

AG
T

*Algebraic & Geometric
Topology*

Volume 25 (2025)

Infinite-type loxodromic isometries of the relative arc graph

CAROLYN ABBOTT

NICHOLAS MILLER

PRIYAM PATEL



Infinite-type loxodromic isometries of the relative arc graph

CAROLYN ABBOTT
NICHOLAS MILLER
PRIYAM PATEL

An infinite-type surface Σ is admissible if it has an isolated puncture p and admits shift maps. This includes all infinite-type surfaces with an isolated puncture outside of two sporadic classes. Given such a surface, we construct an infinite family of intrinsically infinite-type mapping classes that act loxodromically on the relative arc graph $\mathcal{A}(\Sigma, p)$. J Bavard produced such an element for the plane minus a Cantor set, and our result gives the first examples of such mapping classes for all other admissible surfaces. The elements we construct are the composition of three shift maps on Σ , and we give an alternative characterization of these elements as a composition of a pseudo-Anosov on a finite-type subsurface of Σ and a standard shift map. We then explicitly find their limit points on the boundary of $\mathcal{A}(\Sigma, p)$ and their limiting geodesic laminations. Finally, we show that these infinite-type elements can be used to prove that $\text{Map}(\Sigma, p)$ has an infinite-dimensional space of quasimorphisms.

20F65, 57K20

1. Introduction	563
2. Background	569
3. Coding arcs and standard position	576
4. Characterizing loops with trivial image	593
5. Constructing the homeomorphisms	602
6. Arcs that start like $\alpha_i^{(n)}$	606
7. Highways in arcs	613
8. Loxodromic elements for $\mathcal{A}(\Sigma, p)$	621
9. Infinite-type quasimorphisms	631
10. Convergence to a geodesic lamination	637
References	643

1 Introduction

A surface Σ is of finite-type if $\pi_1(\Sigma)$ is finitely generated, and otherwise Σ is of infinite-type. Recently, there has been a surge of interest in infinite-type surfaces and their mapping class groups $\text{Map}(\Sigma)$, which arise naturally in a variety of contexts in low-dimensional topology, dynamics, and even descriptive set

theory. See the overview by Aramayona and Vlamis [4] for a survey of recent results on infinite-type mapping class groups.

For finite-type surfaces Σ , Nielsen [21] and Thurston [27] give a powerful classification of the elements of $\text{Map}(\Sigma)$: every element is periodic, reducible, or pseudo-Anosov. The action of $\text{Map}(\Sigma)$ by isometries on the (infinite-diameter and hyperbolic) curve graph $\mathcal{C}(\Sigma)$ captures a coarser classification of the elements of $\text{Map}(\Sigma)$ since elements are either *elliptic* or *loxodromic*. These two classifications, which are both interesting in their own right, have a strong relationship; the loxodromic elements are exactly the pseudo-Anosovs. In this way, the most interesting and complex mapping classes correspond to the dynamically richest actions.

The situation for infinite-type surfaces is more complicated for a few reasons. First, the exact analog of the Nielsen–Thurston classification is no longer valid in this setting since some elements are neither periodic, reducible, nor pseudo-Anosov in the traditional sense. Second, the curve graph of an infinite-type surface has finite diameter unlike for finite-type surfaces. This paper is motivated by one of the biggest open problems for infinite-type surfaces, which is to give an analog of the Nielsen–Thurston classification for infinite-type mapping classes. We work towards this goal by studying the action of $\text{Map}(\Sigma)$ on a different hyperbolic graph.

When Σ is an infinite-type surface with at least one isolated puncture p , the *relative arc graph*, $\mathcal{A}(\Sigma, p)$, plays the role of $\mathcal{C}(\Sigma)$ and is defined as follows: the vertices correspond to isotopy classes of simple arcs that begin and end at p and edges connect vertices for arcs admitting disjoint representatives. The subgroup $\text{Map}(\Sigma, p)$ of $\text{Map}(\Sigma)$ that fixes the isolated puncture p acts on $\mathcal{A}(\Sigma, p)$ by isometries. This graph was first defined by D Calegari [12], who initiated its study by asking whether, for the plane minus a Cantor set, this graph was infinite diameter and whether any element of $\text{Map}(\Sigma, p)$ acted loxodromically. In [5], J Bavard carried out Calegari’s program for the plane minus a Cantor set and, for that surface, showed that $\mathcal{A}(\Sigma, p)$ is both infinite-diameter and hyperbolic. Aramayona, Fossas, and Parlier [2] then showed that these properties for $\mathcal{A}(\Sigma, p)$ hold more generally for any infinite-type surface with at least one isolated puncture.

Given that the trichotomy of the Nielsen–Thurston classification does not exactly hold for infinite-type surfaces, it is necessary to redefine reducible, and therefore irreducible, mapping classes in this setting. One of the most promising ways to motivate a new definition is to classify the elements of infinite-type mapping class groups that are loxodromic with respect to the action of $\text{Map}(\Sigma)$ on a hyperbolic graph since these elements correspond to infinite-order irreducibles in the finite-type setting. In order to classify these elements, we must first construct them. We restrict to surfaces with an isolated puncture and their associated graphs $\mathcal{A}(\Sigma, p)$ in this paper because the relative arc graph is one of the few known graphs associated to an infinite-type surface that is both infinite-diameter and hyperbolic. In general, proving results for infinite-type surfaces typically involves a piecemeal approach. Constructing loxodromic isometries for infinite-type surfaces without a puncture will require a different graph, and thus, a genuinely different approach.




Figure 1: Above: a handleshift on an infinite-type surface. Below: a shift map on an infinite-type surface.

When Σ is the sphere minus a Cantor set with an isolated puncture p (ie Σ is the plane minus a Cantor set), Bavard [5] constructed an *intrinsically infinite-type* mapping class that is loxodromic with respect to the action of $\text{Map}(\Sigma)$ on $\mathcal{A}(\Sigma, p)$, and for several years, this was the only known such example. In this paper, we give a new construction of mapping classes that are loxodromic with respect to the action of $\text{Map}(\Sigma, p)$ on the relative arc graph $\mathcal{A}(\Sigma, p)$ for all infinite-type surfaces with an isolated puncture (outside of two sporadic classes). This class of surfaces is uncountable.

Theorem 1.1 *For any admissible surface Σ , there is an infinite family of intrinsically infinite-type homeomorphisms $\{g_n\}_{n \in \mathbb{N}}$ in $\text{Map}(\Sigma, p)$ such that each g_n is loxodromic with respect to the action of $\text{Map}(\Sigma, p)$ on $\mathcal{A}(\Sigma, p)$.*

In addition, we explore other dynamical and geometric properties of these mapping classes by demonstrating the convergence of a simple closed curve to a geodesic lamination on Σ under iterates of the maps and constructing an infinite-dimensional space of quasimorphisms of $\text{Map}(\Sigma, p)$ using these elements (see below).

Each mapping class in our construction is the composition of three homeomorphisms called *shift maps*. Shift maps are generalizations of the *handleshift* homeomorphisms constructed by the third author and N Vlamis in [22] (see Figure 1 for examples of both). Roughly, an infinite-type surface Σ with an isolated puncture p is admissible if there is a proper embedding of the biinfinite flute surface containing p into Σ such that certain shift maps on the flute surface induce shift maps on Σ . Such an embedding allows us to reduce the proof of Theorem 1.1 to the case of the biinfinite flute surface. See Section 2.4 for more




Figure 2: Examples of sporadic infinite-type surfaces that are not admissible. The first is a flute with finite genus, the other three are fluted Loch Ness monster surfaces.

details and Figure 3 for some examples of admissible surfaces. In Lemma 2.7, we show that a surface with an isolated puncture is admissible if and only if it admits shift maps. This set of surfaces consists of all infinite-type surfaces with an isolated puncture except a flute surface with finite (possibly zero) genus and a *fluted Loch Ness monster*. We call these two classes *sporadic surfaces* in this context. See Figure 2 for examples of sporadic surfaces and Lemma 2.9 for a proof of this fact. Since sporadic surfaces are exactly the small class of surfaces with an isolated puncture that do not admit shift maps, different methods will need to be developed in order to prove an analog of Theorem 1.1 for these surfaces. It would be interesting to understand how elements of the mapping class groups of sporadic surfaces act on the relative arc graph.

The handleshift homeomorphisms mentioned above have proven to be crucial in understanding various aspects of infinite-type mapping class groups. For example, it is shown in [22] that they are needed to topologically generate the *pure mapping class group* whenever Σ has at least two nonplanar ends, and in [3] Aramayona, Patel, and Vlamis showed that they are used to show that the pure mapping class groups of such surfaces surject onto \mathbb{Z} . With this paper, we emphasize the importance of more general shift maps to the theory of infinite-type mapping class groups. Inspired by Bavard's work in [5], we choose the shift maps in our construction carefully so that their composition mimics some of the behavior of pseudo-Anosov maps in the finite-type setting. In fact, we show that there is an alternative description of our homeomorphisms as the composition of a pseudo-Anosov homeomorphism on a finite-type subsurface and a *standard shift map* on Σ in Theorem 8.5. Additionally, in Section 10, we use the work of D Šarić [26] to prove the following theorem regarding geodesic laminations for the mapping classes constructed in Theorem 1.1.

Theorem 1.2 *If Σ is an admissible surface equipped with its conformal hyperbolic metric that is equal to its convex core, then there exists a simple closed curve c_0 on Σ such that the sequence $(g_n^i(c_0))_{i \in \mathbb{N}}$ converges to a geodesic lamination on Σ .*

In particular, we produce a train track on Σ and show the geodesic lamination from this theorem is weakly carried by this train track.

We emphasize that the elements arising from our construction are of *intrinsically infinite-type*, that is, they do not lie in the closure of the compactly supported mapping class group $\overline{\text{Map}_c(\Sigma)}$, where the closure is taken with respect to the compact-open topology on $\text{Map}(\Sigma)$. These are the first examples of intrinsically




Figure 3: Examples of admissible surfaces, the first of which is the biinfinite flute surface S itself.

infinite-type loxodromic isometries for all admissible surfaces outside of the plane minus a Cantor set. Additionally, we emphasize that our construction does not rely on, and is not a generalization of, one of the few known methods for constructing pseudo-Anosov mapping classes for finite-type surfaces.

The most obvious candidates for mapping classes that are loxodromic with respect to the action on $\mathcal{A}(\Sigma, p)$ are those that are pseudo-Anosov on a finite-type subsurface $\Sigma' \subset \Sigma$ containing the special puncture p , that extend via the identity map to the rest of Σ (these are compactly supported mapping classes). In [7], Bavard and Walker prove that these types of mapping classes do indeed act loxodromically on a graph that is quasi-isometric to $\mathcal{A}(\Sigma, p)$. In that paper they point out that, though their class of examples is interesting, it will be even more important to construct mapping classes of intrinsically infinite-type that act loxodromically on $\mathcal{A}(\Sigma, p)$; this remark was one of the main points of inspiration for writing this paper. The intrinsically infinite-type elements of $\text{Map}(\Sigma)$ are more mysterious since tools from finite-type surface theory do not directly generalize when studying these elements.

Remark 1.3 Morales and Valdez [20], building off of the work of Hooper [16], have also produced noncompactly supported elements that are loxodromic, but their elements are in the closure of the compactly supported mapping class group. Their method is a generalization of the Thurston–Veech construction of pseudo-Anosovs in the finite-type setting.

Aside from the motivation provided by a Nielsen–Thurston classification for infinite-type mapping classes, Bestvina and Fujiwara [10] show that constructing elements of $\text{Map}(\Sigma)$ that act loxodromically on hyperbolic graphs can be used to understand the second bounded cohomology $H_b^2(\text{Map}(\Sigma), \mathbb{R})$ of $\text{Map}(\Sigma)$. In particular, they show that, for a compact surface Σ , there exist elements acting loxodromically on $\mathcal{C}(\Sigma)$ that are *weakly properly discontinuous* (WPD). These elements are used to prove that the space of quasimorphisms of $\text{Map}(\Sigma)$ is infinite-dimensional, which is sufficient to conclude that $H_b^2(\text{Map}(\Sigma), \mathbb{R})$ is, as well.

Along these lines, M Bestvina asked the following question at the AIM workshop on infinite-type surfaces [1, Problem 4.7]: “For $\Sigma = \mathbb{R}^2 - C$ (where C is a Cantor set) is it true that every subgroup of $\text{Map}(\Sigma)$ has either infinite-dimensional space of quasimorphisms or is amenable?”¹ More generally, we would like to characterize the infinite-type mapping classes that can be used to produce quasimorphisms of $\text{Map}(\Sigma)$. In Section 9, we show that the elements constructed in Theorem 1.1 can be used to give a new proof of the following theorem, originally due to Bavard [5] in the case of a plane minus a Cantor set and Bavard and Walker [7] in the general case.

Theorem 1.4 *Let Σ be an admissible surface. The space of nontrivial quasimorphisms on $\text{Map}(\Sigma, p)$ is infinite-dimensional.*

When Σ is infinite type, $\text{Map}(\Sigma)$ does not contain WPD elements; see the demonstration by Bavard and Genevois [6]. Despite this, we are still able to build nontrivial quasimorphisms using a weaker condition on loxodromic elements introduced by Bestvina and Fujiwara [10], called being antialigned, and an approach similar to that of Bavard in [5] which involves defining an intersection pairings on a specific class of arcs on Σ . In [8], Bavard and Walker use the same condition of being antialigned to show that homeomorphisms that are pseudo-Anosov on finite-type subsurfaces Σ' and extend via the identity to the rest of Σ can be used to produce quasimorphisms of $\text{Map}(\Sigma)$. There is also a weaker version of the WPD condition, called WWPD, which was introduced by Bestvina, Bromberg, and Fujiwara in [9]; WWPD elements are always antialigned. A Rasmussen shows in [24] that for a surface Σ with an isolated puncture p , an element of $\text{Map}(\Sigma, p)$ is WWPD with respect to the action on $\mathcal{A}(\Sigma, p)$ if and only if it stabilizes a finite-type subsurface Σ' containing the puncture p and restricts to a pseudo-Anosov on Σ' . In particular, the elements used by Bavard and Walker to construct nontrivial quasimorphisms are WWPD elements. The elements we construct in Theorem 1.1 do not fix any finite-type subsurface and thus are *not* WWPD. Our construction gives subgroups of $\text{Map}(\Sigma)$ that do not contain WWPD elements but do have an infinite-dimensional space of quasimorphisms.

Plan of the paper In order to prove Theorem 1.1, we explicitly compute the images of a particular arc on Σ under iterates of each homeomorphism g_n and prove that these images form a quasigeodesic axis for the action of $\langle g_n \rangle$ on $\mathcal{A}(\Sigma, p)$. Though some of the methods in our paper are inspired by Bavard’s work in [5], we note that there are a variety of additional challenges in proving Theorem 1.1 for such a

¹This question has since been answered by Fournier-Facio, Lodha, and Zaremsky [14].

wide class of surfaces. In fact, we first prove the theorem for the biinfinite flute surface S and then use the fact that the inclusion of $\mathcal{A}(S, p)$ into $\mathcal{A}(\Sigma, p)$ is a $(2, 0)$ -quasi-isometric embedding (see Lemma 2.10) to extend the theorem to all admissible surfaces. One of the first challenges in proving Theorem 1.1 is rigorously coding arcs on S , which we do in Section 3.1, in order to quantify how long two arcs on S fellow travel. We note that the coding of our arcs is done in way that is “geometrically intuitive”. What we mean is that, given the code for an arc, one can easily draw the corresponding arc on the surface. For other ways to code arcs see [5; 7]. We then introduce *standard position* for an arc on S in Section 3.2 so that we can use the code for an arc to find its image under our shift maps in a well-defined way. Most importantly, we must understand when segments of arcs become trivial under our shift maps, and in Section 4.1 we introduce a kind of cancellation in the image of the code for a segment which we call *cascading cancellation*. This kind of cancellation will cause technical problems throughout the paper and much of Section 6 is devoted to understanding how to control it.

The rest of Section 4 is devoted to proving Theorem 4.8, which answers the question of when a segment in an arc becomes trivial under our shift maps. We define the homeomorphisms g_n of Theorem 1.1 in Section 5, show that we have “*starts like*” functions in Section 6, and show that we have *highways* in Section 7. Finally, we prove Theorem 1.1 in Section 8, introduce an intersection pairing for arcs and prove Theorem 1.4 in Section 9, and prove the convergence to a geodesic lamination from Theorem 1.2 in Section 10.

Acknowledgements The authors are grateful to Chris Leininger for pointing out the work of Verberne and to Ekta Patel for helping us with computational work related to this project. The authors would also like to thank Alden Walker and Nick Vlamis for helpful discussions as well as Elizabeth Field for helpful comments on an earlier draft. Abbott was partially supported by NSF DMS–1803368 and DMS–2106906, Miller was partially supported by NSF DMS–2005438 and DMS–2300370, and Patel was partially supported by NSF DMS–1840190 and DMS–2046889. The authors would also like to thank the American Institute of Mathematics and the organizers of the *Infinite-type surfaces* workshop that was held there, during which some of this work was completed.

2 Background

2.1 Space of ends and classification of infinite-type surfaces

Central to the classification of infinite-type surfaces is the definition of the space of ends $E(\Sigma)$ of an infinite-type surface Σ . Informally, an end of Σ is a way to escape or go off to infinity in Σ . More formally we have:

Definition 2.1 An *exiting sequence* in Σ is a sequence $\{U_n\}_{n \in \mathbb{N}}$ of connected open subsets of Σ satisfying

- (1) $U_n \subset U_m$ whenever $m < n$;

- (2) U_n is not relatively compact for any $n \in \mathbb{N}$, that is, the closure of U_n in Σ is not compact;
- (3) the boundary of U_n is compact for each $n \in \mathbb{N}$; and
- (4) any relatively compact subset of Σ is disjoint from all but finitely many U_n .

Two exiting sequences $\{U_n\}_{n \in \mathbb{N}}$ and $\{V_n\}_{n \in \mathbb{N}}$ are equivalent if for every $n \in \mathbb{N}$ there exists $m \in \mathbb{N}$ such that $U_m \subset V_n$ and $V_m \subset U_n$. An *end* of Σ is an equivalence class of exiting sequences.

The *space of ends* $E(\Sigma)$, or simply E , of Σ is the set of ends of Σ equipped with a natural topology for which it is totally disconnected, Hausdorff, second countable, and compact. In particular, $E(\Sigma)$ is homeomorphic to a closed subset of the Cantor set. To describe the topology, let V be an open subset of Σ with compact boundary, define $\hat{V} = \{[U_n]_{n \in \mathbb{N}} \in E : U_n \subset V \text{ for some } n \in \mathbb{N}\}$ and let $\mathcal{V} = \{\hat{V} : V \subset \Sigma \text{ is open with compact boundary}\}$. The set E becomes a topological space by declaring \mathcal{V} a basis for the topology.

We note that ends can be isolated or not and can be *planar* (if there exists an i such that U_i is homeomorphic to an open subset of the plane \mathbb{R}^2) or *nonplanar* (if every U_i has infinite genus). The set of nonplanar ends of Σ is a closed subspace of $E(\Sigma)$ and will be denoted by $E_g(\Sigma)$.

