') \cong_{\mathrm{QI}} C(S')$. So we can consider a coarsely well-defined projection

$$\tau : L(u, v) \to \mathcal{AC}_k(S')$$

defined as follows. If $\alpha \in L(u, v)$ and $i(\alpha, u) = i(\alpha, v) = 0$, then $\tau(\alpha) = \alpha \in \mathcal{AC}_k(S')$. Otherwise, send α to the multiarc representing its intersection with S', which is a simplex in $\mathcal{AC}_k(S)$. It follows from the definition of τ that

$$d_L(\alpha, \beta) \geq d_{\mathcal{AC}_k(S')}(\tau(\alpha), \tau(\beta)). \tag{2}$$

Consider a nonseparating curve $\gamma \in \mathcal{C}(S')$, and let $\phi : S \to S$ be a map fixing u and v pointwise and which restricts to a pseudo-Anosov on S'. Then for any $N \in \mathbb{N}$, there exists n such that

$$d_{\mathcal{C}(S')}(\gamma, \phi^n(\gamma)) \geq N.$$

Since $\mathcal{AC}_k(S')$ is quasi-isometric to $\mathcal{C}(S')$, we can choose appropriate n to make $d_{\mathcal{AC}(S')}(\gamma, \phi^n(\gamma))$ in $\mathcal{AC}_k(S')$ arbitrarily large. By inequality (2) the diameter of $L(u, v)$ is infinite. $\square$

6.2. *Shortcut Sets*

We will now make precise the definition of a *shortcut set*. We then use it to distinguish between the remaining cases.

DEFINITION 6.5. Given $L(u, v)$ with $\mathrm{diam}(L(u, v)) = R < \infty$, a *shortcut set* for $L(u, v)$ is a finite set of curves $\{\gamma_0, \ldots, \gamma_n\}$ with the following properties:

(1) $\gamma_i \in L(u, v)$ for all i, and,
(2) given any $\alpha, \beta \in L(u, v)$ with $d_L(\alpha, \beta) = R$, there exists a path of length R between α and β that passes through at least one of the γ_i.

We can now prove the statement given in the fifth row of Table 1.

PROPOSITION 6.6. *Given $u, v \in \mathcal{C}_k(S)$ such that u, v do not fill S, if (u, v) is a nonzero edge with* $\mathrm{diam}(L(u, v)) = 3$, *then there exists a shortcut set for $L(u, v)$.*

Proof. Let $\Gamma = \{\gamma_0, \ldots, \gamma_n\}$ be the set of curves in $L(u, v)$ entirely contained in $F(u, v)$, which is finite by Lemma 6.1. We claim that they form a shortcut set.

Let $\alpha, \beta \in L(u, v)$ with $d_L(\alpha, \beta) = 3$. We will construct a path of length 3 between α and β that contains at least one $\gamma_i \in \Gamma$. This is trivially true if either α or β is contained in the subsurface $F(u, v)$.

Claim. *Either α or β has distance 1 from Γ.*

Proof. The claim is clear if α or β is contained in $S \setminus F(u, v)$. Consider when α and β intersect both $F(u, v)$ and its complement nontrivially. Assume by contradiction that both α and β are at $L(u, v)$-distance at least 2 from every curve in Γ. We can then replace α with its image α' under a high power of a mapping class ϕ that restricts to the identity on $F(u, v)$ and acts as a pseudo-Anosov on $S \setminus F(u, v)$. By choosing a sufficiently large power we can assume that

$$d_{\mathcal{AC}(S \setminus F(u,v))}(\alpha', \beta) > 4.$$

Observe that α' is still at least distance 2 from Γ because $\alpha \cap F(u, v) = \alpha' \cap F(u, v)$.

Let $\{\alpha', v_0, v_1, \beta\}$ be a path in $L(u, v)$ from α' to β, where v_0 and v_1 are not necessarily distinct (such a path always exists since $\operatorname{diam}(L(u, v)) = 3$). By assumption, β is at least $L(u, v)$-distance 2 from every curve in Γ, so $v_0, v_1 \notin \Gamma$. Then both v_0 and v_1 project nontrivially to $S \setminus F(u, v)$, which yields a path of length at most 3 in $\mathcal{AC}(S \setminus F(u, v))$ between the projections of α' and β. This is a contradiction, which finishes the proof of the claim. □

Now, without loss of generality, assume that α is adjacent to some γ_i. Then $\beta \cap F(u, v) \neq \emptyset$. Otherwise, β is disjoint from γ_i, and $d_L(\alpha, \beta) \leq 2$, a contradiction. By Lemma A.1 there exists a simple closed curve $\eta \in S \setminus F(u, v)$ that intersects β no more than k times. Thus η is adjacent to both γ_i and β in the link, and we have a desired path of length 3 from α to β passing through the shortcut set. □

The following two propositions, together with Lemma 6.3, establish the statements in the first, third, and fourth rows of Table 1.