Kerékjártó [17] and Richards [25] showed that the homeomorphism type of an orientable infinite-type surface is determined by the quadruple

$$(g, b, E_g(\Sigma), E(\Sigma)),$$

where $g \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$ is the genus of Σ and $b \in \mathbb{Z}_{\geq 0}$ is the number of (compact) boundary components of Σ .

Of particular interest to us is the infinite-type surface called the *biinfinite flute* obtained from an infinite cylinder by deleting a countable discrete sequence of points exiting both ends of the cylinder (see Figure 3). By the classification theorem of Kerékjártó and Richards, this surface can also be obtained from \mathbb{S}^2 by deleting $\{x_i\}$, $\{y_i\}$, x , and y , where $\{x_i\}$ and $\{y_i\}$ are countable discrete sequences of points converging to distinct points x and y , respectively. Note that S has two special nonisolated ends.

2.2 Mapping class groups and arc graphs

The *mapping class group*, $\text{Map}(\Sigma)$, of a surface Σ is the group of orientation-preserving homeomorphisms of Σ up to isotopy. The natural topology on any group of homeomorphisms is the compact-open topology and $\text{Map}(\Sigma)$ is endowed with the quotient topology with respect to the compact-open topology on the space of homeomorphisms of Σ . When Σ is a finite-type surface, this topology agrees with the discrete topology on $\text{Map}(\Sigma)$, but when Σ is of infinite type it does not. There are several important subgroups of $\text{Map}(\Sigma)$: $\text{Map}_c(\Sigma)$ is the subgroup consisting of mapping classes with compact support, $\text{PMap}(\Sigma)$ is the pure mapping class group consisting of mapping classes which fix the set of ends pointwise,

$\overline{\text{Map}_c(\Sigma)} < \text{PMap}(\Sigma)$ is the closure of the compactly supported mapping class group with respect to the topology described above, and when Σ has an isolated puncture p , $\text{Map}(\Sigma, p)$ is the subgroup of mapping classes that fix p .

When Σ is finite-type, $\text{Map}(\Sigma)$ is algebraically generated by finitely many Dehn twists [19]. Infinite-type mapping class groups, sometimes called big mapping class groups, are uncountable groups, so there is no countable algebraic generating set. However, one can consider topological generating sets (countable dense subsets of $\text{Map}(\Sigma)$) and in [22], Vlamis and the third author prove that for many infinite-type surfaces, Dehn twists are not sufficient in topologically generating even $\text{PMap}(\Sigma)$. They show that in addition to Dehn twists, a new class of homeomorphisms called *handleshifts* (defined in Section 2.4) are often needed to topologically generate $\text{PMap}(\Sigma)$. In a subsequent paper with Aramayona [3], Vlamis and the third author give an algebraic description of $\text{PMap}(\Sigma)$ that will be relevant in Section 8. When Σ is an infinite-type surface with $n > 1$ nonplanar ends, they prove that $\text{PMap}(\Sigma) = \overline{\text{Map}_c(\Sigma)} \rtimes \mathbb{Z}^{n-1}$, where \mathbb{Z}^{n-1} is generated by $n - 1$ handleshifts with disjoint support. In particular, when Σ has exactly 2 nonplanar ends (for example when Σ is the ladder surface), $\text{PMap}(\Sigma) = \overline{\text{Map}_c(\Sigma)} \rtimes \mathbb{Z}$, where $\mathbb{Z} = \langle H \rangle$ and H is the standard handleshift, shifting each genus of Σ over to the right by one.

In this paper we are primarily concerned with mapping classes of intrinsically infinite-type.

Definition 2.2 An element $f \in \text{Map}(\Sigma)$ is of *intrinsically infinite-type* if $f \notin \overline{\text{Map}_c(\Sigma)}$.

More specifically, we are interested in how such elements act on a particular graph of arcs called the *relative arc graph*.

Let Σ be a connected, orientable surface with empty boundary, and let $\Pi \subset \Sigma$ be the set of punctures of Σ , which we assume to be nonempty. In this subsection, it is convenient to regard Π as a set of marked points on Σ . By a *proper arc* on Σ we mean a map $\alpha: [0, 1] \rightarrow \Sigma$ such that $\alpha^{-1}(\Pi) = \{0, 1\}$. We often conflate an arc with its image in Σ . An arc is *simple* if it is an embedding when restricted to the open interval $(0, 1)$.

The *arc graph* $\mathcal{A}(\Sigma)$ is the simplicial graph whose vertices are isotopy classes of simple arcs on Σ , where we only consider isotopies rel endpoints, and two (isotopy classes of) arcs are connected by an edge if they can be realized disjointly away from Π . The mapping class group $\text{Map}(\Sigma)$ acts on $\mathcal{A}(\Sigma)$ by isometries. Hensel, Przytycki, and Webb [15] show that when Σ has finite-type, the graph $\mathcal{A}(\Sigma)$ is infinite diameter and 7-hyperbolic. On the other hand, when Σ is infinite-type with infinitely many punctures, it is straightforward to see that $\mathcal{A}(\Sigma)$ has diameter 2, and so this graph is not particularly useful for studying $\text{Map}(\Sigma)$.

Assuming that Π contains a nonempty set of isolated punctures, Aramayona, Fossas, and Parlier [2] construct a particular subgraph of the arc graph which has interesting geometry, even when Π is infinite. We are interested in a special case of this construction, involving a single isolated puncture p .

Definition 2.3 The *relative arc graph* $\mathcal{A}(\Sigma, p)$ is the subgraph of $\mathcal{A}(\Sigma)$ spanned by arcs which start and end at p . More precisely, the vertices of $\mathcal{A}(\Sigma, p)$ are isotopy classes of arcs on Σ with endpoints on p ,

where we allow only isotopy rel endpoints. There is an edge between two (isotopy classes of) arcs if they can be realized disjointly away from p .

Aramayona, Fossas, and Parlier show that $\mathcal{A}(\Sigma, p)$ is connected, has infinite diameter, and is 7-hyperbolic (see [2, Theorem 1.1]). While $\text{Map}(\Sigma)$ does not necessarily act on $\mathcal{A}(\Sigma, p)$, the subgroup $\text{Map}(\Sigma, p)$ that fixes the puncture p does act by isometries on this graph. When Σ has only one isolated puncture p , $\text{Map}(\Sigma) = \text{Map}(\Sigma, p)$.

2.3 Metric spaces and loxodromic isometries

We now introduce some basics of metric spaces and isometries of a hyperbolic metric space. Given a metric space X , we denote by d_X the distance function on X . A map $f: X \rightarrow Y$ between metric spaces X and Y is a (K, C) -quasi-isometric embedding if there are constants $K \geq 1$, $C \geq 0$ such that for all $x, y \in X$,

$$\frac{1}{K}d_X(x, y) - C \leq d_Y(f(x), f(y)) \leq Kd_X(x, y) + C.$$

A *geodesic* in X is an isometric embedding of an interval into X and a (K, C) -*quasigeodesic* in X is a (K, C) -quasi-isometric embedding of an interval into X . We call the constants K, C the *quality* of the quasigeodesic. By an abuse of notation, we often conflate a (quasi)geodesic and its image in X .

Definition 2.4 Given an action by isometries of a group G on a hyperbolic space X , an element $g \in G$ is *elliptic* if it has bounded orbits; *loxodromic* if the map $\mathbb{Z} \rightarrow X$ given by $n \mapsto g^n x_0$ for some (equivalently, any) $x_0 \in X$ is a quasi-isometric embedding; and *parabolic* otherwise.

Any biinfinite quasigeodesic in X which is preserved by a loxodromic isometry $g \in G$ is called an *axis* of g . An axis always exists; for any $x_0 \in X$, the set $\{g^n x_0 \mid n \in \mathbb{Z}\}$ is a (discrete) quasigeodesic preserved by g . If X is a geodesic metric space, in the sense that there exists a geodesic connecting any two points of X , then we may construct a continuous quasigeodesic axis as follows. Fix a geodesic $[x_0, gx_0]$ from x_0 to gx_0 . Then g stabilizes the path formed by concatenating the geodesics $g^n[x_0, gx_0]$; this path is a quasigeodesic axis of g in X . Varying the point x_0 will change the quality of the quasigeodesic. Let $g^+ = \lim_{n \rightarrow \infty} g^n x_0$ and $g^- = \lim_{n \rightarrow -\infty} g^n x_0$ be points in the Gromov boundary ∂X of X . The *limit set* of $\langle g \rangle$ is the subset $\{g^+, g^-\} \subseteq \partial X$; this set is fixed pointwise by g . It is straightforward to show that the limit set $\{g^+, g^-\}$ does not depend on the choice of $x_0 \in X$.

2.4 Shift maps and the biinfinite flute surface

A handleshift was first defined in [22] as follows. Consider the surface S' defined by taking the strip $\mathbb{R} \times [-1, 1]$, removing a disk of radius $\frac{1}{2}$ with center $(n, 0)$ for each $n \in \mathbb{Z}$, and attaching a torus with one boundary component to the boundary of each such disk. A handleshift on S' is the homeomorphism that acts like a translation, sending (x, y) in S to $(x + 1, y)$ and which tapers to the identity on $\partial S'$. Given a surface of infinite-genus Σ with at least two nonplanar ends and a proper embedding of S' into

Σ so that the two ends of the strip correspond to two distinct ends of Σ , the handleshift on S' induces a handleshift on Σ , where the homeomorphism acts as the identity on the complement of S' . In this paper, more flexibility is allowed, and we define the following generalization.

Definition 2.5 Let S' be the surface defined by taking the strip $\mathbb{R} \times [-1, 1]$, removing a closed disk of radius $\frac{1}{4}$ with center $(n, 0)$ for $n \in \mathbb{Z}$, and attaching any fixed topologically nontrivial surface with exactly one boundary component to the boundary of each such disk. A *shift* on S' is the homeomorphism that acts like a translation, sending (x, y) in S' to $(x + 1, y)$ and which tapers to the identity on $\partial S'$.

Lanier and Loving use two particular cases of this generalization in [18]. Naming the full generalization a “shift” is in line with their paper. Note that it is essential for the same surface to be glued to the boundary component of each disk in order for the shift to be a homeomorphism of the surface.

As above, given a surface Σ with a proper embedding of S' into Σ so that the two ends of the strip correspond to two different ends of Σ , the shift on S' induces a shift on Σ , where the homeomorphism acts as the identity on the complement of S' . Given a shift h on Σ , the embedded copy of S' in Σ is called the *domain* of h . In this paper, we produce special homeomorphisms that can be obtained as a composition of three shift maps on such a surface Σ with an isolated puncture p and that are loxodromic with respect to the action of $\text{Map}(\Sigma, p)$ on $\mathcal{A}(\Sigma, p)$. Instead of working generally with surfaces that admit shift maps, we begin by letting S be the biinfinite flute surface. Then, S admits shift maps which shift a countable collection of punctures on S . To prove Theorem 1.1 we first construct mapping classes that are loxodromic with respect to the action of $\text{Map}(S, p)$ on $\mathcal{A}(S, p)$. We then use this surface as a template for constructing the desired mapping classes for more general surfaces Σ by extending the shift maps on S to shift maps on Σ as follows.

Definition 2.6 Let S be the biinfinite flute surface. A surface Σ with an isolated puncture p is *admissible* if there exists a proper embedding $S \hookrightarrow \Sigma$ where S contains p , the two nonisolated ends of S correspond to distinct ends of Σ , and such that a countably infinite collection of connected components of $\Sigma \setminus S$ are of the same (nontrivial) topological type. Note that when the components are once-punctured disks, there are countably many isolated punctures of S that remain isolated punctures when embedded in Σ . Denote this special class of connected components of $\Sigma \setminus S$ by \mathcal{U} , so that the elements of \mathcal{U} are all homeomorphic to a fixed surface Σ_0 with one boundary component. See Figure 1 for some examples of admissible surfaces.

Given a shift map h on S , the support of h is a strip $\mathbb{R} \times [-1, 1]$ with countably many punctures. When the set of punctures in the support of h only consists of those corresponding to elements of \mathcal{U} , we can glue copies of Σ_0 onto the punctures of this strip to produce a shift map on a surface S' as in Definition 2.5. The embedding of S in Σ therefore gives an embedding of S' in Σ and the shift on S' in Σ is extended via the identity on $\Sigma \setminus S'$ as usual. From this construction, we immediately have one direction of the following lemma.




Figure 4: Examples showing that surfaces with shift maps are always admissible. For the first surface, only one curve γ is needed to cut away the extra topology of Σ . In the second case, a countable collection of curves $\{\gamma_i\}$ is needed.

Lemma 2.7 *Given a surface Σ with an isolated puncture, Σ is admissible if and only if Σ admits shift maps.*

Proof It is left to show that if Σ has an isolated puncture p and admits a shift map, then Σ is admissible. To see this, we consider the proper embedding of S' into Σ . Recall that S' is obtained from a punctured strip by gluing on countably many copies of any surface Σ_0 with exactly one boundary component. Let $\mathcal{T} = \{T_i\}$ denote the corresponding countable collection of subsurfaces homeomorphic to Σ_0 in Σ , indexed by \mathbb{Z} .

Note that $E(S')$ is a closed subset of $E(\Sigma)$, as is $X = E(S') \cup \{p\}$, and thus $X^c = E(\Sigma) \setminus X$ is open in $E(\Sigma)$. The second countability of the topology on $E(\Sigma)$ implies that X^c is the union of countably many basis elements. If X^c is in fact clopen in $E(\Sigma)$, then X^c is compact and is therefore a finite union of basis elements. In this case, there exists a simple closed curve γ in Σ with the following property: there exists a connected component K of $\Sigma \setminus \gamma$ such that the end space of $\bar{K} = K \cup \gamma$ is exactly X^c . In this way, γ cuts away the ends of Σ that are in X^c (see Figure 4). We then have that

$$\Sigma \setminus (\bar{K} \cup (\bigcup_i T_i))$$

is homeomorphic to the biinfinite flute surface and Σ is admissible with \mathcal{T} playing the role of \mathcal{U} in the definition of an admissible surface.

In general, we can only assume that X^c is open, not clopen, so that it can be expressed as the union of countably many basis elements for the topology. Then, there exists a countable collection of simple closed curves $\{\gamma_i\}$ and a countable collection of connected components K_i of $\Sigma \setminus \gamma_i$ such that the end space of $\bigcup_i \bar{K}_i = \bigcup_i K_i \cup \gamma_i$ is exactly X^c (see Figure 4). In this case,

$$\Sigma \setminus ((\bigcup_i \bar{K}_i) \cup (\bigcup_i T_i))$$

is homeomorphic to the biinfinite flute surface and Σ is admissible. \square

Given this equivalent definition for an admissible surface, we can show that this class includes all infinite-type surfaces with an isolated puncture outside of two sporadic classes that do not admit shift maps. We will need the following definition.

Definition 2.8 The *Loch Ness monster* is the infinite-type surface with no planar ends and exactly one nonplanar end. An infinite-type surface is a *fluted Loch Ness monster* if it is obtained from the Loch Ness monster in one of the two following ways:

- (1) by deleting a finite, nonzero collection of isolated points, or
- (2) deleting a countably infinite collection of isolated points accumulating to exactly one point, which we also delete from the surface, or accumulating onto the end of the Loch Ness monster.

See Figure 2 for examples of fluted Loch Ness monsters.

Lemma 2.9 Let Σ be an infinite-type surface with an isolated puncture. Then Σ is admissible unless Σ is a flute surface with finite (possibly zero) genus or is a fluted Loch Ness monster surface.

Proof Let Σ be an infinite-type surface with an isolated puncture p . If Σ has at least two nonplanar ends, then Σ admits a shift map (in fact a handleshift). Similarly, if Σ has at least two nonisolated planar ends, then Σ admits a shift map with these two ends corresponding to the two ends of the strip S' in Definition 2.5. Thus, if Σ does not admit a shift map, Σ has exactly one nonisolated planar end and finite genus, ie a flute surface with finite genus, or has exactly one nonplanar end and up to one nonisolated planar end, ie a fluted Loch Ness monster. \square

Going back to the original definition of an admissible surface, there are a few more notable remarks regarding the relationship between S and Σ . First, there is not necessarily an embedding of $\text{Map}(S, p)$ into $\text{Map}(\Sigma, p)$ since if the support of a shift h on S contains punctures that do not correspond to elements of \mathcal{U} , then there may not be a way to extend that shift to Σ . In particular, if h shifts one puncture x to another puncture x' but the topology of the surfaces glued to x and x' are different, there is no extension of h to a shift of Σ . This will not affect our arguments since there are countably many punctures of S corresponding to the elements of \mathcal{U} which we move to the *front* of the cylinder for S along with p , and we move all other

punctures to the *back* of the cylinder. Here we are choosing one nonisolated end of S to correspond to the left direction on the surface, and the other nonisolated end to correspond to moving right on the surface so that there is a well-defined notion of the front and back of S . In our constructions, we use shift maps on S whose support only contains the punctures on the front of S so that all of these shifts extend to Σ .

Second, and most importantly, we now show that proving Theorem 1.1 for admissible surfaces Σ can be reduced to the case of the biinfinite flute S . In fact, this is the motivation for the original definition of an admissible surface. For simplicity, given a surface M with an isolated puncture p and any points $a, b \in \mathcal{A}(M, p)$, we write $d_M(a, b)$ for the distance between a and b in $\mathcal{A}(M, p)$.

Lemma 2.10 *Let Σ be an admissible infinite-type surface with an isolated puncture p . Then the inclusion of $\mathcal{A}(S, p)$ into $\mathcal{A}(\Sigma, p)$ is a $(2, 0)$ -quasi-isometric embedding.*

Proof As $S \subset \Sigma$, it is clear that $d_\Sigma(a, b) \leq d_S(a, b)$ for any $a, b \in \mathcal{A}(S, p)$.

To obtain the other inequality, let $S_{a,b} \subset S$ be a finite-type subsurface of S which contains a, b , the puncture p , and has complexity at least 2. Note that $S_{a,b}$ is then a finite-type subsurface of Σ as well. Thus by [2, Corollary 4.3] applied to $S_{a,b} \subset \Sigma$ and to $S_{a,b} \subset S$, we have

$$d_S(a, b) \leq d_{S_{a,b}}(a, b) \leq 2d_\Sigma(a, b).$$

Together, these imply that

$$d_\Sigma(a, b) \leq d_S(a, b) \leq 2d_\Sigma(a, b),$$

completing the proof. \square

In particular, let $g \in \text{Map}(S, p)$ be loxodromic with respect to the action of $\text{Map}(S, p)$ on $\mathcal{A}(S, p)$ with a (K, C) -quasigeodesic axis. If g can be extended to an element of $\text{Map}(\Sigma, p)$, then this extension is loxodromic with respect to the action of $\text{Map}(\Sigma, p)$ on $\mathcal{A}(\Sigma, p)$, and the extension will have a $(2K, C)$ -quasigeodesic axis.

3 Coding arcs and standard position

Let S be the biinfinite flute surface with a distinguished isolated puncture p , and let $\{p_i\}_{i \in \mathbb{Z}}$ be any countably infinite discrete collection of punctures on S which exits both ends of the cylinder and does not contain p . As described in Section 2.4, we choose one nonisolated end of S to correspond to the left direction and one to correspond to the right direction, which gives a well-defined notion of a front and back of the cylinder for S . We move all of the punctures in $\{p_i \mid i \in \mathbb{Z}\} \cup \{p\}$ to the front of the cylinder for S and all other punctures to the back. We also move the distinguished puncture p so that it lies to the right of p_{-1} and to the left of p_0 . We will consider the collection $\{p_i \mid i \in \mathbb{Z}\} \cup \{p\}$ of punctures. We index this set with $\mathbb{Z} \cup \{P\}$, which we give the ordering consisting of the usual ordering on \mathbb{Z} with the additional requirement that $-1 < P < 0$. The index P corresponds to the distinguished puncture p .