Proposition 6.7. *Let $u, v \in \mathcal{C}_k(S)$ be nonseparating. If (u, v) is a 0-edge such that u and v jointly separate S, then:*

(i) *If neither component of $S \setminus (u \cup v)$ is a 3-holed sphere, then $L(u, v)$ has diameter 3 and does not have a shortcut set.*

(ii) *Otherwise, $L(u, v)$ has infinite diameter.*

Proof. Denote the two components of $S \setminus (u \cup v)$ by S_1 and S_2. Let $\alpha, \beta \in L(u, v)$ be distinct.

We begin by considering (i). By assumption we know that neither S_1 nor S_2 is a three-holed sphere. Note that if α and β are both contained in the same component of $S \setminus (u \cup v)$, say S_1, then $d_L(\alpha, \beta) \leq 2$, since there is an essential curve contained in S_2 and thus disjoint from both α and β. If α and β are contained in different components of $S \setminus (u \cup v)$, then $d_L(\alpha, \beta) = 1$ since they are disjoint.

Next, suppose that $\alpha \subset S_1$ and that β intersects both S_1 and S_2 nontrivially. Then by Lemma A.1 there is an essential curve η contained entirely in either S_1

or S_2 such that $i(\eta, \beta) \leq k$. If $\eta \subset S_2$, then it follows that $d_L(\alpha, \beta) \leq 2$. If $\eta \subset S_1$, then let $\rho \subset S_2$ be any essential curve, which exists because S_2 is not a 3-holed sphere. Then $\{\alpha, \rho, \eta, \beta\}$ is a length 3 path in $L(u, v)$ from α to β.

Finally, suppose α and β intersect both S_1 and S_2 nontrivially.

We now construct a path $\{\alpha, \rho, \eta, \beta\}$ of length 3 between α and β in $L(u, v)$. There are three possibilities:

(1) the projection $\pi_{S_1}(\alpha)$ consists of a single weighted arc. In this case, we let ρ be a curve on S_1 disjoint from α. This is always possible since S_1 is not a 3-holed sphere. See Figure 19.

(2) $\pi_{S_1}(\alpha)$ consists of multiple nonhomotopic essential arcs, but every arc in $\pi_{S_1}(\alpha)$ begins and ends at two different boundary components (i.e., each arc intersects u and v only once). In this case, we construct ρ in the following manner: take two nonhomotopic arcs c_1 and c_2 in $\pi_{S_1}(\alpha)$ whose endpoints on u are adjacent to each other among all arcs in $\pi_{S_1}(\alpha)$. Concatenate c_1 and c_2 first with the subarc of u that contains no endpoints of other arcs, and then with a subarc of v that connects the two other endpoints of c_1 and c_2. The concatenation ρ is an essential simple closed curve since we assume that c_1 and c_2 are nonhomotopic. See Figure 20.

(3) $\pi_{S_1}(\alpha)$ consists of multiple nonisotopic essential arcs, but some arc in $\pi_{S_1}(\alpha)$ begins and ends at the same component. Without loss of generality, we assume that there exists an arc $c \in \pi_{S_1}(\alpha)$ intersecting u twice. In this case,

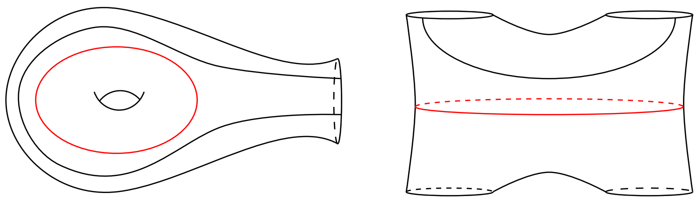

Figure 19 Since S_1 is not a 3-holed sphere, it either has genus or at least 4 punctures. In either case, given any arc, there always exists an essential curve disjoint from it.

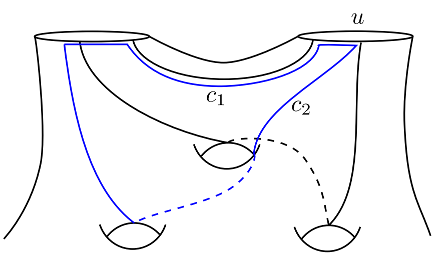

Figure 20 Construction of ρ in the case that each arc of $\pi_{S_1}(\alpha)$ has endpoints on distinct boundary components and there exist at least 2 arcs in the projection that are not parallel.