Figure 5: The curves B_i are in red. The blue region is the domain of the shift map H .

Fix the simple closed curve B_0 bounding the puncture p_0 on S shown in Figure 5. More formally, to define B_0 we fix a complete hyperbolic metric on S and let B_0 be a horocycle at a height sufficiently far out the cusp. Fix a shift map H on S whose domain contains exactly the collection $\{p_i\}$ for $i \in \mathbb{Z} \cup P$ and which shifts p_i to p_{i+1} for all $i \in \mathbb{Z} \cup P$.

Definition 3.1 Define the simple closed curves $B_i := H^i B_0$ for $i \in \mathbb{Z} \cup P$. Then B_i is a simple closed curve bounding the puncture p_i , where $p_P = p$.

Our choice of left/right also gives a well-defined notion of an arc passing *over* or *under* a puncture (or equivalently some B_i). In all pictures of S throughout the paper, we denote the special puncture p by an “X”, and rather than drawing the punctures p_i , we draw the simple closed curves B_i in S . We will use these simple closed curves to put arcs into *standard position* as described later in this section.

3.1 Coding arcs

We use the simple closed curves B_i to describe a way to code simple arcs on S starting and ending at p . We will use this code to quantify how long two arcs fellow travel, which will be essential for proving the results of this paper.

Suppose that γ is an oriented arc on S starting and ending at p such that γ can be homotoped to be completely contained on the front of S . We code γ as follows. First homotope γ so that it is disjoint from all B_i with $i \in \mathbb{Z} \cup P$, with the exception that γ starts and ends at the puncture p and therefore intersects B_P exactly twice. The code for γ always starts and ends with the character P_s (which stands for “puncture start”) and contains either the character k_o or the character k_u , where $k \in \mathbb{Z} \cup \{P\}$, whenever γ




Figure 6: Pictured are arcs on the front of the surface S . The X denotes the puncture p and the elements $k \in \mathbb{Z} \cup \{P\}$ shown under S denote the subscript on the simple closed curves B_k . The codes for arcs α, β, γ are given in Example 3.2.

passes over or under the simple closed curve B_k for $k \in \mathbb{Z} \cup P$. These characters appear in the code for γ in the same order in which γ passes over/under the curves B_k . For example, since $-1 < P < 0$, the second character of the code for γ must be either $0_o, 0_u, (-1)_o, (-1)_u, P_o$, or P_u , because if γ doesn't immediately wrap around p (which would lead to the second character being P_o or P_u), it must pass over or under either B_0 or B_{-1} before it can pass over or under B_k for any $k \neq 0, -1$. Similarly, if the character 2_o or 2_u appears in the code, each adjacent character must be one of $1_o, 1_u, 2_o, 2_u, 3_o$, or 3_u . To simplify notation, we write $k_{o/u}$ to mean that the character could be k_o or k_u . We will write $k_{o/u}k_{u/o}$ to mean that the two adjacent characters are either $k_o k_u$ or $k_u k_o$; the $k_{u/o}$ is used to emphasize that the second character has the opposite subscript as the first one.

Example 3.2 Consider the arcs shown in Figure 6. The elements $k \in \mathbb{Z} \cup \{P\}$ shown under S denote the subscript on the simple closed curves B_k . The code for α is $P_s 0_o 1_u 2_o 2_u 1_u 0_u P_s$, the code for β is $P_s P_u P_o 0_o 1_o 2_o 2_u 1_o 0_o P_s$, and the code for γ is $P_s (-1)_o (-2)_o (-2)_u (-1)_u P_u 0_u 1_u 1_o 0_o P_s$.

Now suppose γ is an oriented arc on S starting and ending at p such that no arc in its homotopy class is contained on the front of S . Since γ starts and ends at p , which is on the front of the surface, every time γ leaves the front of S it must eventually reenter the front. We give the code C to any subpath of γ which is on the back of S . Up to homotopy, we may assume that each time γ exits then enters the front of S , it does so “between” two simple closed curves B_k and B_{k+1} . In other words, there is an arc γ' in the homotopy class of γ whose code contains either $k_{o/u}C(k+1)_{o/u}$ or $k_{o/u}Ck_{o/u}$ each time γ' leaves the front of S . We give γ the same code as γ' . We emphasize that this implies that the code of an arc does not distinguish the behavior of arcs γ on the back of S .

By an abuse of notation, we typically blur the distinction between an arc and its code, writing, for example, $\alpha = P_s 0_o 1_u 2_o 2_u 1_u 0_u P_s$.

Definition 3.3 Let γ be an oriented arc on S starting and ending at p . A code for γ is *reduced* if no two adjacent characters in the code are the same and if the character immediately following the initial P_s or preceding the terminal P_s is not $P_{o/u}$.




Figure 7: The arcs γ and γ' are homotopic and have the same reduced code: γ' is formed from γ by the removal of the pair $1_o 1_o$.

The appearance of repeated characters in the code of an arc indicates backtracking in the arc. The following lemma is immediate.

Lemma 3.4 *If there are two arc γ and γ' , starting and ending at p , whose codes differ only by the removal of two adjacent characters which are equal, ie $k_o k_o$ or $k_u k_u$, then γ and γ' are homotopic.*

Example 3.5 The arcs γ and γ' in Figure 7, with codes $P_s 0_o 1_o 1_o 1_u 2_u 2_o 1_o 0_o P_s$ and $P_s 0_o 1_u 2_u 2_o 1_o 0_o P_s$, respectively, are homotopic.

Note that if a triple appears in the code for an arc, it is reduced to a single character according to our convention, as only *pairs* of repeated characters are removed. For example, $P_s 0_o 1_o 1_o 1_u 0_o P_s$ is reduced to $P_s 0_o 1_o 1_u 0_o P_s$.

Each homotopy class of curves on S determines a reduced code, in the sense that any two homotopic curves have the same reduced code. We write that two codes are equal if they determine homotopic arcs. For example, we write

$$P_s 0_o 1_o 1_o 1_u 2_u 1_u 0_u P_s = P_s 0_o 1_u 2_u 1_u 0_u P_s.$$

The converse of this fact is not true, however, because the code does not encode the behavior of arcs on the back of S ; hence there can be nonhomotopic arcs with the same reduced code. This will not cause any problems in this paper.

Definition 3.6 The *code length* of an arc γ , denoted $\ell_c(\gamma)$, is the number of characters in a reduced code for γ .

Convention 3.7 When giving the code of an arc for which the numerical values of the characters are unimportant (or unknown), we will use variables in the code. Our convention is to use Roman letters to represent single characters and Greek letters to represent strings of characters whose length is (possibly) greater than one. For example, $\ell_c(a_1 a_2 a_3) = 3$ while $\ell_c(a \gamma b) = \ell_c(\gamma) + 2$.

Given a string of characters $\alpha = a_1 a_2 \dots a_n$, we denote by $\bar{\alpha}$ the reverse of α , so that $\bar{\alpha} = a_n a_{n-1} \dots a_2 a_1$. If α is an arc, then $\bar{\alpha}$ is the same arc with the opposite orientation.

3.2 Standard position

In this section, we describe how to use the code for an arc to find its image under a general class of shifts which we call “permissible”.

Definition 3.8 We say a shift shifting to the right is a *right shift*, while a shift shifting to the left is a *left shift*. A right shift is *permissible* if its domain D stays on the front of our subsurface and contains a *turbulent region* (n_1, n_2) , that is, there exist $n_1, n_2 \in \mathbb{Z} \cup \{P\}$ with $n_1 < n_2$ such that D contains B_k for all $k \in (-\infty, n_1] \cup [n_2, \infty)$ but does not contain B_k for any $k \in (n_1, n_2)$. We call $(-\infty, n_1] \cup [n_2, \infty)$ the *shift region* of h . See Figure 11. Analogously, a left shift is *permissible* if its domain D stays on the front of our subsurface and contains a *turbulent region* (n_2, n_1) , that is, there exist $n_1, n_2 \in \mathbb{Z} \cup \{P\}$ with $n_2 < n_1$ such that D contains B_k for all $k \in (-\infty, n_2] \cup [n_1, \infty)$ but does not contain B_k for any $k \in (n_2, n_1)$. The shift region for a left shift is $(-\infty, n_2] \cup (n_1, \infty)$.

Convention 3.9 Throughout the paper, we will use both left and right shifts. For notational simplicity, all general results about shifts will be stated for right shifts. All statements of results, proofs, and figures will make this assumption as well. However, all of our definitions and results (and their proofs) also hold for left shifts, by modifying any proof for a right shift so that we essentially replace all instances of n_1 with n_2 and vice versa and replace all instances of the word “increasing” by the word “decreasing” and vice versa. The only subtleties are that:

- We retain the convention that $h(B_{n_1}) = B_{n_2}$.
- For the shift region intervals $(-\infty, n_1] \cup [n_2, \infty)$ that appear for a right shift, we use $(-\infty, n_2] \cup (n_1, \infty)$ for the left shift. In particular, the n_2 is always contained in the shift region.

Remark 3.10 It is worthwhile to mention that Convention 3.9 is equivalent to simply redefining the order, given by the symbol $<_{\text{rev}}$, on $\mathbb{Z} \cup \{P\}$ to be the opposite of the standard meaning of the inequality sign $<$. For example, in this “reversed order” we would have $5 <_{\text{rev}} 3$ and so on. Given this and using the standard meanings for “increasing” and “decreasing” with respect to $<_{\text{rev}}$, all of the proofs for shifts that shift to the left would go through identically as shifts that shift to the right when one replaces each instance of $<$ with an $<_{\text{rev}}$. Despite the simplicity of this reversed order, we found writing proofs with it to be more confusing to the reader than applying the above convention.

In order to find the image of an arc using only its code, we will need to consider paths whose endpoints are not on p .

Definition 3.11 A *segment* is a simple path with at least one endpoint which is not a puncture, and no endpoints on a puncture other than p . We code a segment in an analogous way as we did arcs in Section 3.1. If a segment begins or ends on p , then the initial or terminal character of the code is P_s , respectively.

Note that a segment can have at most one instance of P_s in its code. Given a segment γ , we denote the initial and terminal character of its code by γ^i and γ^t , respectively. A segment is *supported on* an interval $(a, b) \subset \mathbb{Z} \cup \{P\}$ if the numerical value of every character in its reduced code is contained in (a, b) . A subsegment of γ which is supported on an interval (a, b) is denoted $\gamma|_{(a,b)}$, with similar notation for half-open and closed intervals. A segment is (strictly) *monotone* if the numerical value of the characters in its reduced code are (strictly) monotone as a subset of $\mathbb{Z} \cup \{P\}$. A segment with code C is called a *back loop*.

The notion of left/right on the front of S induces an orientation on strictly monotone segments contained on the front of S in the following way. If the terminal endpoint of such a segment γ is to the right of the initial endpoint, then γ is *oriented to the right*. Similarly, if the terminal endpoint is to the left of the initial endpoint, then γ is *oriented to the left*. Since γ is strictly monotone, one of the above two possibilities must occur. We note that single characters of a code represent strictly monotone segments and so can be oriented in this way.

We will use the code for an arc or segment to find the image of the arc or segment under certain homeomorphisms of S . The process can be complicated. We now introduce a new way of concatenating strings of characters which will be more suited to finding the image of an arc or segment in certain situations.

Definition 3.12 Given two segments α and β such that the terminal character of α agrees with the initial character of β and such that these two characters have the same orientation, the *efficient concatenation* of α and β , denoted $\alpha + \beta$, is formed by removing the terminal character of α to form a new string α' and concatenating this new string with β , resulting in $\alpha'\beta$.

For example,

$$P_s 0_o + 0_o 1_o = P_s 0_o 1_o,$$

and

$$P_s 0_o 1_o 2_o 2_u + 2_u 2_u 2_o 1_o 1_u = P_s 0_o 1_o 2_o 2_u 2_u 2_o 1_o 1_u = P_s 0_o 1_u,$$

where the middle term is an unreduced code and the final term is a reduced code. See Figure 8. We note that if α and β can be efficiently concatenated, then they cannot be concatenated, because α^t and β^i have the same orientation. By a similar reasoning, if two segments can be concatenated, then they cannot be efficiently concatenated. Throughout the paper, we only (efficiently) concatenate two segments when it is possible.

As written, the code of a segment is not well behaved under homotopy because every segment is homotopically trivial or homotopic into a puncture. We will introduce a standard position for segments on S with the property that any two segments that are homotopic rel endpoints will, in standard position, have the same reduced code. Standard position will also allow us to find the image of a segment under a permissible shift using only its code.




Figure 8: The two examples of efficient concatenation following Definition 3.12. In each, the (reduced) code for the blue segment is the efficient concatenation of the codes of the two black segments.

Definition 3.13 Fix a simple closed curve S_j as in Figure 9 for each $j \in \mathbb{Z} \cup \{P\}$. Let

$$\mathcal{C} = \{S_j \mid j \in \mathbb{Z} \cup \{P\}\}.$$

For each $j \in \mathbb{Z} \cup \{P\}$, we orient the simple closed curve B_j clockwise and identify B_j with the subset \mathbb{S}^1 of \mathbb{C} by a homeomorphism which preserves this orientation. Fix points $b_L^j, b_R^j \in B_j$ corresponding to $-1, 1 \in \mathbb{S}^1$, respectively. Here L and R stand for *left* and *right*.

To describe standard position, we will sometimes move endpoints of segments γ to lie on various boundary components. When we do this, we will use the following convention. Suppose $\gamma^i = k_{o/u}$ and we want to move the initial endpoint of γ onto the boundary component B_k . If γ is oriented to the right, then we move the initial endpoint of γ to b_L^k , and if it is oriented to the left, we move the initial endpoint to b_R^k . On the other hand, suppose $\gamma^t = (k')_{o/u}$ and we want to move the terminal endpoint of γ onto the boundary component $B_{k'}$. If γ is oriented to the right, we move the terminal endpoint of γ to $b_R^{k'}$, and if it is oriented to the left, we move the terminal endpoint to $b_L^{k'}$. See Figure 10. Moving endpoints in this way does not change the code for γ .

It will be useful to understand how a given segment interacts with the domain of a permissible shift.




Figure 9: Some of the simple separating curves in \mathcal{C} defined in Definition 3.13.




Figure 10: Adjusting the endpoints of γ (black) results in the red segment.

Definition 3.14 Let h be a permissible right shift with domain D and turbulent region (n_1, n_2) . Then both (the upper and lower) boundary components of D have the same reduced code on (n_1, n_2) . For any $k, k' \in (n_1, n_2)$, we let $\partial D|_{[k, k']}$ be the reduced code of the boundary components of D on the interval $[k, k']$ with similar notation for open and half-open intervals. If a segment γ has the same code as ∂D on an interval $(k_1, k_2) \subset (n_1, n_2)$, we say that γ follows ∂D or agrees with ∂D on that interval. As noted in the convention above, there is an analogous definition when h is a left shift. When a segment γ supported on $[n_1, n_2]$ intersects one component of ∂D , we call this a *half crossing*. When such a γ intersects both components of ∂D so that the code for the subsegment of γ between these two half crossings is empty, in the sense that this subsegment does not pass over or under B_k for any k , we call this a *full crossing*. See Figure 11.

Even though S is a straightforward surface, standard position for segments on S is necessarily complicated. Before introducing it formally in the next two subsections, we briefly give an intuitive idea of standard position in the following remark. On a first reading of this paper, we strongly suggest the reader read this remark and study Figures 12–15 instead of reading the formal definition of standard position given in Sections 3.2.1 and 3.2.2. The reader may then safely skip to Section 3.3. If, later in the paper, the image of a particular segment seems counterintuitive, likely this is because we put the segment in standard position before taking its image. This would be a good time to look back at Sections 3.2.1 and 3.2.2 with that example of a segment in mind.

Remark 3.15 Let h be a permissible right shift with domain D and turbulent region (n_1, n_2) , and let γ be a segment. To put γ in standard position with respect to h , we first homotope subsegments of γ that are contained in the region $(-\infty, n_1) \cup (n_2, \infty)$ to be completely contained in D . For the subsegments of γ contained in the region $[n_1, n_2]$, we homotope γ so that crossings are full crossings whenever possible. We will always be able to make crossings full except near n_1 or n_2 , because n_1 and n_2 are where γ leaves




Figure 11: The domain D for a permissible shift is shown in red. The segment γ follows ∂D on $[n_1 + 1, n_1 + 3]$, has a half crossing between B_{n_2-1} and B_{n_2} , and has a full crossing between B_{n_1+1} and B_{n_1+2} .

the turbulent region and enters the shift region, where we have already ensured that it is contained in D . We also homotope γ to minimize the number of full crossings. If γ contains a loop $(n_1)_{o/u}(n_1)_{u/o}$ or $(n_2)_{o/u}(n_2)_{u/o}$, we homotope γ so that it crosses one connected component of ∂D immediately before the loop and the other connected component immediately after. In other words, γ has to enter one side of D before the loop and exit the other side of D after the loop. If γ contains the character C , we homotope it so that the subsegment of γ represented by C exits and reenters the front of S at the same place, in the sense that the endpoints of this subsegment lie on the same simple closed curve S_i . Finally, we also homotope each endpoint of γ to lie either on the nearest B_k (if it is in the shift region) or the nearest S_k (if it is in the turbulent region).

Standard position will be slightly different for those segments which contain back loops. We first discuss standard position for segments that do not contain back loops.

3.2.1 Segments without back loops Given a segment γ whose reduced code does not contain C and a permissible right shift h with domain D and turbulent region (n_1, n_2) , we put γ into *standard position with respect to h* as follows.

If the endpoints of γ are not contained in (n_1, n_2) and it is possible to homotope γ completely inside of the domain of h , we do so, and we move the endpoints of γ to lie on the B_k curve numbered by the first and last characters of the code as described above.

Otherwise, γ can be written as the concatenation $\gamma_1 \dots \gamma_k$ of disjoint connected maximal subsegments such that each γ_i is either

- (a) supported on either $(-\infty, n_1]$ or $[n_2, \infty)$; or
- (b) supported on (n_1, n_2) .

We now homotope each γ_i individually. If γ_i satisfies (a), then we move the endpoints of γ_i onto the B_k curve numbered by the first and last character of the code as above and homotope the interior of γ_i to lie completely inside the domain of h . We homotope segments γ_i satisfying (b) using the following procedure:

Step (i) If the initial character of the segment is $k_{o/u}$, move the initial endpoint of the segment onto the separating curve S_k if $k_{o/u}$ is oriented to the right and onto S_{k+1} if it is oriented to the left. If the terminal character of the segment is $k'_{o/u}$, move the terminal endpoint of the segment onto $S_{k'}$ if $k'_{o/u}$ is oriented to the left and onto $S_{k'+1}$ if it is oriented to the right. Now move the endpoints along the S_j containing them to reduce the number of full and half crossings, if possible, without creating any self-intersections. In particular, the endpoints should not lie in the domain of h . See Figure 12.

There is one caveat to the rule above. Note that when γ_i is type (b), then γ_{i+1} and γ_{i-1} are type (a), otherwise γ_i would not be maximal. In the case that either of these neighboring segments is exactly a

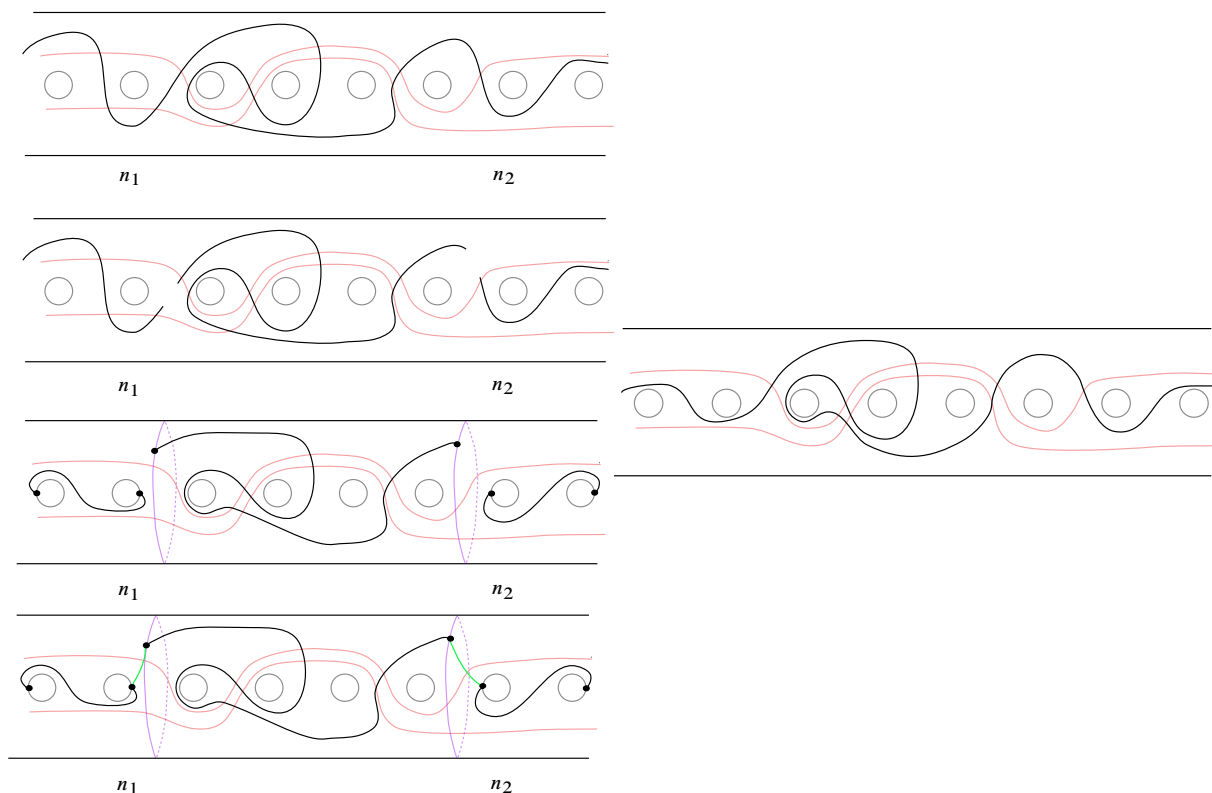


Figure 12: The domain D for a permissible shift is shown in red. In black is a segment γ that has no $(n_i)_o(u)(n_i)_u(o)$ for $i = 1, 2$. Left: putting γ in standard position. Right: a simple segment with the same code as γ which we say is in standard position.

loop around n_1 or n_2 we must adjust the position of the endpoints of γ_i along S_j . If $\gamma_{i+1} = (n_k)_o(n_k)_u$, for $k = 1$ or 2 , we require that the terminal endpoint of γ_i is above D and require that the terminal endpoint is below D when $\gamma_{i+1} = (n_i)_u(n_i)_o$. Similarly, if $\gamma_{i-1} = (n_i)_o(n_i)_u$, for $i = 1$ or 2 , we require that the initial endpoint of γ_i is below D and require that the initial endpoint of γ_i is above D when $\gamma_{i-1} = (n_i)_u(n_i)_o$. See Figure 13. Note that this repositioning of endpoints can cause additional crossings of D as in Figure 14, but this is the appropriate configuration for our calculations.

Step (ii) Homotope the segment rel endpoints to make all crossings full. Since Step (i) ensures that the endpoints of γ_i are always outside the domain D , this is always possible.

Step (iii) Homotope the segment rel endpoints to reduce the number of crossing by removing all bigons that bound disks and have one side on the segment γ_i and the other side on ∂D .

Step (iv) Finally, there may be a choice of where a full crossing occurs. If there is such a choice, then it will always be possible to homotope rel endpoints so that the crossing occurs between two adjacent characters $k_o(u), (k')_o(u)$ with $k, k' \in (n_1, n_2)$ such that the o/u pattern of k and/or k' does not match that of ∂D , and our convention is to make this choice for the largest possible k, k' . For example, in Figure 12,

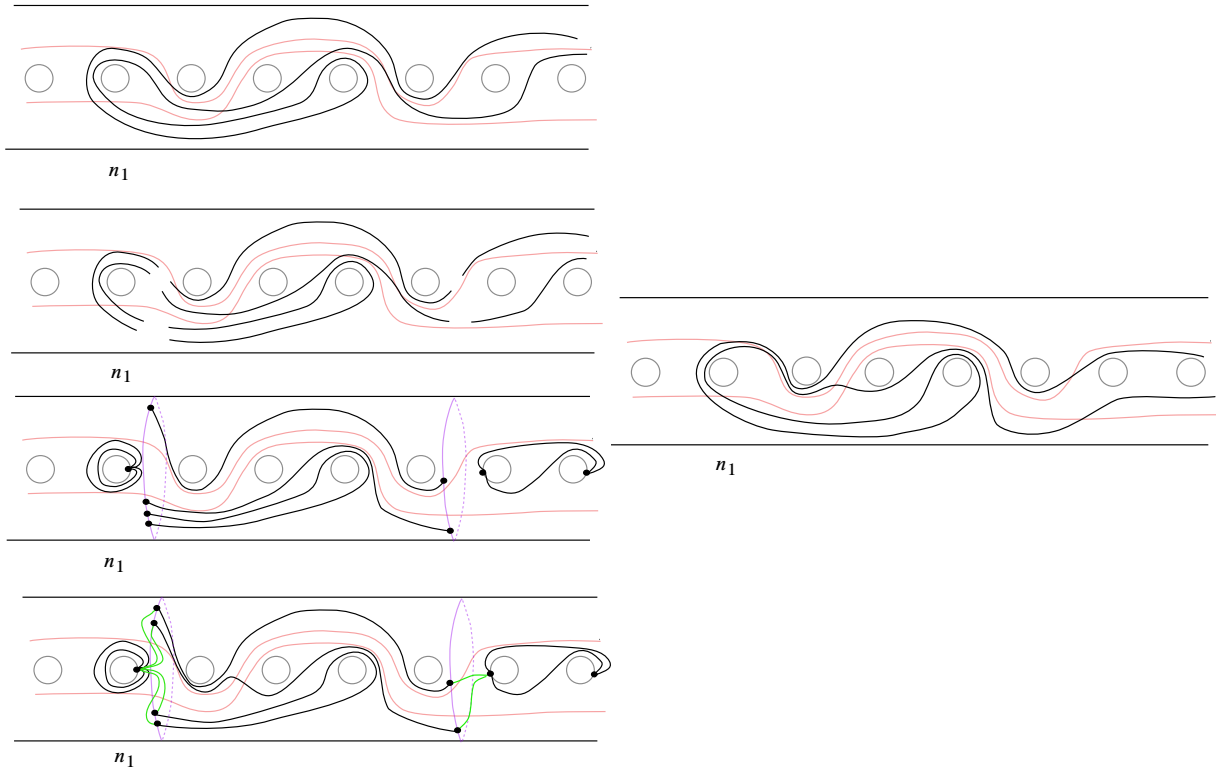


Figure 13: A segment γ that contains multiple copies occurrences of $(n_1)_o(n_1)_u$. Left: putting γ in standard position. Right: a simple segment with the same code as γ which we say is in standard position.

the bottom strand could fully cross the domain between $(n_1 + 1)_o$ and $(n_1 + 1)_u$ or between $(n_1 + 1)_u$ and $(n_1 + 2)_u$; we make the latter choice.

At this point, we have a collection of disjoint subsegments, each in standard position. The final step is to connect the endpoints of these segments, in the order they were originally connected, by segments called *connectors*. Connectors always occur between characters with numerical value n_1 and $n_1 + 1$ or n_2 and $n_2 - 1$. For the purposes of the code, we picture these segments extended slightly in either direction to overlap with the segments on either side so that the code of a connector will always be a pair $j_o u(j + 1)_o u$ or $(j + 1)_o u j_o u$. Thus, if a connector α connects disjoint subsegments δ_1 and δ_2 (in that order), then $\alpha = (\delta_1)^t (\delta_2)^i$ and we can write $\delta_1 + \alpha + \delta_2$. Note that this code is equivalent to the concatenation $\delta_1 \delta_2$. By construction, connectors always have one endpoint inside and one endpoint outside the domain D of the shift. After applying this procedure, the resulting segment in standard position may no longer be simple. However, it will have the same code as our original γ .

The following lemma summarizes the above procedure.

Lemma 3.16 *A segment γ in standard position with respect to a permissible right shift h which is not completely contained in the domain of h can be written as the efficient concatenation of (possibly empty)*

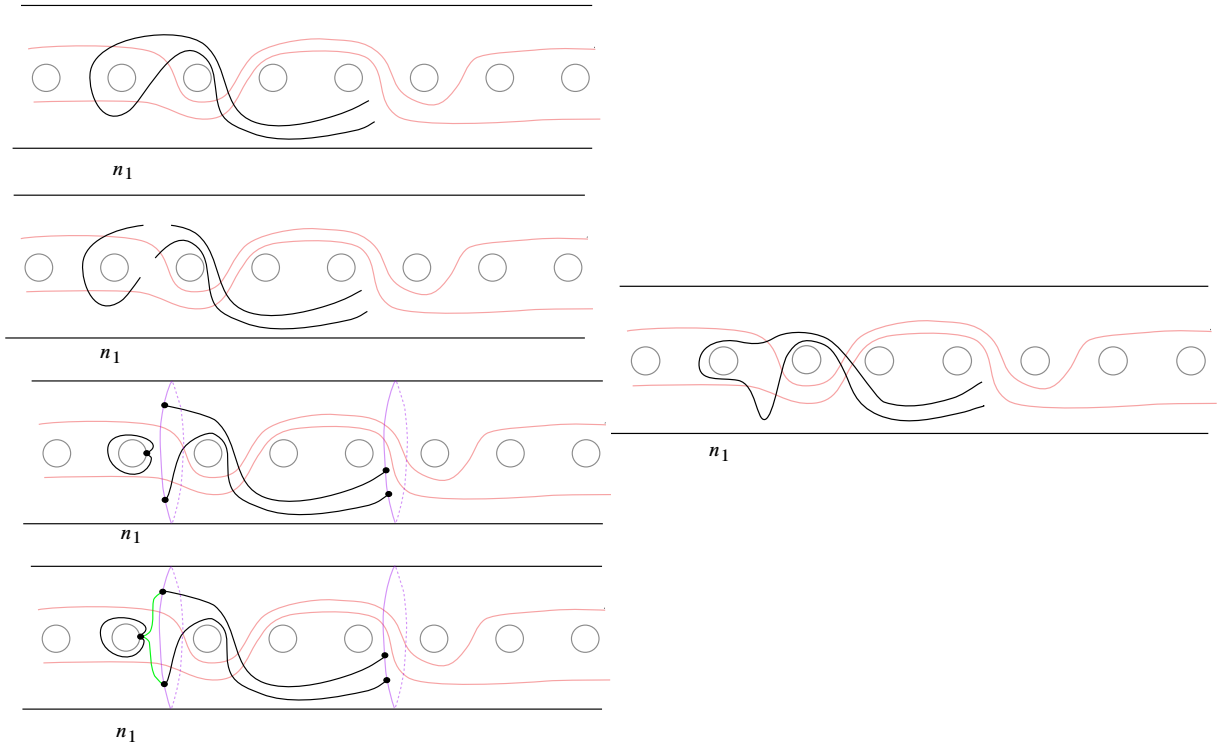


Figure 14: A segment γ that contains $(n_1)_o(n_1)_u$. Note that the requirements on the endpoints of γ_i in step (i) of the procedure for standard position require that we have one connector cross the top of ∂D and one connector cross the bottom of ∂D . Left: putting γ in standard position. Right: a simple segment with the same code as γ which we say is in standard position.

segments in the following way,

$$(1) \quad \gamma = \gamma_1^{\text{tu}} + \gamma_1^{c1} + \gamma_1^{\text{sh}} + \gamma_1^{c2} + \gamma_2^{\text{tu}} + \gamma_2^{c1} + \cdots + \gamma_n^{\text{sh}} + \gamma_n^{c2},$$

where for each i ,

- γ_i^{tu} is supported on the turbulent region (n_1, n_2) , has been put into standard position following steps (i)–(iv) above, and has both endpoints outside of the domain D ;
- γ_i^{sh} is supported on the shift region $(-\infty, n_1] \cup [n_2, \infty)$, is completely contained in the domain D , and has both endpoints on the B_k curves; and
- $\gamma_i^{c1}, \gamma_i^{c2}$ are connectors, each of which has code length 2, is supported on either $[n_1, n_1 + 1]$ or $[n_2 - 1, n_2]$, and has one endpoint on B_{n_1} or B_{n_2} and one endpoint outside the domain D on S_{n_1+1} or S_{n_2} .

We will sometimes use the notation γ_i^{conn} for a connector if it is not important if this subsegment is the first connector γ_i^{c1} or the second connector γ_i^{c2} .




Figure 15: A segment γ that has endpoints in $(-\infty, n_1] \cup [n_2, \infty)$ and can be homotoped to be completely inside D .

It is often convenient to abuse notation and say that a simple segment is in “standard position” even if it is not the result of the above procedure because these segments are easier to draw and think about. What we mean by this is that the segment intersects the boundary of D in the same place as it would in standard position. In Figures 12–15, we give a segment, the steps to put it in standard position, and an example of a simple segment with the same code which we also say is in standard position.

Because each segment in (1) has fixed endpoints, its image under h is well-defined up to homotopy rel endpoints. Thus we may use the decomposition of γ to find its image under h as

$$h(\gamma) = h(\gamma_1^{\text{tu}}) + h(\gamma_1^{c1}) + h(\gamma_1^{\text{sh}}) + h(\gamma_1^{c2}) + h(\gamma_2^{\text{tu}}) + h(\gamma_2^{c1}) + \cdots + h(\gamma_n^{\text{sh}}) + h(\gamma_n^{c2}).$$

Since in standard position γ may not be simple, its image under h may not be simple. However, this code corresponds to a unique (homotopy class of) simple segment with the same endpoints. It is important that we use *efficient* concatenation when calculating $h(\gamma)$. Using regular concatenation, the code for γ is simply $\gamma_1^{\text{tu}}\gamma_1^{\text{sh}} \dots \gamma_n^{\text{tu}}\gamma_n^{\text{sh}}$. However, it is not true that $h(\gamma_1^{\text{tu}})h(\gamma_1^{\text{sh}}) \dots h(\gamma_n^{\text{tu}})h(\gamma_n^{\text{sh}})$ is a code for γ ; in fact, much of the time this code does not define a segment. See Example 3.17. Most of the interesting behavior in the image of an arc or segment under a permissible shift actually comes from the full and half crossings. Since each connector contributes a half crossing, they are essential for determining the image of γ .

Example 3.17 Consider the permissible shift h shown in Figure 11 along with the segment γ . In Figure 16, we put γ in standard position and find its image under the shift. Using code, we have $\gamma = (n_1 + 1)_u(n_1 + 2)_o(n_1 + 3)_o(n_2 - 1)_o(n_2)_u$. Note that every character in this code except the terminal $(n_2)_u$ is fixed by h , and $h((n_2)_u) = (n_2 + 1)_u$. If we simply compute $h(\gamma)$ character by character, we




Figure 16: The segment γ from Example 3.17 in standard position (left) and its image under the shift whose domain is shown (right).

get $h(\gamma) = (n_1 + 1)_u(n_1 + 2)_o(n_1 + 3)_o(n_2 - 1)_o(n_2 + 1)_u$. However, this is not a well-defined code since $n_2 - 1$ and $n_2 + 1$ are not adjacent. On the other hand, using efficient concatenation, we see that

$$\begin{aligned} h(\gamma) &= h(\gamma^{\text{tu}}) + h(\gamma^{\text{conn}}) + h(\gamma^{\text{sh}}) \\ &= h((n_1 + 1)_u(n_1 + 2)_o(n_1 + 3)_o(n_2 - 1)_o) + h((n_2 - 1)_o(n_2)_u) + h((n_2)_u) \\ &= (n_1 + 1)_u(n_1 + 2)_o(n_1 + 3)_o(n_2 - 1)_o + (n_2 - 1)_o(n_2)_o(n_2 + 1)_u + (n_2 + 1)_u \\ &= (n_1 + 1)_u(n_1 + 2)_o(n_1 + 3)_o(n_2 - 1)_o(n_2)_o(n_2 + 1)_u. \end{aligned}$$

Note that all but the final pair $(n_2 - 1)_o(n_2)_u$ are fixed by h . This final pair is a half crossing and in standard position it is a connector.

In fact, given a segment $\gamma = \gamma^{\text{tu}}$ supported on (n_1, n_2) , we can further decompose γ into subsegments which are disjoint from D and pairs which fully cross D in such a way that makes it straightforward to find its image under h . Write

$$(2) \quad \gamma = \gamma_1^d + \gamma_1^e + \cdots + \gamma_s^d + \gamma_s^e,$$

where each γ^d is a maximal subsegment disjoint from D , each γ^e fully crosses D , and $\ell_c(\gamma^d) \geq 2$, $\ell_c(\gamma^e) = 2$, when nonempty.

Using the above decomposition, in an unreduced code we have

$$h(\gamma) = \gamma_1^d + h(\gamma_1^e) + \cdots + \gamma_s^d + h(\gamma_s^e).$$

Every nonempty γ_j^e will have image which follows ∂D , loops around n_2 , and follows ∂D back, so that

$$h(\gamma_j^e) = \partial D|_{[(\gamma_j^e)^i, n_2)}(n_2)_o/u(n_2)_u/o\overline{\partial D|_{[(\gamma_j^e)^t, n_2)}}.$$

As in Example 3.17, if we don't use efficient concatenation then we can write $\gamma = \gamma_1^d \gamma_2^d \cdots \gamma_s^d$. Applying h to each of these subsegments individually would yield $h(\gamma) = \gamma$, since each of these subsegments is fixed by h , which is not the correct image.

3.2.2 Segments with back loops If γ is a segment with code equal to C , we require that γ has both endpoints on some separating curve S_i in our collection \mathcal{C} . We also assume that the endpoints of C lie outside the domain, and, moreover, that γ does not intersect D . There are two possibilities for γ , either γ both enters and exits the front of S at the top or bottom or (up to taking inverses) γ enters at the top and exits at the bottom of the front of the surface. Recall that we define the top/bottom of the front of S with respect to the notion of right/left on the front of S . In the first case, this implies that the endpoints of γ are both above or both below D , respectively, while in the second case one will be above and one will be below. In either case, this convention implies that $\gamma \cap D = \emptyset$ and $h(\gamma) = \gamma$.