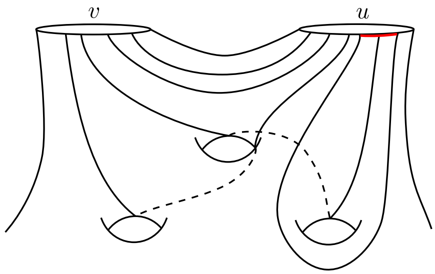

Figure 21 In the event that at least one arc in $\pi_{S_1}(\alpha)$ has both endpoints on u, we concatenate it with the red subarc of u to form the desired ρ.

we construct ρ by concatenating c with the subarc of u that makes ρ essential. This is always possible since S_1 is not a three-holed sphere as in Figure 19. Since ρ is disjoint from both u and v, $\rho \in L(u, v)$. Meanwhile, $d_{L(u,v)}(\rho, \alpha) = 1$ since $i(\alpha, \rho) \leq i(\alpha, u) \leq k$. See Figure 21.

Note that the curve ρ will not be peripheral except in the case that S_1 is a punctured annulus bounded by u and v. In this case, we choose ρ so that it differs from u by a single puncture on the interior of S_1. Since S_1 is not a three-holed sphere, u and v are not homotopic.

We therefore have constructed a curve ρ that is completely contained in S_1 and adjacent to α in $L(u, v)$. A similar curve $\eta \subset S_2$ adjacent to β may be constructed. This gives a path of length 3 between α and β in $L(u, v)$ as desired.

Now consider partial pseudo-Anosov φ_i that act as pseudo-Anosov on S_i and as the identity on $S \setminus S_i$ for $i = 1, 2$. Let $\alpha' = \varphi_2^{n_2}(\varphi_1^{n_1}(\alpha))$ where n_1, n_2 are chosen sufficiently large so as to guarantee $d_{\mathcal{AC}_k(S_i)}(\alpha', \beta) > 3$. We can do this since the diameter of $\mathcal{AC}_k(S_i)$ is infinite for $i = 1, 2$ by Proposition 6.4.

Note that we can construct a path $\{\alpha', \rho, \eta, \beta\}$ between α' and β in the same way as above. We will now show that such a path is minimal in $L(u, v)$. Suppose to the contrary that there exists a path of length 2 between α' and β, say $\{\alpha', \psi, \beta\}$. Without loss of generality, $\pi_{S_1}(\psi)$ is nontrivial, so $d_{\pi_{S_1}(L)}(\alpha', \beta) = 2$, a contradiction. A path of length 1 is ruled out for the same reason. Thus the shortest path between α' and β has length 3, and $\mathrm{diam}(L(u, v)) = 3$.

It remains to show that there does not exist a shortcut set in $L(u, v)$. Let $\Gamma = \{\gamma_1, \gamma_2, \ldots, \gamma_n\}$ be any finite set of curves in $L(u, v)$. Since Γ is a finite set, its image in $\mathcal{AC}_k(S_i)$ has finite diameter.

Let ρ_i be an essential arc in S_i that has one end point on u and the other on v. Fix $\beta_i = \varphi_i^n(\rho_i)$ with $n \in \mathbb{N}$ large enough so that $d_{\mathcal{AC}_k(S_i)}(\beta_i, \pi_i(\Gamma)) > 3$, where φ_i are defined as above. Now let $\alpha_i = \varphi_i^m(\rho_i)$ where $m \in \mathbb{N}$ is chosen sufficiently large so that $d_{\mathcal{AC}_k(S_i)}(\alpha_i, \beta_i) > 6$. Note that this implies $d_{\mathcal{AC}_k(S_i)}(\alpha_i, \pi_i(\Gamma)) > 3$.

Let α and β be the curves obtained by concatenating α_1 with α_2 and β_1 with β_2, respectively. Note that α and β intersect each of u and v exactly once and are therefore contained in $L(u, v)$ and $d_{L(u,v)}(\alpha, \beta) = 3$, by construction. Additionally, if α or β were adjacent to some $\gamma_j \in \Gamma$ in $L(u, v)$, then one of the α_i or β_i would be adjacent to $\pi_i(\gamma_j)$ in $\mathcal{AC}_k(S_i)$, a contradiction. Therefore α and β are at least distance 2 from Γ, and therefore there is no path of length 3 from α to β that goes through Γ. We conclude that there does not exist a shortcut set on S in $L(u, v)$.