Suppose next that γ is a segment whose code contains C but also contains other characters. Suppose for simplicity that the code for γ contains a single instance of C , so that $\gamma = \tau_1 C \tau_2$, where τ_i does not contain C for $i = 1, 2$. We put γ in standard position as follows. First note that by definition of the code, we must either have $(\tau_1)^t = (\tau_2)^i$ (if τ_1 and τ_2 have opposite orientations) or $(\tau_1)^t = (\tau_2)^i \pm 1$ (if $(\tau_1)^t$ and $(\tau_2)^i$

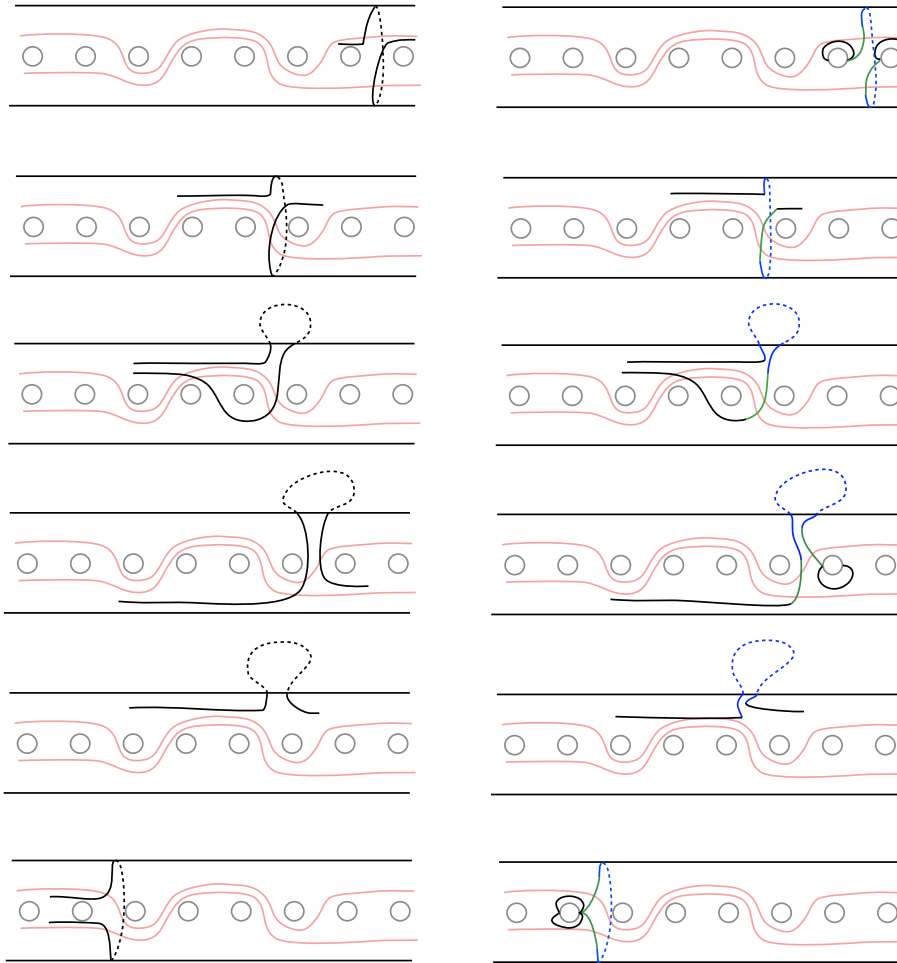


Figure 17: Various segments γ containing back loops. The segments on the right are in standard form. Back loops are blue and back loop connectors are green. Note that in the second, third, and fifth examples, at least one of C^\pm is empty.

have the same orientation). See Figure 17. Without loss of generality, suppose τ_1^t is oriented to the right. Consider the (disjoint) segments τ_1 and τ_2 , put them in standard position as in the previous section, and homotope the endpoints of C as in the previous paragraph. Note the endpoints of C will lie on the curve $S_{(\tau_1)^t+1}$. We now have three disjoint segments with codes τ_1 , τ_2 , and C . We will add (possibly empty) segments called *back loop connectors* from the terminal point of τ_1 to the initial point of C and from the terminal point of C to the initial point of τ_2 , respectively, to form a connected segment. See Figure 17. We code these back loop connectors with the characters C^- , C^+ , respectively, so that we can easily discuss their image. In particular, by a slight abuse of notation, we replace C in the code for γ with C^-CC^+ .

Recall that without loss of generality, $(\tau_1)^t$ is oriented to the right. If the terminal endpoint of τ_1 lies outside D and on the same side of D as the initial point of the back loop C , then the back loop connector



Figure 18: A segment γ (left) is split into subsegments $\gamma = \gamma_1 + \gamma_2$ and each γ_i is individually put into standard position with respect to the shift h whose domain is shown (right). This causes a loss of information when we take the images of γ_i individually under h : since $h(\gamma_i) = \gamma_i$ for $i = 1, 2$, we have $h(\gamma_1) + h(\gamma_2) = \gamma_1 + \gamma_2 = \gamma$, but γ is clearly not fixed by h .

C^- is empty. An analogous statement holds for C^+ , where we use the initial endpoint of τ_2 in place of the terminal endpoint of τ_1 . Each nonempty back loop connector C^\pm either

- (i) has one endpoint on $B_{(\tau_1)^t}$ or $B_{(\tau_2)^i}$ (depending on whether it is C^- or C^+), the other endpoint on $S_{(\tau_1)^t+1}$, and half-crosses D ; or
- (ii) has both endpoints on $S_{(\tau_1)^t+1}$ which are outside of D and fully crosses D .

3.3 Gaps in segments

When we use the code of a segment to find its image under a permissible shift, we first break it into smaller subsegments using standard position. When we do this, we always use efficient concatenation, so that the codes of the individual pieces overlap in a single character. The goal of efficient concatenation is to ensure that we do not lose any information about the segment by breaking it into pieces. However, we need to be careful when we do this. If we first break a segment into subsegments and then put each subsegment into standard position, it is possible that we will lose some information. In particular, we may cause there to be a “gap” in the segment. Based on standard position, these gaps can only occur when breaking a segment in the interior of the turbulent region (n_1, n_2) .

To see this, suppose we break a segment $\gamma = \gamma_1 + \gamma_2$ into two subsegments such that the numerical value of $\gamma_1^t = \gamma_2^i$ is $j \in (n_1, n_2)$. If we put each γ_i into standard position individually, it is possible that γ_1^t and γ_2^i lie on opposite sides of D (see Figure 18). In this case, we have lost the full crossing between them. Recall that a shift fixes the surface outside of its domain. In the region (n_1, n_2) , a segment in standard position is disjoint from the domain of the shift *except* where there is a full crossing (see the decomposition in equation (2)), so the full crossings are essential for determining the image of a segment. Thus we cannot use γ_1 and γ_2 to find the correct image of γ .

In order to ensure that we do not lose any information when working with a segment and subsegments in the turbulent region, we always first homotope the whole segment γ to be a simple segment in standard position. We then fix this particular representative of the homotopy class of γ for the remainder of the time we work with it. Thus, when we break γ into subsegments γ_1 and γ_2 , we do *not* put these into standard position individually. To be precise, we choose the endpoints γ_1^t and γ_2^i to be the intersection of




Figure 19: A segment γ and its image $h(\gamma)$ are shown. Left: a subsegment which does not have a preimage. The pink subsegment of $h(\gamma)$ does not have an inverse; it is a proper subpath of the image of the purple subsegment of γ . Right: a subsegment which has a preimage. The pink subsegment of $h(\gamma)$ has an inverse; it is the image of the purple subsegment of γ .

γ and the appropriate separating curves in our collection \mathcal{C} and we do not allow any further homotopies of γ_1 or γ_2 . This will always ensure that there are no gaps between γ_1 and γ_2 .

In certain cases, it may be simpler to break γ into subsegments in a different way, and we do this whenever it will avoid technicalities. For example, in Figure 18, we may simply choose not to divide γ into subsegments at all. On the other hand, we could also choose to make γ_1 or γ_2 longer than is strictly necessary in a particular calculation in order to avoid a potential loss of information.

3.4 Taking inverses of segments

In general, if we have a segment γ and a subsegment ζ of $h(\gamma)$, there is not necessarily a subsegment of γ which we may call $h^{-1}(\zeta)$. In other words, not every subsegment of $h(\gamma)$ is the image of a subsegment of γ . It is important here that when we think of a subsegment of γ , we are fixing γ in standard position. That is, we are thinking of a *reduced* code for γ , rather than any (unreduced) code representing γ .

For example, consider the shift shown in Figure 19, left. Here, the pink subsegment ζ_1 of $h(\gamma)$ is not the image of any subsegment of γ . Rather, it is properly contained in the image of the purple subsegment of γ , specifically because it is in the image of the full crossing of the purple segment with D . Explicitly,

$$\gamma = 0_o 1_u 2_o 2_u 1_u, \quad h(\gamma) = 0_o 1_u 2_o 2_u 1_u 0_o 0_u 1_u 2_u 2_o 1_u 1_o 2_o 2_u 1_u 0_u,$$

and one can see that if the subsegment $\zeta_1 = 1_u 2_o 2_u 1_u$ of $h(\gamma)$ then there is no subsegment γ' of γ for which $h(\gamma') = \zeta_1$.

However, we may take an inverse image of a subsegment ζ of $h(\gamma)$ whenever we *know* that ζ is the image of a subsegment of γ . This is the case, for example, when $\zeta = h(\gamma)$; in other words, it is true that $\gamma = h^{-1}(h(\gamma))$. This can also happen when the initial and terminal characters of a subsegment of $h(\gamma)$

are the images of the initial and terminal characters of a subsegment of γ . For example, in Figure 19, right, a direct computation will show that the initial and terminal characters of ζ_2 (in pink) are the images of the initial and terminal characters of the purple subsegment of γ . Therefore ζ_2 is the image of the purple subsegment of γ , which is precisely the result of calculating $h^{-1}(\zeta_2)$.

4 Characterizing loops with trivial image

As in Section 3, we let S be the biinfinite flute surface with a distinguished puncture p and fix the collection of simple closed curves $\{B_i \mid i \in \mathbb{Z} \cup \{P\}\}$ as in Definition 3.1. Let h be a permissible shift (see Definition 3.8) and $k \in \mathbb{Z} \cup \{P\}$. By an abuse of notation, we may occasionally write $h(k)$, by which we mean that $h(k)$ is the label of $h(B_k)$. Thus, given any segment whose code is $k_{o/u}$, we have $h(k_{o/u}) = h(k)_{o/u}$.

Recall that a segment is a path which does not have both endpoints on p .

Definition 4.1 A segment in standard position is *trivial* if it can be homotoped rel endpoints to one of the following:

- (a) a segment contained in one of the separating curves $S_i \in \mathcal{C}$; or
- (b) a point.

We will use the notation \emptyset to denote the reduced code for a trivial segment. For example, we write $k_o k_o = \emptyset$.

Definition 4.2 A *loop* is a segment that has one of the following forms:

- (1) $\delta_1 a_{o/u} a_{u/o} \delta_2$ for some a where $\delta_1^i = \delta_2^t$,
- (2) $a_{o/u} a_{u/o}$ for some a , or
- (3) C .

See Figure 20 for examples of loops. A loop satisfying (1) is called a *regular loop*, a loop satisfying (2) is called an *over-under loop*, and, as in Definition 3.11, a loop satisfying (3) is called a *back loop*. Note




Figure 20: Some examples of loops from Definition 4.2. The loop in red fits case (1), the loop in blue fits case (2), and the loop in purple fits case (3).




Figure 21: The permissible shift h translates to the right. The first two loops on the left are homotopic via homotopies which keep the endpoints on the same fixed separating curve. However, one image is trivial while the other is not.

that regular loops always contain an over-under loop but not all over-under loops can be extended to a regular loop. A *single loop* is a back loop, an over-under loop, or a regular loop such that δ_i does not contain a loop for $i = 1, 2$.

With the above definitions, the rest of this section is devoted to analyzing the following question.

Question 4.3 *Let h be a permissible shift. When does h send a loop to a trivial segment?*

The reason we introduced standard position is to ensure that this question is well defined. The issue is that homotopies of a loop can change whether or not its image is trivial, even if those homotopies keep the endpoints on a fixed separating closed curve (see Figures 21 and 22). Thus it is important that, given a loop, we first put it in standard position before applying the permissible shift h . This will remove any possible ambiguity in the image of the loop.

In this section, we first introduce a kind of cancellation in the image of a segment and its code which we call *cascading cancellation*. This kind of cancellation will cause technical problems throughout the paper, and much of Section 6 is devoted to understanding how to control it. We then prove Theorem 4.8, which answers Question 4.3. We end the section with a discussion of several technical consequences of the Theorem 4.8 which will be useful later.




Figure 22: The second loop from Figure 21 is in standard position. We show why its image under h is trivial.

4.1 Cascading cancellation

Definition 4.4 We call an arc γ *symmetric* if $\gamma = \delta q_1 q_2 \bar{\delta}$ for any characters q_1, q_2 . In other words, a reduced code for γ is palindromic with the exception of the middle two characters. Note in particular that this implies that q_1, q_2 have the same numerical value.

Recall that we find the image of a path with code $\alpha\beta$ under a permissible shift f as follows. Let q be the last character of α and q' be the first character of β . Then

$$f(\alpha\beta) = f(\alpha) + f(qq') + f(\beta).$$

While $f(\alpha)$, $f(qq')$, and $f(\beta)$ are all reduced codes, it is possible that the efficient concatenation will cause there to be cancellation. For example, if $f(\alpha) = 1_u 1_o 2_o 3_o$ and $f(qq') = 3_o 3_o 2_o 1_o 0_o P_o$, then $f(\alpha) + f(qq') = 1_u 0_o P_o$. When this type of cancellation occurs, that is, when a character of $f(qq')$ does not appear in a reduced code of the image, we say there is *cancellation involving $f(qq')$* . In our example, there is cancellation involving $f(qq')$ and $f(\alpha)$. When it is necessary to be more precise, we may also say there is (respectively, is not) cancellation involving a character s , if s appears in the unreduced code but not the reduced code (respectively, appears in both the unreduced code and the reduced code) of the image. Thus in our example, there is no cancellation involving 1_u but there is cancellation involving 2_o . Our goal is to understand, in general, when there is cancellation with a particular character in a path under a permissible shift.

Suppose $\alpha = \gamma q_1$ and h is a permissible shift. Let q be the terminal character of γ . It is tempting to believe that if we can show that there is no cancellation involving $h(q_1)$ within $h(qq_1)$, then there is no cancellation involving $h(q_1)$ at all. However, this is not sufficient, for it is possible that there is “cascading cancellation”. Before giving a formal definition, we illustrate this phenomenon with an example.

Example 4.5 Consider the segment $\delta = 3_u 2_u 1_o 0_u$ and the permissible left shift h whose domain is shown in Figure 23. Then

$$h(\delta) = h(3_u 2_u) + h(2_u 1_o) + h(1_o 0_u).$$

We have

$$h(1_o 0_u) = 1_o 0_u, \quad h(2_u 1_o) = 2_u 1_o \quad \text{and} \quad h(3_u 2_u) = 3_u 2_u 1_o 0_u (-1)_u (-1)_o 0_u 1_o 2_u 2_u.$$

Putting this together, we obtain $h(\delta) = 3_u 2_u 1_o 0_u (-1)_u (-1)_o$.

There is no cancellation involving either of the terms when computing $h(2_u 1_o) + h(1_o 0_u)$. However, when computing $h(3_u 2_u) + h(2_u 1_o)$, we see that $h(2_u 1_o)$ completely cancels with an initial segment of $h(3_u 2_u)$, and $h(1_o 0_u)$ completely cancels with the next subsegment of $h(3_u 2_u)$. Therefore, there is in fact cancellation involving $h(0_u)$ in $h(\delta)$.

Definition 4.6 Formally, given an arc $\delta_1 \dots \delta_n$ and a permissible shift f , we say there is *cascading cancellation involving $f(\delta_n)$* if there is cancellation involving $f(\delta_n)$ in $f(\delta_1 \delta_2 \dots \delta_n)$ but there is no cancellation involving $f(\delta_n)$ in $f(\delta_{n-1} \delta_n)$.




Figure 23: Above: the segment δ in standard position. The domain of the left shift h is shown in green. Below: the image $h(\delta)$.

Understanding and controlling cascading cancellation is the difficult part of many of the proofs in this paper. The remainder of this subsection is devoted to theorems that will allow us to control cascading cancellation for permissible shifts. We will often be in the following situation: there is a segment $\gamma = \gamma_1 \gamma_2$ whose image under a shift h we would like to understand and we know that $h(\gamma_1)$ has some desired quality (such as containing a loop, for example). In order to show that the desirable behavior of $h(\gamma_1)$ persists in $h(\gamma)$, we need to ensure that there is no cancellation involving $h(\gamma_1)$ and $h(\gamma_2)$ by controlling cascading cancellation. In Section 6, we will revisit this topic and prove some additional results that allow us to control cascading cancellation for the particular homeomorphisms we construct, which are compositions of shifts.

Lemma 4.7 *Let h be a permissible right shift with domain D and turbulent region (n_1, n_2) . Let α be a strictly monotone segment supported on (n_1, n_2) . Then in a reduced code $h(\alpha)$ has n loops around n_2 , where n is the number of times α fully crosses D .*

Proof Without loss of generality, we will assume that α is strictly monotone increasing, that is, the numerical value α^i is strictly less than that of α^t , as the conclusion is invariant under replacing α by $\bar{\alpha}$. As in Section 3.2.1, put α in standard position. Since no segment in standard position which is supported on (n_1, n_2) will be completely contained in the domain of the shift, we may write

$$\alpha = \alpha_1^d + \alpha_1^e + \cdots + \alpha_s^d + \alpha_s^e,$$

where each α_j^d is a maximal subsegment disjoint from D , each α_j^e fully crosses D , and $\ell_c(\alpha_j^d) \geq 2$, $\ell_c(\alpha_j^e) = 2$ when nonempty. Notice that if $\alpha_j^d \neq \emptyset$ and $\alpha_{j+1}^d \neq \emptyset$, then also $\alpha_j^e \neq \emptyset$ by maximality.

Under the above decomposition, in an unreduced code we have

$$h(\alpha) = \alpha_1^d + h(\alpha_1^e) + \cdots + \alpha_s^d + h(\alpha_s^e),$$




Figure 24: Above: The case where $\alpha_j^e + \alpha_j^d + \alpha_{j+1}^e = \alpha_j^e + \alpha_{j+1}^e$ in the proof of Lemma 4.7, where α_j^e is in purple and α_{j+1}^e is in red. Below: the nontrivial image of this segment under h .

where every nonempty α_j^e will have image

$$h(\alpha_j^e) = \partial D|_{[(\alpha_j^e)^i, n_2)} (n_2)_{o/u} (n_2)_{u/o} \overline{\partial D|_{[(\alpha_j^e)^t, n_2)}}.$$

Thus each full crossing α_j^e contributes a loop around n_2 in an *unreduced* code for $h(\alpha)$. We must show that such loops persist in a *reduced* code for $h(\alpha)$.