We now finish the proof by considering (ii). Suppose exactly one of the connected components of $S \setminus (u \cup v)$ is a three-holed sphere, say S_1. Then we can choose essential simple closed curves α, β contained in S_2 such that $d_{\mathcal{AC}_k(S_2)}(\alpha, \beta)$ is arbitrarily large. Let $P = \{\alpha, v_1, \ldots, v_n, \beta\}$ be any path of shortest length between α and β in $L(u, v)$. Then each v_i for $1 \leq i \leq n$ projects nontrivially to S_2. Therefore $\pi_{S_2}(P) = \{\pi_{S_2}(\alpha), \pi_{S_2}(v_1), \ldots, \pi_{S_2}(v_n), \pi_{S_2}(\beta)\}$ is a path in $\mathcal{AC}_k(S_2)$. Observe that $\pi_{S_2}(\alpha) = \alpha$ and $\pi_{S_2}(\beta) = \beta$, which implies $|P| \geq d_{\mathcal{AC}_k(S_2)}(\alpha, \beta)$. So there exists curves $\alpha, \beta \in L(u, v)$ that are arbitrarily far apart in $L(u, v)$, and so $L(u, v)$ has infinite diameter. □

It remains to consider the possibility that (u, v) is a 0-edge with at least one of u, v separating:

PROPOSITION 6.8. *Let $u, v \in \mathcal{C}_k(S)$ such that at least one of them is separating. If (u, v) is a 0-edge, then:*

(i) *If at least two components of $S \setminus (u \cup v)$ are not 3-holed spheres, then $L(u, v)$ has diameter 3 and does not have a shortcut set.*
(ii) *Otherwise, $L(u, v)$ has infinite diameter.*

Proof. This argument follows the logic of the proof of Proposition 6.7 very closely; we include it in full for the ease of the reader.

Suppose that at least two components of $S \setminus (u \cup v)$ are not 3-holed spheres. Note that there can be at most three components of $S \setminus (u \cup v)$; if there are exactly two components that are not 3-holed spheres, then denote them by S_1, S_2. In this scenario the exact argument used in (i) of Proposition 6.7 applies here.

We also note that if we are in scenario (ii), then the argument used in Proposition 6.7 applies here verbatim.

It remains to consider the possibility that there are three components of $S \setminus (u \cup v)$ all of which are not 3-holed spheres. Denote these components by S_1, S_2, S_3. Let α, β be curves in $L(u, v)$. Suppose that α is disjoint from some S_i. Using Lemma A.1, there exists some η intersecting β at most k times and contained entirely in some S_j. If $j = i$, then $d_L(\alpha, \beta) \leq 2$. Otherwise, choose $\rho \subset S_i$ arbitrarily and note that $\{\alpha, \rho, \eta, \beta\}$ is a length 3 path.

Assume next that α, β both intersect each S_i. This is analogous to the case in the previous proof corresponding to α, β intersecting both S_1 and S_2. The argument used in this setting in the previous proposition yielded a curve ρ contained entirely inside S_1 and disjoint from α. The exact same argument can be used here to produce a curve contained entirely inside either S_1, S_2, S_3 and disjoint from α. By symmetry we can also produce a curve disjoint from β and contained entirely in one of the other components.

For the lower bound on diameter, choose partial pseudo-Anosovs ϕ_i on S_i and choose some $\alpha \in L(u, v)$ intersecting each S_i. Then, as in the previous proof, choose n_i, $i = 1, 2, 3$, sufficiently large so as to guarantee that

$$d_{\mathcal{AC}_k(S_i)}(\alpha, \phi_1^{n_1}(\phi_2^{n_2}(\phi_3^{n_3}(\alpha))) > 3.$$

Then as in the previous argument, no length 2 path can exist between α and $\beta := \phi_1^{n_1}(\phi_2^{n_2}(\phi_3^{n_3}(\alpha)))$ since any curve ψ realizing this path has to project nontrivially to at least one S_i and this would yield a path of length 2 between the projections of α and β in that subsurface.

Finally, we will show that there does not exist a shortcut set in $L(u, v)$. Consider any finite set of curves $\Gamma = \{\gamma_1, \gamma_2, \ldots \gamma_n\}$ with each $\gamma_i \in L(u, v)$. Note that Γ has finite diameter in $\mathcal{AC}_k(S_i)$. The idea is to construct curves $\alpha, \beta \in L(u, v)$ that are at distance exactly 3 from each other, but are also at distance at least 2 from Γ. This rules out any curve in Γ from appearing in any length of path 3 between α and β, and we can conclude that there is no shortcut set for $L(u, v)$. In fact, the proof can be carried out in the exact same way as in the previous proposition, since we are assuming that all components of $S \setminus (u \cup v)$ are not 3-holed spheres. □

We are now in a position to prove our main result.