If $\alpha_j^e, \alpha_{j+1}^e \neq \emptyset$ and $\alpha_{j+1}^d = \emptyset$ then a calculation shows that in a reduced code

$$h(\alpha_j^e + \alpha_{j+1}^d + \alpha_{j+1}^e) = h(\alpha_j^e + \alpha_{j+1}^e) = h(\alpha_j^e) + h(\alpha_{j+1}^e).$$

In particular, there is no cancellation between $h(\alpha_j^e)$ and $h(\alpha_{j+1}^e)$ and the loops around n_2 persist. See Figure 24.

Now assume that α_{j+1}^d is nonempty. The maximality of α_{j+1}^d implies that α_j^e and α_{j+1}^e are nonempty as well. We must show that there is no cancellation between the loops around n_2 in $h(\alpha_j^e)$ and $h(\alpha_{j+1}^e)$ so that both loops persist in a reduced code for $h(\alpha_j^e) + \alpha_{j+1}^d + h(\alpha_{j+1}^e)$. Recall that in standard position, a full crossing occurs between two adjacent characters $k_{o/u}, (k')_{o/u}$ with $k, k' \in (n_1, n_2)$ such that the o/u pattern of k and/or k' does not match that of ∂D , and our convention is to make this choice for the largest possible k, k' .

The subtlety arises because the code for $h(\alpha_j^e) + \alpha_{j+1}^d$ may not be reduced if an initial subsegment of $\alpha_{j+1}^d = h(\alpha_{j+1}^d)$ agrees with ∂D , in which case $(\alpha_j^e)^t$ agrees with ∂D . If $(\alpha_j^e)^t$ is the character of the full crossing α_j^e that does *not* agree with ∂D , then this character will block cancellation between $h(\alpha_j^e)$ and α_{j+1}^d , so that the loops around n_2 in $h(\alpha_j^e)$ and $h(\alpha_{j+1}^e)$ cannot cancel. On the other hand, if $(\alpha_{j+1}^e)^i$ is the character of the full crossing α_{j+1}^e that does *not* agree with ∂D , then this character will block cancellation between the two n_2 loops, even if α_{j+1}^d fully cancels in a reduced code for $h(\alpha_j^e) + \alpha_{j+1}^d + h(\alpha_{j+1}^e)$.

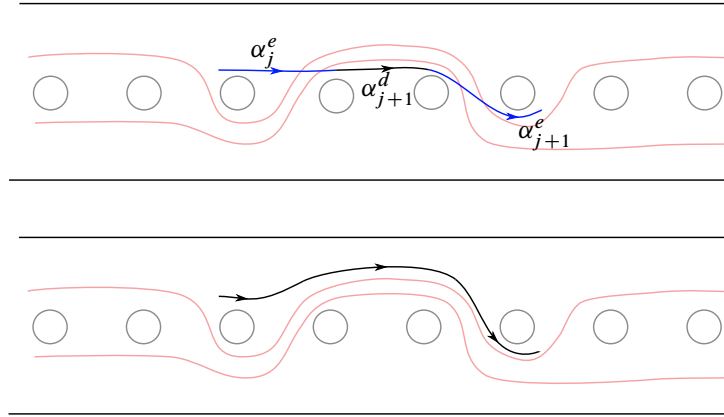


Figure 25: Above: A picture of $\alpha_j^e + \alpha_{j+1}^d + \alpha_{j+1}^e$, not in standard position, where $(\alpha_j^e)^t$ and $(\alpha_{j+1}^e)^i$ both agree with ∂D . Below: The same subsegment of α in standard position.

In the last case, where $(\alpha_j^e)^t$ and $(\alpha_{j+1}^e)^i$ both agree with ∂D , the largest possible choice of k, k' for the full crossing α_j^e would in fact result in a segment $\alpha_j^e + \alpha_{j+1}^d + \alpha_{j+1}^e$ that is fully disjoint from D , contradicting the maximality of α_{j+1}^d , so that this case cannot occur. See Figure 25.

Thus, there is no cancellation between the loops around n_2 in $h(\alpha_j^e)$ and $h(\alpha_{j+1}^e)$, which proves the lemma. \square

4.2 Loops with trivial image

In this subsection, we describe the form a loop must have if its image under a shift is trivial. In particular, the image of any loop which does not have the form as stated in the theorem is nontrivial. In addition, any segment that is not a loop cannot have trivial image under a shift since the numerical values the initial and terminal characters, and thus their images, differ.

Theorem 4.8 *Let h be a permissible right shift with turbulent region (n_1, n_2) . Suppose β is a nontrivial loop such that $h(\beta) = \emptyset$. Then either $\beta = k_o \delta k_o$ or $\beta = k_u \delta k_u$, where*

- (i) $k \in (n_1, n_2)$,
- (ii) $\delta = \gamma(n_1)_{o/u}(n_1)_{u/o}\bar{\gamma}$, and
- (iii) γ follows ∂D between k and n_1 .

Proof Put β in standard position. We will consider the image $h(\beta)$. By the discussion in Section 3.4, we have $\beta = h^{-1}(h(\beta))$. By assumption, $h(\beta)$ is trivial, and therefore it is homotopic rel endpoints to either a segment contained in the separating curve S_k for some $k \in [n_1, n_2]$ or a point. If $h(\beta)$ is homotopic rel endpoints to a point, then $\beta = h^{-1}(h(\beta))$ is also homotopic rel endpoints to a point, in which case β is trivial, which is a contradiction. So suppose that $h(\beta)$ is homotopic rel endpoints to an embedded subsegment σ of the separating curve S_k for some $k \in (n_1, n_2]$. Then β is homotopic rel endpoints to $h^{-1}(\sigma)$.

Let D denote the domain of h . Since β is in standard position, the endpoints of σ (which are the endpoints of β) are not contained in D . There are two possibilities. If $\sigma \cap D = \emptyset$, then β is homotopic rel endpoints to $h^{-1}(\sigma) = \sigma$ and so β is trivial, which is a contradiction. On the other hand, if $\sigma \cap D \neq \emptyset$, then σ must fully cross D . A direct computation shows that, up to reversing orientation,

$$h^{-1}(\sigma) = \overline{\partial D}|_{[k, n_1)}(n_1)_{o/u}(n_1)_{u/o}\partial D|_{(n_1, k]},$$

for $k \in (n_1, n_2)$, as desired. \square

We note one immediate consequence of Theorem 4.8.

Corollary 4.9 *Suppose β is any of the following:*

- (1) *a loop containing more than one single loop;*
- (2) *a loop containing a back loop; or*
- (3) *a loop with endpoints outside of the turbulent region.*

Then $h(\beta)$ is nontrivial.

Some of these consequences may be surprising, so we now give a more intuitive (and informal) discussion of why the image of loops as in the corollary are nontrivial.

We first consider the case of segments which contain more than one single loop. At first glance, it may seem that if a segment β is composed of loops which each satisfy the conclusion of Theorem 4.8, then $h(\beta)$ will be trivial. The problem arises in how these loops fit together to form β . Suppose we have two loops, β_1 and β_2 which each satisfy the conclusion of Theorem 4.8. If the numerical value of β_1^i is not the same as that of β_2^i , then in order to “connect” β_1 to β_2 , we must add a segment between the terminal point of β_1 and the initial point of β_2 . This segment is either disjoint from D , in which case it is fixed by h , or it fully crosses D . In the former case, this segment will appear in the image of β , while in the latter case, the image of this full crossing will be nontrivial by Lemma 4.7 applied to the full crossing. If the numerical values are the same and $\beta = \beta_1\beta_2$, then β will be trivial. On the other hand, if there is some segment connecting the terminal point of β_1 to the initial point of β_2 , then as before, the image of this segment will not be trivial and will force $h(\beta)$ to be nontrivial. See Figure 26 for an example of this.

Suppose next that β contains a back loop C . Intuitively, it seems reasonable that if $\beta = C\beta'\overline{C}$ and β' is a loop as in Theorem 4.8, then $h(\beta)$ is trivial. However, this is not the case. To see this, notice that if we put β' in standard position then it will have one endpoint above D and one below D . However, C and \overline{C} are the same loop. If C is either above or below D , then this will cause β to have a full crossing between β' and one of C or \overline{C} . The image of this full crossing will be nontrivial by Lemma 4.7 applied to the full crossing, which will prevent any cancellation between $h(C) = C$ and $h(\overline{C}) = \overline{C}$. On the other hand, if C exits at the top and enters at the bottom, then \overline{C} exits at the bottom and enters at the top. This will again force β to have a full crossing between β' and one of C or \overline{C} , preventing cancellation between the images of the back loops.

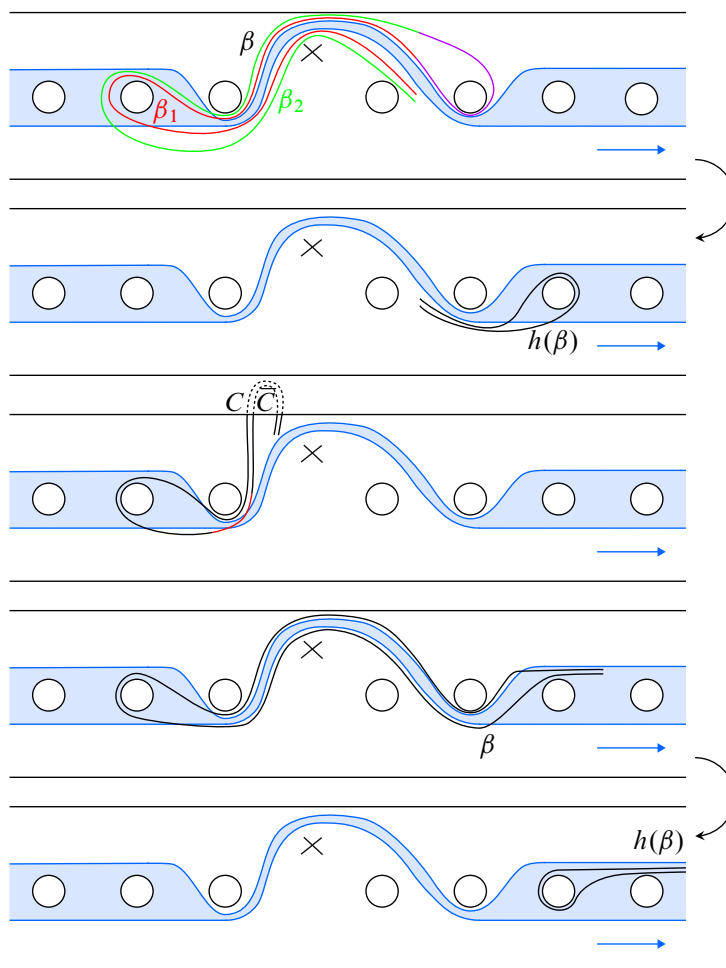


Figure 26: Examples of loops whose images are nontrivial. From top to bottom, these correspond to Corollary 4.9 (1), (2), and (3).

Finally, if β is a loop whose endpoints are outside the turbulent region, say both in $[n_2, \infty)$, and $\beta|_{(n_1, n_2)}$ has the form in the conclusion of Theorem 4.8, it seems possible that $h(\beta)$ is trivial. For example, suppose that $\beta = (n_2)_o \beta' (n_2)_o$, where β' is as in Theorem 4.8. Then since $h(\beta')$ is trivial, $h(\beta')$ will be a segment contained in the separating curve S_{n_2} . On the other hand, the images of the connectors $(n_2)_o (\beta')^i$ and $(\beta')^t (n_2)_o$ will each be of the form $(n_2 + 1)_o (n_2)_{o/u} (n_2 - 1)_{o/u}$ or $(n_2 - 1)_{o/u} (n_2)_{o/u} (n_2 + 1)_o$. In particular, we have $h(\beta) = (n_2 + 1)_o (n_2)_{o/u} (n_2)_{u/o} (n_2 + 1)_o$, which is nontrivial.

4.3 Consequences of Theorem 4.8

In this subsection, we record several (technical) consequences of Theorem 4.8 which will be useful in later sections. The first shows that loops in the shift region persist under the image of shifts.

Lemma 4.10 *Let h be a permissible right shift and γ any nontrivial, simple arc with no back loops. If $k \in (-\infty, n_1) \cup [n_2, \infty)$ and $k_{o/u}k_{u/o}$ appears in a reduced code for γ , then there is no cancellation involving $h(k_{o/u}k_{u/o})$ in an unreduced code for $h(\gamma)$. In particular, such a pair $k_{o/u}k_{u/o}$ will yield a pair $(k+1)_{o/u}(k+1)_{u/o}$ in a reduced code for $h(\gamma)$.*

Proof We argue for $k_{o/u}k_{u/o}$ without loss of generality. Fix k as above and put γ into standard position. If in standard position γ is completely contained in the domain of h , then the result is clear, so suppose this is not the case. We assume for contradiction that there is cancellation with $h(k_{o/u}k_{u/o})$ and, without loss of generality, that it is with $h(k_u) = (k+1)_u$, as otherwise we replace γ with $\bar{\gamma}$ and argue identically for k_o instead. Let γ' be the minimal subsegment of γ beginning with $k_{o/u}k_{u/o}$ so that $h(\gamma')$ had cancellation with $h(k_u)$. The only way cancellation in an unreduced code can involve $h(k_u)$ is if a subsegment of the form $(k+1)_o(k+1)_u\delta(k+1)_u$ appears in an unreduced $h(\gamma')$ where a reduced code for δ is trivial. Since $k \in (-\infty, n_1) \cup [n_2, \infty)$, any $(k+1)_u$ in $h(\gamma')$ must be the image under h of k_u . Therefore, $\gamma' = k_{o/u}k_{u/o}\eta k_u$, where $h(k_u\eta k_u) = (k+1)_u\delta(k+1)_u$, which has trivial reduced code. This contradicts Theorem 4.8 as $k \notin (n_1, n_2)$. \square

The second consequence of Theorem 4.8 shows that characters of a segment that lie in the turbulent region which disagree with the domain of a shift persist in a reduced code of the image of the segment under that shift.

Lemma 4.11 *Let h be a permissible right shift with domain D . Let δ be a simple segment whose support intersects (n_1, n_2) nontrivially, and suppose δ contains a character b with numerical value in (n_1, n_2) which disagrees with ∂D . Then b persists in a reduced code for $h(\delta)$.*

In other words, if $\delta = \delta_1 b \delta_2$, then $h(\delta) = \sigma_1 b \sigma_2$, where $\sigma_1 b = h(\delta_1 b)$ and $b \sigma_2 = h(b \delta_2)$.

Proof Write $b = k_{o/u}$, where by assumption $k \in (n_1, n_2)$.

Claim 4.12 *Since b disagrees with ∂D , any occurrence of b in $h(\delta)$ must also appear in δ .*

Proof Since b disagrees with ∂D , the segment δ cannot be homotoped rel endpoints to be completely contained in D . Thus in standard position, δ can be written as the efficient concatenation of subsegments which are disjoint from D , subsegments of length two which are either full or half crossings, and subsegments supported on $(-\infty, n_1] \cup [n_2, \infty)$ (see Lemma 3.16). We consider the images of each type of subsegment in turn. The subsegments which are disjoint from D are fixed by h . The image of subsegments supported on $(-\infty, n_1] \cup [n_2, \infty)$ have empty intersection with (n_1, n_2) , and so cannot contain b . The remaining subsegments are half crossings, full crossings, or segments that include a character whose numerical value is n_1 . Any subsegment of the image of any of these subsegments supported on (n_1, n_2) agrees with ∂D and so cannot contain b . This proves the claim. \square

Now suppose towards a contradiction that b does not persist in $h(\delta)$. Then there must be another instance of b in an unreduced code for $h(\delta)$ which cancels with b . By the claim, δ must contain a subsegment of the

form $b\sigma b$ whose image is trivial. Theorem 4.8 then implies that $b\sigma b = \overline{\partial D|_{(n_1, k]}}(n_1)_{o/u}(n_1)_{u/o}\partial D|_{(n_1, k]}$. However, $b = k_{o/u}$ disagrees with ∂D by assumption, and therefore $b\sigma b$ cannot have this form, which is a contradiction. \square

5 Constructing the homeomorphisms

As in Section 3, let S be the biinfinite flute with an isolated puncture p , and let $\{B_i \mid i \in \mathbb{Z} \cup P\}$ be the simple closed curves from Definition 3.1 bounding a collection of punctures $\{p_i \mid i \in \mathbb{Z} \cup P\}$, where $p_P = p$. As in that section, we move the punctures $\{p_i \mid i \in \mathbb{Z} \cup P\}$ to the front of S and all other punctures to the back of S .

In this section, we define a countable collection of elements $\{g_n\}_{n \in \mathbb{N}}$ in $\text{Map}(S, p)$. In the following sections we will show that these elements are of intrinsically infinite type and are loxodromic with respect to the action of $\text{Map}(S, p)$ on $\mathcal{A}(S, p)$.

Definition 5.1 For each $n \in \mathbb{N}$, define

$$(3) \quad g_n := h_3^{(n)} \circ h_2^{(n)} \circ h_1,$$




Figure 27: The various shift maps used in the construction of $g_n = h_3^{(n)} \circ h_2^{(n)} \circ h_1$. Top: the shift map h_1 . Bottom left: the shift map $h_2^{(n)}$ for various n . In order from top to bottom, the first two pictures are the case of $n = 1, 2$ (respectively) and the bottom two are for general n odd and general n even (respectively). Bottom right: the shift map $h_3^{(n)}$ for various n . In order from top to bottom, the first two pictures are the case of $n = 1, 2$ and the bottom picture is the case of general n .

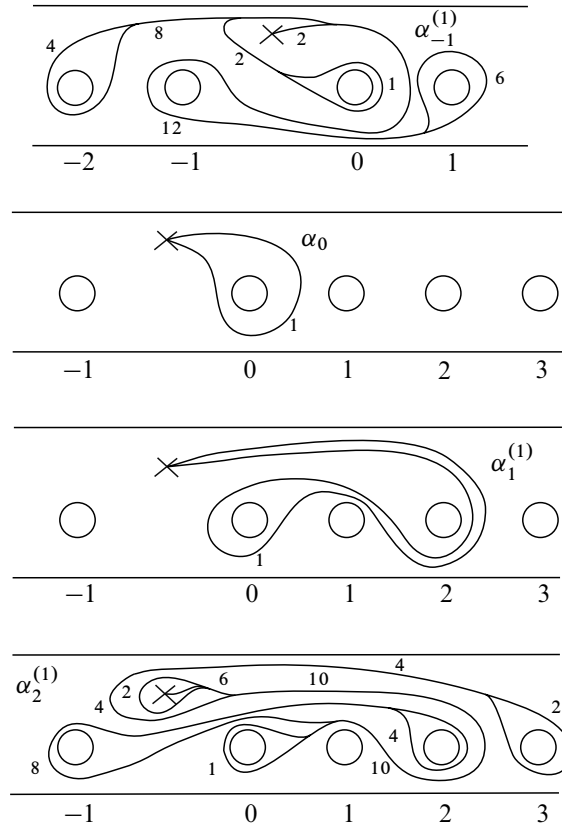


Figure 28: The arcs $\alpha_i^{(1)}$ for $i = -1, 0, 1, 2$ drawn as train tracks for simplicity.

where h_1 and $h_i^{(n)}$ for $i = 2, 3$ are permissible shifts defined as follows. Let h_1 be the right shift whose domain includes all simple closed curves B_j for $j \in \mathbb{Z} \setminus \{0\}$ and passes under p and B_0 . Let $h_2^{(n)}$ be the left shift whose domain includes only $\{B_j \mid j \in (-\infty, -n-1] \cup [n+1, \infty)\}$, passes over p and B_0 , and, when $j \in [-n, -1] \cup [1, n]$, passes under B_j for j odd and over B_j for j even. Finally, let $h_3^{(n)}$ be the right shift whose domain includes only $\{B_j \mid j \in (-\infty, -n-1] \cup [n+1, \infty)\}$ and passes under p and over B_j for all $j \in [-n, P) \cup (P, n]$. See Figure 27.

In the language of the previous sections, we have that $n_1 = -1$ and $n_2 = 1$ for h_1 . For $h_2^{(n)}$, $n_1 = n+1$ and $n_2 = -n-1$. For $h_3^{(n)}$, $n_1 = -n-1$ and $n_2 = n+1$.

In order to prove that each g_n is loxodromic with respect to the action of $\text{Map}(S, p)$ on $\mathcal{A}(S, p)$, we introduce a collection of arcs $\{\alpha_i^{(n)} \mid i \in \mathbb{Z}\}$ which are invariant under g_n . We will show in Proposition 8.2 that for $i \geq 0$, these arcs form a quasigeodesic half-axis for g_n in $\mathcal{A}(S, p)$.

Definition 5.2 For each fixed n , we define a sequence of arcs $\{\alpha_i^{(n)}\}_{i=-\infty}^{\infty}$ in $\mathcal{A}(S, p)$ as the images under successive applications of g_n of the fixed initial arc $\alpha_0 = P_s 0_o 0_u P_s$. That is to say that $\alpha_i^{(n)} = g_n^i(\alpha_0)$ for any $i \in \mathbb{Z}$. See Figure 28 for the case $n = 1$ and Figure 29 for general $n \geq 1$.

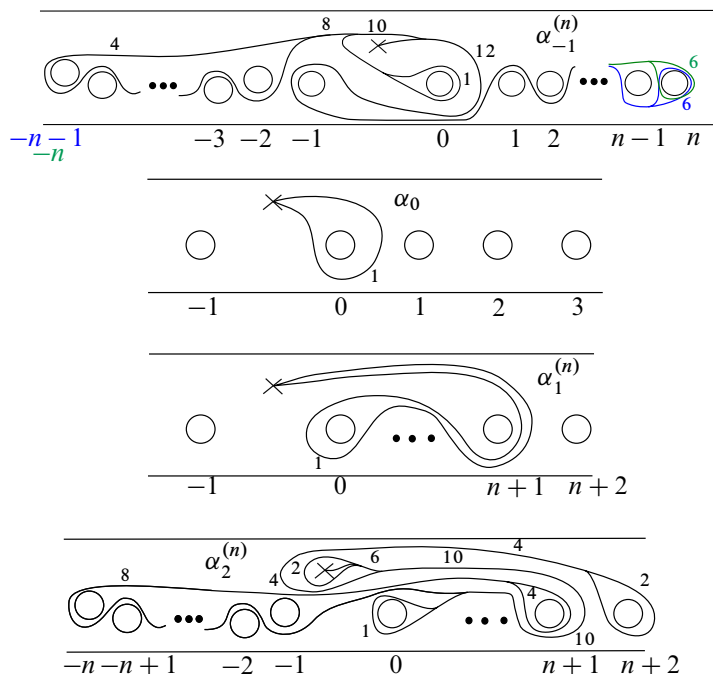


Figure 29: The arcs $\alpha_i^{(n)}$ for $i = -1, 0, 1, 2$ and $n \geq 1$. In $\alpha_{-1}^{(n)}$, the blue appears when n is odd, while the green appears when n is even.

When $n = 1$, the arcs $\alpha_i^{(1)}$ are the most straightforward and thus most useful for building intuition. We suggest that the reader keep these arcs in mind while reading the remainder of the paper. We now give the code for $\alpha_1^{(1)}$ and $\alpha_2^{(1)}$, which the reader can compare to Figure 28:

$$\begin{aligned} \alpha_1^{(1)} &= P_s 0_o 1_o 2_o 2_u 1_o 0_o 0_u 1_o 2_u 2_o 1_o 0_o P_s, \\ \alpha_2^{(1)} &= P_s 0_o 1_o 2_o 2_u 1_o 0_o P_u (-1)_u (-1)_o P_u 0_o 1_o 2_o 2_u 1_o 0_o P_u (-1)_o (-1)_u P_u 0_o 1_o 2_u 2_o 1_o \\ &\quad 0_o P_u P_o 0_o 1_o 2_o 3_u 3_u 2_o 1_o 0_o P_o P_u 0_o 1_o 2_o 2_u 1_o 0_o P_u (-1)_u (-1)_o P_u 0_o 1_o 2_u 2_o 1_o 0_o \\ &\quad P_u (-1)_o (-1)_u P_u 0_o 1_o 2_u 2_o 1_o 0_o P_u P_o 0_o 1_o 2_o 2_u 1_o 0_o 0_u 1_o 2_u 2_o 1_o 0_o P_o P_u 0_o 1_o 2_o \\ &\quad 2_u 1_o 0_o P_u (-1)_u (-1)_o P_u 0_o 1_o 2_o 2_u 1_o 0_o P_u (-1)_o (-1)_u P_u 0_o 1_o 2_u 2_o 1_o 0_o P_u P_o 0_o \\ &\quad 1_o 2_o 3_u 3_u 2_o 1_o 0_o P_o P_u 0_o 1_o 2_o 2_u 1_o 0_o P_u (-1)_u (-1)_o P_u 0_o 1_o 2_u 2_o 1_o 0_o P_u (-1)_o (-1)_u \\ &\quad P_u 0_o 1_o 2_u 2_o 1_o 0_o P_s. \end{aligned}$$

While it is possible to compute the images of arcs under $g_n = h_3^{(n)} \circ h_2^{(n)} \circ h_1$ by hand, we have also written a computer program to implement this. The interested reader should contact the authors for more details.

Remark 5.3 In light of how complicated the arcs $\alpha_i^{(n)}$ are when $i \neq 0, 1$, it may be surprising that only single loops can have trivial image under each of the shifts $h_1, h_2^{(n)}$, and $h_3^{(n)}$ (see Theorem 4.8 and Corollary 4.9). Figure 30 gives the intermediate steps so that one can see how the image under g_1 of a complicated arc such as $\alpha_{-1}^{(1)}$ becomes the straightforward arc α_0 . The figure shows the collection of single loops with trivial image at each stage of computing $g_1(\alpha_{-1}^{(1)}) = \alpha_0$. A similar phenomenon happens for $n > 1$.




Figure 30: Each of the strands in blue are single loops whose image is trivial under the shift whose domain is shown in red. Each blue strand fits the form of Theorem 4.8.

Our method for using the arcs $\alpha_i^{(n)}$ to prove that g_n is a loxodromic isometry of $\mathcal{A}(S, p)$ is inspired by the foundational work of Bavard in [5]. Bavard constructs an element f of the mapping class group of the plane minus a Cantor set which is loxodromic with respect to the associated relative arc graph. We give a brief outline of Bavard's methods here.

Bavard constructs a collection of simple paths β_i which start at an isolated puncture and end at some point of the Cantor set. These paths are invariant under the action of a chosen homeomorphism f and have the property that $\beta_{i+1} = f(\beta_i)$. They are constructed so that for all i , the path β_{i+1} begins by following the same path as β_i . Roughly, if a path γ begins by following the same path as β_i then Bavard says the path γ “begins like” β_i (see [5, Section 2]). Bavard uses this definition to define a function from the vertex set of a certain graph of paths (defined similarly to the relative arc graph but with paths instead of arcs) to $\mathbb{Z}_{\geq 0}$ by sending a path γ to the maximal $i \in \mathbb{Z}_{\geq 0}$ so that γ begins like β_i . Aramayona, Fossas, and Parlier [2] show that the relative arc graph is quasi-isometric to the graph whose vertex set is isotopy class

of paths with *at least one* endpoint on the distinguished isolated puncture. Therefore this function can then be used to estimate distances in the relative arc graph and ultimately to show that the collection of paths $\{\beta_i \mid i \in \mathbb{Z}_{\geq 0}\}$ forms a geodesic half-axis for the element f . The key fact that Bavard uses is that if δ is a path which begins like β_i and γ is any path disjoint from δ , then γ begins like β_{i-1} [5, Lemma 2.4].

Our arcs $\{\alpha_i^{(n)}\}_{i \in \mathbb{Z}}$ do not satisfy the same property as Bavard's paths. One notable difference is that our arcs start and end at the puncture p . Because of this, for a fixed n , $\alpha_{i+1}^{(n)}$ does not begin by following the *entirety* of $\alpha_i^{(n)}$. However, we will show that it does begin by following the *first half* of $\alpha_i^{(n)}$. In light of this, we modify Bavard's notion of "begins like" and make the following definition.

Definition 5.4 Fix an $n \in \mathbb{N}$. An arc δ *starts like* $\alpha_i^{(n)}$ if the maximal initial or the maximal terminal segment of δ which agrees with an initial or terminal segment of $\alpha_i^{(n)}$ (not necessarily respectively) has code length at least $\lfloor \frac{1}{2} \ell_c(\alpha_i^{(n)}) \rfloor$. Recall that code length was defined in Definition 3.6.

In Section 6, we investigate the properties of arcs that start like $\alpha_i^{(n)}$ for some i , and in Section 7 we use these properties to prove Theorem 7.10, which is an analog of [5, Lemma 2.4]. In our more general context, proving this result is quite a bit more involved than the proof of [5, Lemma 2.4]. One reason for this is that the behavior of our arcs $\alpha_i^{(n)}$ is much more complicated than the paths from [5].

6 Arcs that start like $\alpha_i^{(n)}$

The main goal of this section is to prove that the images of arcs which start like $\alpha_i^{(n)}$ under g_n must start like $\alpha_{i+1}^{(n)}$ (see Corollary 6.9). This will follow from Proposition 6.2 and Theorem 6.4. Proposition 6.2 states that the image of the first half of $\alpha_i^{(n)}$ is the first half of $\alpha_{i+1}^{(n)}$. We need to be careful about what we mean by this because of the possibility that we cause there to be a gap when we break $\alpha_i^{(n)}$ into its first and second half (see Section 3.3). Given a segment γ , let $\overset{\circ}{\gamma}$ be the initial subsegment of γ with code length $\lfloor \frac{1}{2} \ell_c(\gamma) \rfloor$. For a fixed i and n , we have $\alpha_i^{(n)} = \overset{\circ}{\alpha}_i^{(n)} + \alpha'$ for some α' . As in Section 3.3, it is possible that when we put $\overset{\circ}{\alpha}_i^{(n)}$ and α' (or their images) into standard position individually with respect to one of the shifts $h_1, h_2^{(n)}$, or $h_3^{(n)}$, we will lose information when we take its image. In fact, this does happen when we put $h_2^{(n)}(h_1(\overset{\circ}{\alpha}_i^{(n)}))$ and $h_2^{(n)}(h_1(\alpha'))$ into standard position with respect to $h_3^{(n)}$ (see Figure 31). We avoid this issue by extending $\overset{\circ}{\alpha}_i^{(n)}$ by one character and considering $\overset{\circ}{\alpha}_i^{(n)} 0_u$. Note that 0_u is fixed by g_n . Thus when we say that the image of the first half of $\alpha_i^{(n)}$ is the first half of $\alpha_{i+1}^{(n)}$ we precisely mean that $g_n(\overset{\circ}{\alpha}_i^{(n)} 0_u) = \overset{\circ}{\alpha}_{i+1}^{(n)} 0_u$ as a subsegment of $g_n(\alpha_i^{(n)}) = \alpha_{i+1}^{(n)}$, in a reduced code. In other words, g_n sends everything before the central 0_u in $\alpha_i^{(n)}$ to $\overset{\circ}{\alpha}_{i+1}^{(n)}$.

Remark 6.1 When dealing with the arcs $\alpha_i^{(n)}$, the floor function in $\overset{\circ}{\alpha}_i^{(n)}$ is actually unnecessary, as $\ell_c(\alpha_i^{(n)})$ will always be even with the central two characters being $0_o 0_u$. However, at this point in the paper we have not proven this fact, which is a consequence of Proposition 6.2, so we will not use it in what follows.

Proposition 6.2 For any n and any i , we have $g_n(\overset{\circ}{\alpha}_i^{(n)} 0_u) = \overset{\circ}{\alpha}_{i+1}^{(n)} 0_u$ in a reduced code.




Figure 31: Here we consider $h_1(\tilde{\alpha}_1 0_u) = h_1(\tilde{\alpha}_1) + h_1(0_o 0_u)$ and take the images of each subsegment separately. The images are drawn as train tracks for clarity, but we omit the weights as they are immaterial. When we put the images $h_2^{(n)}(h_1(\tilde{\alpha}_1))$ and $h_2^{(n)}(h_1(0_o 0_u)) = 0_o 0_u$ in standard position individually with respect to $h_3^{(n)}$ (the second picture), there is a gap in the segment between p and B_1 , which would have been a full crossing. This results in a loss of information in the image under $h_3^{(n)}$.

In order to prove Proposition 6.2, we need to understand how to control cascading cancellation for our homeomorphisms g_n (see Section 4.1 for the definition and a discussion of cascading cancellation). The following lemma is almost a direct consequence of Lemma 4.10, the difference being that g_n is not itself a shift.

Lemma 6.3 *For all $k \geq n + 1$, if $k_{o/u}k_{u/o}$ appears in a reduced code for $\alpha_i^{(n)}$, then there is no cancellation involving $g_n(k_{o/u}k_{u/o})$ in an unreduced code for $\alpha_{i+1}^{(n)}$. In particular, such a pair $k_{o/u}k_{u/o}$ will yield $(k + 1)_{o/u}(k + 1)_{u/o}$ in a reduced code for $\alpha_{i+1}^{(n)}$.*

Proof The union of the turbulent regions for h_1 , $h_2^{(n)}$, and $h_3^{(n)}$ is $[-n - 1, n + 1]$. Therefore we may apply Lemma 4.10 three times once we remark that $h_2^{(n)}$ shifts left, $h_1(k_{o/u}k_{u/o}) = (k + 1)_{o/u}(k + 1)_{u/o}$, and $h_2^{(n)}$ fits the hypothesis of Lemma 4.10, as $k + 1 > n + 1$. \square

We can now use this control on cascading cancellation to prove Proposition 6.2.

Proof of Proposition 6.2 As $\alpha_i^{(n)}$ is symmetric, the only way this proposition could be false is if there is cancellation in $g_n(\alpha_i^{(n)})$ involving the image of the central $0_o 0_u$ (which is also $0_o 0_u$) in $\alpha_i^{(n)}$. Equivalently, we need to show that there is no cancellation involving the final $0_o 0_u$ in $g_n(\tilde{\alpha}_i^{(n)} 0_u)$. We begin by considering the case $i = 1$, where

$$(4) \quad \tilde{\alpha}_1^{(n)} 0_u = P_s 0_o 1_o 2_o \dots (n + 1)_o (n + 1)_u n_o (n - 1)_o \dots 2_o 1_o 0_o 0_u,$$

and all of the characters that are not displayed are k_o for the appropriate k . By Lemma 6.3, there is no cancellation in $g_n(\alpha_1^{(n)})$ involving $g_n((n + 1)_o (n + 1)_u)$, and so there can be no cascading cancellation

involving both the final $0_o 0_u$ and the image of any character before $(n+1)_o(n+1)_u$ in (4). Thus it suffices to consider $g_n(\mu 0_u)$, where

$$\mu = (n+1)_o(n+1)_u n_o(n-1)_o \dots 2_o 1_o 0_o.$$

A direct computation yields, in a reduced code,

$$(5) \quad g_n(\mu 0_u) = (n+2)_o(n+2)_u(n+1)_o n_o \dots (n-1)_o n_o \mu 0_u,$$

and from this computation we can see that there is no cancellation in $g_n(\alpha_1^{(n)})$ involving the final characters, $0_o 0_u$. Thus $g_n(\alpha_1^{(n)} 0_u) = \alpha_2^{(n)} 0_u$ in a reduced code.

By (5), we also see that $g_n(\mu 0_u)$ also ends in $\mu 0_u$. Hence, to show that the result holds for any index i , we simply repeatedly apply g_n and Lemma 6.3 as in the previous paragraph. \square

We can deduce from Proposition 6.2 that if an arc γ starts like $\alpha_i^{(n)}$, then an *unreduced code* for its image $g_n(\gamma)$ will start like $\alpha_{i+1}^{(n)}$. We will use the following technical theorem to show that this is true for a *reduced code* for $g_n(\gamma)$, as well.

Theorem 6.4 *Let γ be a simple arc of the form $\gamma = P_s \eta_1 0_o \eta_2 P_s$ in standard position. Assume the following two conditions hold:*

- (1) *The numerical value of η_2^i is at most 0, ie η_2^i is either 0_u or $P_{o/u}$.*
- (2) *Either the first two characters of η_2 are not a loop around P or, if they are, the initial segment of η_2 is given by $P_u P_o 0_o 1_o$.*

Then there is no cancellation with the initial 0_o in a reduced code for $g_n(0_o \eta_2 P_s)$.

We note that the first condition means that the segment 0_o is oriented to the left.

Remark 6.5 Theorem 6.4 is written with a particular orientation in mind, but such an orientation is arbitrary. That is to say, the exact same statement is true applied to the image of $\bar{\gamma}$ under g_n . We will use both the original and this “reverse” version of Theorem 6.4 later in the paper, so we make its statement precise: Suppose, as in Theorem 6.4, that $\gamma = P_s \eta_1 0_o \eta_2 P_s$ and

- (1) the numerical value of η_1^t is at most 0;
- (2) either the final two characters of η_1 are not a loop around P or the final segment of η_1 is given by $1_o 0_o P_o P_u$.

Then there is no cancellation with the terminal 0_o in a reduced code for $g_n(P_s \eta_1 0_o)$.




Figure 32: Situation in the proof of Lemma 6.6.

We will prove Theorem 6.4 through a series of lemmas. Fix an arc γ which satisfies the hypothesis of Theorem 6.4. We must show that there is no cancellation with the initial 0_o in a reduced code for $g_n(0_o\eta_2P_s)$. For contradiction, suppose there is cancellation with the initial 0_o in a reduced code for $g_n(0_o\eta_2P_s)$. Consider the subsegment of γ given by $0_o\delta$, where $0_o\delta$ is the minimal subsegment such that $g_n(0_o\delta)$ has cancellation with 0_o .