THEOREM 1.1. *Suppose $|\chi(S)| \geq k + 512$. Then the natural map*

$$\mathrm{Mod}^{\pm}(S) \to \mathrm{Aut}(\mathcal{C}_k(S))$$

is an isomorphism.

Proof. For any edge (u, v) in $\mathcal{C}_k(S)$, note that any automorphism of $\mathcal{C}_k(S)$ preserves the cardinality of $L(u, v)$, the diameter of $L(u, v)$, and the existence of a shortcut set for $L(u, v)$. Therefore by Lemmas 6.2 and 6.3 and Propositions 6.6 and 6.7 we obtain all the statements in Table 1. This implies that every automorphism of $\mathcal{C}_k(S)$ sends 0-edges to 0-edges and nonzero edges to nonzero edges. Hence any automorphism of $\mathcal{C}_k(S)$ induces an automorphism of the curve graph $\mathcal{C}(S)$ and therefore corresponds to a mapping class. The other direction of the isomorphism is clear, and hence the theorem follows. □

A. Where Does 512 Come from?

The following result is necessary for the proofs of Propositions 6.6 and 6.7 and introduces the restriction $|\chi(S)| \geq k + 512$, which appears in Theorem 1.1 (see Remark A.3).

LEMMA A.1. *There exists some constant $D > 0$ satisfying the following. If $Y \subset S$ is an essential subsurface (i.e., all boundary components are essential in S) satisfying $|\chi(S \setminus Y)| > D$ and α is a simple closed curve on S in minimal position with ∂Y and with $i(\alpha, \partial Y) \leq 4k$ for some $k \in \mathbb{N}$, then there exists an essential simple closed curve β on $S \setminus Y$ such that $i(\alpha, \beta) \leq k$.*

Lemma A.1 appears (in a different context and stated in a different way) as Proposition 3.1 of [2]. For completeness, we include a sketch of the proof here, which requires the following fact.

LEMMA A.2 (Lemma 3.2, [11]). *Let $\varepsilon > 0$. There exists a decreasing function $f : (0,1) \longrightarrow \mathbb{R}_+$ such that if $G = (V, E)$ is any graph with average degree greater than $2+\varepsilon$, then G has girth no larger than $g(\varepsilon) \cdot \log_2(|V|)$.*

REMARK A.3. The proof of Lemma 3.2 of [11] gives the bound of 18 for $f(1/2)$. This implies that any choice of $D > 296$ is sufficient for the statement of Theorem 1.1 to hold. However, we choose $D = 512$ to simplify our computation.

We are now ready to sketch the proof of Lemma A.1.

Proof of Lemma A.1. If there exists an essential simple curve in $S \setminus Y$ disjoint from α, then we are done. If no such curve exists, then $\alpha \cap (S \setminus Y)$ is a filling weighted multiarc on $S \setminus Y$, and therefore the complement of α in $S \setminus Y$ is a collection of polygons. Abusing notation slightly, we refer to this weighted multiarc $\alpha \cap (S \setminus Y)$ as α. Note that

$$|\chi(S \setminus Y)| \leq |\alpha| \leq 3|\chi(S \setminus Y)|. \tag{3}$$

Let α' be the collection of all arcs in α of weight $\geq \frac{2k}{\sqrt{|\chi(S \setminus Y)|}}$, which we will call *large mass arcs*. Since $i(\alpha, \partial Y) \leq 4k$, there are at most $\sqrt{|\chi(S \setminus Y)|}$ such arcs in α.

Let S' be the complement of α' in $S \setminus Y$. Cutting along any arc can decrease the absolute value of the Euler characteristic by at most 2, and therefore

$$|\chi(S')| \geq |\chi(S \setminus Y)| - 2\sqrt{|\chi(S \setminus Y)|}, \tag{4}$$

which is positive so long as $|\chi(S \setminus Y)| > 4$.

Let G denote the dual graph to α on $S \setminus Y$: one vertex for each complementary polygon, two of which are connected by an edge when the corresponding polygons share a boundary edge. The average degree $\bar{d}(G)$ of G is at least 3, since otherwise two arcs in α would be homotopic. Hence

$$\bar{d}(G) = \frac{2|E(G)|}{|V(G)|} \geq 3.$$

We now define G' to be the dual graph on S' to the collection of arcs in α that do not have large mass. Using (3) and (4), calculating the average degree of a vertex yields

$$\frac{2E(G')}{V(G')} = \frac{2(|\alpha| - |\alpha'|)}{V(G')} \geq 2.5,$$

where the last inequality holds so long as $|\chi(S \setminus Y)| > 36$.

By Lemma A.2 there exists a cycle β on G' of edge-length at most

$$f\left(\frac{1}{2}\right) \cdot \log_2(|V(G')|) \leq f\left(\frac{1}{2}\right) \cdot \log_2(2|\chi(S \setminus Y)|).$$

Without loss of generality, β is simple (otherwise, there exists a shorter cycle). To prove that β is essential, we show that inessential intersections between β and arcs of α imply the existence of inessential intersections between α and ∂Y (see [2] for details).

Therefore

$$i(\beta,\alpha) \leq \frac{f(\frac{1}{2}) \cdot \log_2(2|\chi(S \setminus Y)|) \cdot 2k}{\sqrt{|2\chi(S \setminus Y)|}}.$$

Hence it suffices to choose $D > 36$ sufficiently large so that

$$\frac{1}{2 \cdot f(\frac{1}{2})} > \frac{\log_2(D)}{D}. \qquad \square$$

We note that the conclusion of Lemma A.1 is not implied by the argument of Hempel [13]. Indeed, the more standard surgery argument will only reduce the intersection number by a factor of 2, as opposed to 4.

B. The Case $p = 6$ of Schaller's Conjecture

To resolve Schaller's conjecture when $(g, p) = (0, 6)$, we use an argumentation motivated by the tools in Section 5 of [26]. In particular, we will prove the following:

PROPOSITION B.1. *An automorphism of the* 0*-edge subgroup of* $\mathcal{SC}(S_{0,6})$ *can be uniquely extended to an automorphism of the full curve graph.*

Thus by Proposition B.1 we obtain an injection of the automorphism group of $\mathcal{SC}(S)$ into the extended mapping class group, as desired. To this end, recall that a *join* is a graph formed by two collections $\mathcal{V}$, $\mathcal{U}$ of vertices such that a pair of vertices $x, y \in \mathcal{V} \cup \mathcal{U}$ spans an edge if an only if either $x \in \mathcal{V}$, $y \in \mathcal{U}$ or $x \in \mathcal{U}$, $y \in \mathcal{V}$.

Using the language of [26], a join is 2*-sided* whenever $\mathcal{V}$ and $\mathcal{U}$ are both infinite sets, and a join is *maximal* when (i) neither $\mathcal{V}$ nor $\mathcal{U}$ can be replaced with proper supersets while maintaining the join property and (ii) there does not exist a vertex x simultaneously disjoint from every vertex in $\mathcal{V} \cup \mathcal{U}$. It is immediate that a graph automorphism must preserve the collection of maximal 2-sided joins.

Let $\mathcal{J} = \{\mathcal{V}, \mathcal{U}\}$ be a maximal 2-sided join in the 0-edge subgraph of $\mathcal{SC}(S_{0,6})$. Given a vertex v, let $b_1(v)$, $b_2(v)$ denote the two boundary components of $S_{0,6}$ cut off by v. It follows that if $v \in \mathcal{V}$ and $u \in \mathcal{U}$, then

$$\{b_1(v), b_2(v)\} \cap \{b_1(u), b_2(u)\} = \emptyset,$$

since otherwise the curves corresponding to u and v would intersect. Therefore $\mathcal{J}$ determines two disjoint subsets T_1, T_2 of the six boundary components of S such that each vertex $v \in \mathcal{V}$ (resp., $u \in \mathcal{U}$) has the property the $b_i(v) \in T_1$ (resp., $b_i(u) \in T_2$).

We claim that $|T_1| = |T_2| = 3$, for suppose that $|T_1| = 2$. Then maximality of $\mathcal{J}$ implies that there must be a boundary component b that is neither in T_1 nor in T_2. Indeed, consider a pair of separating curves that jointly encircle the same two boundary components but which are arranged on opposite sides of a third boundary component that they jointly block from the rest of the surface, as in Figure 22. If no such pair is present in $\mathcal{V}$, then either 2-sidedness or maximality is contradicted, and as soon as such a pair exists, the boundary component they jointly block cannot be in T_2. However, this also contradicts maximality, because we can then consider a curve that encircles the blocked boundary component and one of the two boundary components in T_1, and such a curve can be added to $\mathcal{V}$.

It follows that a maximal 2-sided join $\mathcal{J}$ corresponds to two infinite collections $\mathcal{V}, \mathcal{U}$ and a partitioning of the six boundary components of $S_{0,6}$ into two subsets T_1, T_2 of three boundary components such that each vertex in $\mathcal{V}$ encircles two boundary components of T_1 (and, respectively, for $u \in \mathcal{U}$ and T_2). It follows that exactly one component of $S \setminus \{\mathcal{V} \cup \mathcal{J}\}$ is an essential annulus separating T_1 and T_2, and thus there is a unique (up to isotopy) simple closed curve $\delta_{\mathcal{J}}$ determined by $\mathcal{J}$. See Figure 23.

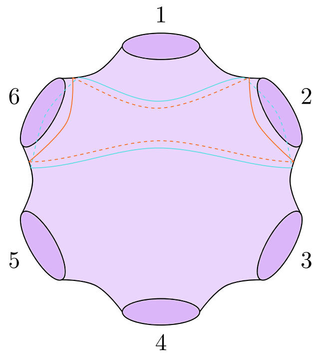

Figure 22 The orange and blue curves both encircle boundary components 2 and 6 while jointly blocking boundary component 1 from the rest of the surface.

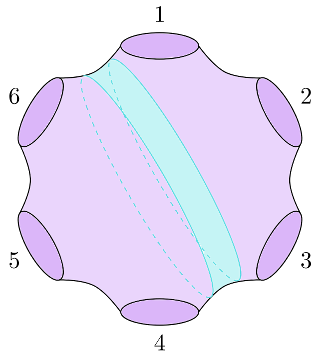

Figure 23 The simple closed curve $\delta_{\mathcal{J}}$ separating $T_1 = \{1, 2, 3\}$ and $T_2 = \{4, 5, 6\}$ is the core curve of the annulus shown in light blue.

Now let Φ be an automorphism of the 0-edge subgraph of $\mathcal{SC}(S_{0,6})$, and let α be a vertex of the curve graph that is not a vertex of $\mathcal{SC}(S_{0,6})$. It follows that α separates S into two subsurfaces each possessing three of the original boundary components. Then to α there corresponds a 2-sided maximal join $\mathcal{J}_\alpha = \{\mathcal{V}_\alpha, \mathcal{U}_\alpha\}$ of the 0-edge subgraph of $\mathcal{SC}(S)$: for instance, $\mathcal{V}_\alpha$ consists of all vertices of $\mathcal{SC}(S)$ disjoint from α and to the left of α. The automorphism Φ then sends $\mathcal{J}_\alpha$ to a maximal 2-sided join $\Phi(\mathcal{J})$, which, by the above paragraphs, uniquely determines a separating curve, which we define to be $\Phi(\alpha)$.

This extends Φ to the entire curve graph, and it remains only to check that Φ is an automorphism. If $v \in \mathcal{SC}(S)$, $\alpha \notin \mathcal{SC}(S)$, and $i(v, \alpha) = 0$, then $i(\Phi(v), \Phi(\alpha)) = 0$ since $\Phi(\alpha)$ will be the curve corresponding to a join for which v lives in one of the two vertex sets. If α, β are both not in $\mathcal{SC}(S)$, then they must intersect. This implies that the corresponding 2-sided maximal joins $\mathcal{J}_\alpha$, $\mathcal{J}_\beta$ are distinct and therefore that $\Phi(\mathcal{J}_\alpha) \neq \Phi(\mathcal{J}_\beta)$, which in turn implies that $\Phi(\alpha) \neq \Phi(\beta)$, and so these curves must also intersect.

Acknowledgments. We sincerely thank ICERM for its hospitality. Thanks to Moira Chas and Jonah Gaster for many helpful conversations and Nick Bell, Maxime Fortier Bourque, Dan Margalit, Alan McLeay, Athanase Papadopoulos, Joe Scull, and Davide Spriano for their comments on a draft of this paper. We would also like to thank the anonymous referee whose careful reading and thorough comments greatly improved the exposition of this paper.

References

[1] J. W. Anderson, H. Parlier, and A. Pettet, *Small filling sets of curves on a surface,* Topology Appl. 158 (2011), no. 1, 84–92.

[2] T. Aougab, *Uniform hyperbolicity of the graphs of curves,* Geom. Topol. 17 (2013), no. 5, 2855–2875.

[3] J. Aramayona, *Simplicial embeddings between pants graphs,* Geom. Dedicata 144 (2010), 115–128.

[4] J. Aramayona, T. Koberda, and H. Parlier, *Injective maps between flip graphs,* Ann. Inst. Fourier (Grenoble) 65 (2015), no. 5, 2037–2055.

[5] J. Aramayona and C. J. Leininger, *Finite rigid sets in curve complexes,* J. Topol. Anal. 5 (2013), no. 2, 183–203.

[6] T. E. Brendle and M. Dan, *Normal subgroups of mapping class groups and the metaconjecture of Ivanov,* J. Amer. Math. Soc. 32 (2019), no. 4, 1009–1070.

[7] M. Dan, *Problems, questions, and conjectures about mapping class groups,* Breadth in contemporary topology, Proc. Sympos. Pure Math., 102, pp. 157–186, Am. Math. Soc., Providence, 2019.

[8] V. Disarlo, *Combinatorial rigidity of arc complexes,* J. Amer. Math. Soc. 32 (2015), no. 4, 1009–1070.

[9] B. Farb ed., *Problems on mapping class groups and related topics,* Proc. Sympos. Pure Math., 74, Am. Math. Soc., Providence, 2006.

[10] B. Farb and M. Dan, *A primer on mapping class groups,* Princeton Math. Ser., 49, Princeton University Press, Princeton, 2012.

[11] S. Fiorini, G. Joret, D. O. Theis, and D. R. Wood, *Small minors in dense graphs,* European J. Combin. 33 (2012), no. 6, 1226–1245.

[12] A. Hatcher, *On triangulations of surfaces,* Topology Appl. 40 (1991), no. 2, 189–194.

[13] J. Hempel, *3-manifolds as viewed from the curve complex,* Topology 40 (2001), no. 3, 631–657.

[14] E. Irmak, *Superinjective simplicial maps of complexes of curves and injective homomorphisms of subgroups of mapping class groups,* Topology 43 (2004), no. 3, 513–541.

[15] ———, *Complexes of nonseparating curves and mapping class groups,* Michigan Math. J. 54 (2006), no. 1, 81–110.

[16] E. Irmak and M. Korkmaz, *Automorphisms of the Hatcher–Thurston complex,* Israel J. Math. 162 (2007), 183–196.

[17] E. Irmak and J. D. McCarthy, *Injective simplicial maps of the arc complex,* Turkish J. Math. 34 (2010), no. 3, 339–354.

[18] N. V. Ivanov, *Automorphism of complexes of curves and of Teichmüller spaces,* Int. Math. Res. Not. 14 (1997), 651–666.

[19] M. Korkmaz and A. Papadopoulos, *On the arc and curve complex of a surface,* Math. Proc. Cambridge Philos. Soc. 148 (2010), no. 3, 473–483.

[20] ———, *On the ideal triangulation graph of a punctured surface,* Ann. Inst. Fourier (Grenoble) 62 (2012), no. 4, 1367–1382.

[21] W. B. R. Lickorish, *A representation of orientable combinatorial* 3*-manifolds,* Ann. of Math. (2) 76 (1962), 531–540.

[22] F. Luo, *Automorphisms of the complex of curves,* Topology 39 (2000), no. 2, 283–298.

[23] H. A. Masur and Y. N. Minsky, *Geometry of the complex of curves. I. Hyperbolicity,* Invent. Math. 138 (1999), no. 1, 103–149.

[24] ———, *Geometry of the complex of curves. II. Hierarchical structure,* Geom. Funct. Anal. 10 (2000), no. 4, 902–974.

[25] J. D. McCarthy and A. Papadopoulos, *Simplicial actions of mapping class groups,* Handbook of Teichmüller theory. Volume III, IRMA Lect. Math. Theor. Phys., 17, pp. 297–423, Eur. Math. Soc., Zürich, 2012.

[26] A. McLeay, *Geometric normal subgroups in mapping class groups of punctured surfaces,* New York J. Math. 25 (2019), 839–888.

[27] K. Rafi and S. Schleimer, *Curve complexes are rigid,* Duke Math. J. 158 (2011), no. 2, 225–246.

[28] S. Schleimer, *Notes on the complex of curves,* ⟨http://homepages.warwick.ac.uk/~masgar/Maths/notes.pdf⟩.

[29] P. Schmutz Schaller, *Mapping class groups of hyperbolic surfaces and automorphism groups of graphs,* Compos. Math. 122 (2000), no. 3, 243–260.

S. Agrawal
731 Lexington Ave
New York, NY 10022
USA

shuchi4112@gmail.com

T. Aougab
Department of Mathematics and Statistics
Haverford College
Lancaster, PA 19041
USA

taougab@haverford.edu

Y. Chandran
Department of Mathematics
City University of New York
New York, NY 10016
USA

ychandran@gradcenter.cuny.edu

M. Loving
Department of Mathematics
Georgia Institute of Technology
Atlanta, GA 30332
USA

mloving6@gatech.edu

J. R. Oakley
Department of Mathematics
Temple University
Philadelphia, PA 19122
USA

oakley@temple.edu

R. Shapiro
Department of Mathematics
Georgia Institute of Technology
Atlanta, GA 30332
USA

shapiro32@gatech.edu

Y. Xiao
Data Scientist
Robinhood
Mountain View, CA 94039
USA

yang_xiao@alumni.brown.edu