Lemma 6.6 *The segment δ does not contain any back loops.*

Proof Suppose towards a contradiction that δ contains at least one back loop and write δ in standard form as

$$\delta = \tau_1 C_1^- C_1^+ \tau_2 \dots \tau_s C_s^- C_s^+ \tau_{s+1},$$

where the τ_i are possibly empty for $i \leq s+1$. Then

$$g_n(\delta) = g_n(\tau_1 C_1^-) C_1 g_n(C_1^+ \tau_2 C_2^-) C_2 \dots g_n(C_{s-1}^+ \tau_s C_s^-) C_s g_n(C_s^+ \tau_{s+1}).$$

In order to have cascading cancellation involving 0_o , there must be an instance of 0_o in $g_n(C_s^+ \tau_{s+1})$ and the initial subsegment of $g_n(\delta)$ which ends immediately before that 0_o must have trivial image. By the minimality of our choice of δ , this initial subsegment must contain $g_n(C_i) = C_i$ for all $i = 1, \dots, s$. Since its image is trivial, there must be cancellation involving each C_i . The only way that this can occur is if there is some j such that $C_j = \bar{C}_{j+1}$ and $C_j g_n(C_j^+ \tau_j C_{j+1}^-) \bar{C}_j = \emptyset$. Consequently, we must have $g_n(C_j^+ \tau_j C_{j+1}^-) = \emptyset$. This implies that exactly one of the following holds:

- (1) $h_1(C_j^+ \tau_j C_{j+1}^-) = \emptyset$.
- (2) $h_1(C_j^+ \tau_j C_{j+1}^-) \neq \emptyset$ in a reduced code while $h_2^{(n)}(h_1(C_j^+ \tau_j C_{j+1}^-)) = \emptyset$.
- (3) $h_2^{(n)}(h_1(C_j^+ \tau_j C_{j+1}^-)) \neq \emptyset$ in a reduced code while $h_3^{(n)}(h_2^{(n)}(h_1(C_j^+ \tau_j C_{j+1}^-))) = \emptyset$.

Let i correspond to which of the above cases we are in and let $h = h_i^{(n)}$. Define

$$\tau'_j = \begin{cases} C_j^+ \tau_j C_{j+1}^-, & h = h_1, \\ h_1(C_j^+ \tau_j C_{j+1}^-), & h = h_2^{(n)}, \\ h_2^{(n)}(h_1(C_j^+ \tau_j C_{j+1}^-)), & h = h_3^{(n)}. \end{cases}$$

By assumption, $\tau'_j \neq \emptyset$ and $h(\tau'_j) = \emptyset$.

By Theorem 4.8, this implies that τ'_j is a loop which begins and ends in the turbulent region and does not fully cross D in (n_1, n_2) . Moreover, in δ , the path τ'_j is preceded by C_j and followed by $C_{j+1} = \bar{C}_j$ and so one of the back loop connectors C_j^+ or C_{j+1}^- must fully cross D . However this is a contradiction, since these back loop connectors must occur in (n_1, n_2) . See Figure 32. \square

Notice that 0_o does not appear in the image under g_n of any segment supported on

$$(-\infty, -n-1) \cup [n+1, \infty).$$

Therefore, if δ contains any subsegment supported in this region then, by minimality, it must also contain a subsegment of the form $q_{o/u}q_{u/o}$ for $q \geq n+1$ or $q < -n-1$; in other words, the segment δ must “turn around” in the shift region. Moreover, by Lemma 6.6, δ can have no back loops, and hence Lemma 4.10 implies that such a $q_{o/u}q_{u/o}$ blocks cascading cancellation, so this cannot occur. Thus, we conclude that a reduced code for $0_o\delta$ is supported on $[-n-1, n+1]$.

We will now show that there is no cancellation involving 0_o under each of h_1 , $h_2^{(n)}$, and $h_3^{(n)}$. By the assumptions of the theorem, the 0_o in $0_o\delta$ is oriented to the left. If there is no cancellation involving 0_o under a shift, then 0_o persists in the image and is still oriented to the left. This implies that the character after 0_o in the image still satisfies condition (1) of the theorem, that is, it is either 0_u or $P_{o/u}$.

A straightforward calculation shows that $h_1(0_o\delta)$ cannot cancel 0_o for any δ . This follows from the fact that ∂D_1 contains 0_u , not 0_o , where D_1 is the domain for h_1 . Write $0_o\delta'$ for the reduced code of $h_1(0_o\delta)$, which is another segment completely contained in the region $(-n-1, n+1]$ with no back loops. The first character of δ' is still either 0_u or $P_{o/u}$.

It remains to rule out any cancellation involving 0_o under $h_2^{(n)}$ and $h_3^{(n)}$. We do this in the following two lemmas.

Lemma 6.7 *There is no cancellation involving 0_o under $h_2^{(n)}$.*

Proof For contradiction, assume that there is cancellation involving 0_o , and let ζ be the minimal subsegment of $0_o\delta'$ which has cancellation with 0_o under $h_2^{(n)}$.

If $\zeta^t = 0_o$, then ζ is a loop whose image is trivial. By Theorem 4.8, ζ must be of the form

$$\partial D_2|_{[0, n_1)}(n_1)_{o/u}(n_1)_{u/o}\overline{\partial D_2|_{[0, n_1)}},$$

where D_2 is the domain of $h_2^{(n)}$. In particular, $\zeta^i = 0_o$ and this 0_o is oriented to the right. However, as noted above, the initial 0_o of $0_o\delta$ is oriented to the left, which contradicts that ζ is an initial subsegment of $0_o\delta'$.

Therefore, $\zeta = 0_o\zeta'a_1a_2$, where the 0_o which cancels with $\zeta^i = 0_o$ appears in $h_2^{(n)}(a_1a_2)$ and is not the terminal character. This is the case exactly when a_1a_2 is either a full or half crossing. Moreover, in order for 0_o to appear in the image of a_1a_2 , we must have that the numerical values of a_1 and a_2 are greater than or equal to zero, since $h_2^{(n)}$ shifts to the left.

There are two cases to consider: either $(\zeta')^i = 0_u$ or $(\zeta')^i = P_{o/u}$. If $(\zeta')^i = 0_u$, then ζ' contains a character which disagrees with ∂D_2 , namely 0_u . On the other hand, if $(\zeta')^i = P_{o/u}$, then it must be oriented to the left. Since the numerical values of a_1 and a_2 are at least 0 and $P < 0$, ζ' must “turn around” and contain an over-under loop in turbulent region. In particular, it must again contain a character which disagrees with ∂D_2 . In either case, call the character which disagrees with the boundary of the domain b .

By Lemma 4.11, b persists in a reduced code for $h_2^{(n)}(\zeta')$ and such a b must block any cascading cancellation. However, by assumption, there is cancellation between two instances of 0_o in an *unreduced* code for $h_2^{(n)}(\zeta')$, one which precedes b and one which succeeds b . This is a contradiction. \square

It remains only to check for cancellation under $h_3^{(n)}$.

Lemma 6.8 *There is no cancellation involving 0_o under $h_3^{(n)}$.*

Proof The argument is similar to the proof of Lemma 6.7. For contradiction assume that there is cancellation involving 0_o and write $\beta = 0_o\delta''$ for the minimal subsegment of a reduced code for $h_2^{(n)}(0_o\delta')$ which has cancellation with 0_o under $h_3^{(n)}$. Note that $(\delta'')^i$ is still either 0_u or $P_{o/u}$. By the same reasoning as in the proof of Lemma 6.7, $\beta = 0_o\beta'c_1c_2$, and c_1c_2 is a full or half crossing. Recall that for the shift $h_3^{(n)}$, $n_1 = -n - 1$, $n_2 = n + 1$, and $h_3^{(n)}$ shifts to the right.

If $\beta'c_1$ contains a character which does not agree with ∂D_3 , then the contradiction follows by applying Lemma 4.11 as in the proof of Lemma 6.7. However, it is possible that $\beta'c_1$ does not contain such a character. This occurs exactly when $\beta = \overline{\partial D_3|_{(n_1, 0]}}(n_1)_{o/u}$. This possibility did not arise when considering h_2 because h_2 is a left shift, and so $n_1 > 0$.

In this case, β must begin with $0_oP_u(-1)_o$. We will now show that β could not be the image of $0_o\delta$ under $h_2^{(n)} \circ h_1$. In the notation above, $(-1)_o$ is in the image of $\xi = 0_o\delta'$ under $h_2^{(n)}$. Since $(-1)_o$ disagrees with ∂D_2 , there must be an instance of $(-1)_o$ in δ' by Claim 4.12. Let $\xi' = 0_o\sigma(-1)_o$ be the subsegment of ξ so that $h_2^{(n)}(\xi') = 0_oP_u(-1)_o$. We can find ξ directly by computing $(h_2^{(n)})^{-1}(0_oP_u(-1)_o)$. Recall that we cannot always take the preimage of a segment under a permissible shift (see the discussion in Section 3.4). However, in this case we are able to take the preimage precisely because Claim 4.12 ensures that $0_oP_u(-1)_o$ is the image of a subsegment of $h_2^{(n)}(\xi)$. Computing $(h_2^{(n)})^{-1}(0_oP_u(-1)_o)$ results in

$$\xi' = 0_oP_u\partial D_2|_{[P, n+1)}(n+1)_u(n+1)_o\overline{\partial D_2|_{[P, n+1)}}(-1)_o.$$

In particular, ξ' starts with $0_oP_uP_o0_o1_u$. To conclude the proof, we will show that ξ' cannot occur as the image under h_1 of the segment $0_o\delta$.

Assume for contradiction that $h_1(0_o\delta) = \xi'$. Since P_o does not agree with ∂D_1 , there must be an instance of P_o in δ by Claim 4.12. Let $0_o\tau P_o$ be the initial subsegment of $0_o\delta$ so that $h_1(0_o\tau P_o) = 0_oP_uP_o$. As above, we may directly compute the preimage of $0_oP_uP_o$ under h_1 . This shows that $0_o\tau P_o = 0_oP_uP_o$. The assumptions of the theorem then imply that $0_o\delta$ starts with $0_oP_uP_o0_o1_o$, where 1_o is oriented to

the right. Since ξ' starts with $0_o P_u P_o 1_u$, the only way $h_1(0_o \delta) = \xi'$ is if there is cancellation with the terminal 1_o in $0_o P_u P_o 0_o 1_o$ when we apply h_1 to $0_o \delta$. Since h_1 shifts to the right and ξ' contains $(-1)_o$, it must be the case that $0_o \delta$ contains a loop $k_{o/u} k_{u/o}$ with $k \geq 1$. In particular, $0_o \delta$ starts with either $0_o P_u P_o 0_o 1_o 1_u$ (if $k = 1$) or $0_o P_u P_o 0_o 1_o \eta k_{o/u} k_{u/o}$ where η is strictly monotone increasing (if $k > 1$). Since $n_2 = 1$ for h_1 , Lemma 4.10 then implies that this $k_{o/u} k_{u/o}$ blocks cancellation between the terminal 1_o in $0_o P_u P_o 0_o 1_o$ and any characters in $0_o \delta$ that appear *after* $k_o k_u$. Since the segment η of $0_o \delta$ between 1_o and $k_o k_u$ is either empty (if $k = 1$) or a strictly monotone increasing segment (if $k > 1$), using the fact that h_1 shifts to the right, it is straightforward to check that there can be no cancellation with the terminal 1_o in $0_o P_u P_o 1_o$ when we apply h_1 , which is a contradiction. This completes the proof. \square

Lemmas 6.6, 6.7, and 6.8 show that $g_n(0_o \delta)$ has no cancellation with 0_o , which completes the proof of Theorem 6.4.

As promised at the beginning of this section, we have the following corollary.

Corollary 6.9 *For any fixed n and any $i \geq 0$, $\alpha_{i+1}^{(n)}$ starts like $\alpha_i^{(n)}$. More generally, for any $i \geq 1$, if γ starts like $\alpha_i^{(n)}$ then $g_n(\gamma)$ starts like $\alpha_{i+1}^{(n)}$.*

Proof First, a direct computation shows that for any fixed n , $\alpha_1^{(n)}$ starts like $\alpha_0^{(n)}$. For all $i \geq 1$, we will prove both statements simultaneously, using strong induction. We will show that for each n and any $i \geq 1$, if γ is a simple arc that starts like α_i , then $g_n(\gamma) = \alpha_{i+1}^{(n)} \eta'$ for some reduced η' which has the form of η_2 in the hypotheses of Theorem 6.4.

For the base case $i = 1$, suppose γ starts like $\alpha_1^{(n)}$ so that we may write

$$\gamma = \alpha_1^{(n)} \eta,$$

in a reduced code. Applying g_n to both sides of this equation, we have, in an *unreduced* code,

$$g_n(\gamma) = g_n(\alpha_1^{(n)}) \eta',$$

for some reduced η' . By Proposition 6.2, $g_n(\alpha_1^{(n)}) = \alpha_2^{(n)}$ in a reduced code. Thus to conclude that $g_n(\gamma)$ starts like $\alpha_2^{(n)}$, we need to show that $g_n(\alpha_1^{(n)}) = \alpha_2^{(n)}$ persists in a reduced code for $g_n(\gamma)$. Since $\alpha_1^{(n)}$ ends with 0_o , it suffices to show that η has the form of η_2 in the hypotheses of Theorem 6.4 so that there is no cancellation with this 0_o . We have $\alpha_1^{(n)} = P_s 0_o 1_o 2_o \dots (n+1)_o (n+1)_u n_o (n-1)_o \dots 2_o 1_o 0_o$, and so the terminal 0_o is oriented to the left. It follows that η satisfies Theorem 6.4(1). Moreover, since γ is simple and starts like $\alpha_1^{(n)}$, if η begins with a loop around p , it must begin with $P_u P_o 0_o 1_o$ (see Figure 33). Thus η satisfies Theorem 6.4(2) and so Theorem 6.4 shows that there is no cancellation in $g_n(\gamma)$ between the final 0_o in $g_n(\alpha_1^{(n)}) = \alpha_2^{(n)}$ and η' . Therefore, $g_n(\gamma)$ starts like $\alpha_2^{(n)}$. In particular, a direct computation shows that $\alpha_2^{(n)}$ starts like $\alpha_1^{(n)}$ and so by setting $\gamma = \alpha_2^{(n)}$, we have shown that $g_n(\alpha_2^{(n)}) = \alpha_3^{(n)}$ starts like $\alpha_2^{(n)}$.

To finish the base case, we will now show that η' has the form of η_2 in the hypotheses of Theorem 6.4. A direct computation shows that $\alpha_2^{(n)}$ ends with $(n+1)_o (n+1)_u n_o (n-1)_o \dots 2_o 1_o 0_o$, and so that η'




Figure 33: The arc $\alpha_1^{(n)}$ is in red; the gray portions may vary depending on n . One possibility for γ and the subsegment η is shown in blue; the gray portions may vary depending on n .

satisfies Theorem 6.4(1). As above, since $g_n(\gamma) = \alpha_2^{(n)}\eta'$ is simple and $\alpha_2^{(n)}$ (and therefore $\alpha_2^{(n)}$) starts like $\alpha_1^{(n)}$, it follows that if η' starts with a loop around p , it must begin with $P_u P_o 0_o 1_o$, so η' satisfies Theorem 6.4(2). This completes the base case.

Now assume the results hold for all $j < i$. For the induction step, suppose that $\gamma = \alpha_i^{(n)}\eta$ for some η which has the form of η_2 in the hypotheses of Theorem 6.4. The argument for this case goes through exactly as in the base case except we need one additional argument to show that η' has the form of η_2 in the hypotheses of Theorem 6.4, where $g_n(\gamma) = \alpha_{i+1}^{(n)}\eta'$ in a reduced code. By the proof of Proposition 6.2, $\alpha_{i+1}^{(n)}$ ends with $\mu = (n+1)_o(n+1)_u n_o(n-1)_o \dots 2_o 1_o 0_o$ and therefore η' satisfies Theorem 6.4(1). To see that η' satisfies Theorem 6.4(2), notice that $\alpha_{i+1}^{(n)}$ starts like $\alpha_i^{(n)}$, which starts like $\alpha_{i-1}^{(n)}$, etc, so that $\alpha_{i+1}^{(n)}$ (and therefore $\alpha_{i+1}^{(n)}$) starts like $\alpha_1^{(n)}$. In particular, γ must start like $\alpha_1^{(n)}$. Since γ is simple, if η' starts with a loop around p , it must begin with $P_u P_o 0_o 1_o$ and Theorem 6.4(2) is satisfied. This completes the induction step, and the result is proved. \square

7 Highways in arcs

In this section, we introduce and examine the prevalence of certain segments of the code for $\alpha_i^{(n)}$ that we call *highways*. The presence of highways forces arcs disjoint from $\alpha_i^{(n)}$ to have very specific initial and terminal subsegments. This will be instrumental in proving Theorem 7.10, which shows that if δ is any arc which starts like $\alpha_i^{(n)}$ and γ is an arc disjoint from δ , then γ starts like $\alpha_{i-1}^{(n)}$, provided i is large enough. In Section 8, we will use Theorem 7.10 to show that the arcs $\alpha_i^{(n)}$ lie on a quasigeodesic in the modified arc graph.

In Section 7.1, we will give general preliminary definitions and results for general arcs, and then in Section 7.2 we will analyze highways in the arcs $\alpha_i^{(n)}$.

7.1 Preliminaries on highways

Recall our convention that all arcs are assumed to be simple and start and end at p .

Definition 7.1 Given an arc $\delta = P_s q_1 \eta q_2 P_s$, where q_1, q_2 are single characters which are not C and η is a segment, we say that δ has *highways* if either $q_1 P_o/u P_u/o q_1$ or $q_2 P_o/u P_u/o q_2$ is a subsegment of δ .




Figure 34: At the top is an arc where $q_1 = q_2$ and the red segment shows that this arc has highways. Below are two arcs where $q_1 \neq q_2$ and the blue segments must appear if the arc has highways. However, this arrangement contradicts the simplicity of each arc (see Lemma 7.2) so neither arc has highways.

The following lemma is an almost immediate corollary of the definition and the fact that δ is simple.

Lemma 7.2 *If δ is an arc that has highways, then*

$$\delta = P_s \tau \dots \bar{\tau} P_s,$$

for some segment τ with $\ell_c(\tau) > 0$. That is to say, the first part of the code for δ always overlaps with the reverse of the last part of the code in at least two characters, one of which is P_s .

Proof It suffices to consider the case when τ is a single character, so that $\tau = a = \bar{\tau}$ for some a . Since $\delta = P_s q_1 \eta q_2 P_s$, if the conclusion does not hold, then $q_1 \neq q_2$. In this case, either the subsegment $P_s q_1$ intersects $q_2 P_o P_u q_2$ or the subsegment $q_2 P_s$ intersects $q_1 P_o P_u q_1$, contradicting the fact that δ is simple. See Figure 34. \square

Recall that in the code for an arc, the character P_s does not correspond to any subsegment. Since the first and last characters of every arc are always P_s , we use the first *two* characters of δ in the statement of the above lemma to ensure that there is an initial subsegment of δ which fellow travels a terminal subsegment.

In the future, we will need to use a refined notion of highways to constrain the beginnings of certain arcs. For this we define a notion of right lane and left lane. See Figure 35 for examples and nonexamples.

Definition 7.3 Given an arc δ that has highways, a subsegment λ of δ is called a *left lane* if one of the following holds:

- (1) $\lambda = P_o P_u \gamma$, where γ does not contain C and γ is maximal with respect to the property that the code for $P_s \gamma$ coincides with the initial $\ell_c(P_s \gamma)$ many characters in δ and $\bar{\gamma} P_s$ coincides with the terminal $\ell_c(P_s \gamma)$ many characters in δ ; or
- (2) $\lambda = \bar{\gamma} P_u P_o$ is the reverse of the segment of the form in (1).




Figure 35: Some examples and nonexamples of left and right lanes are shown in blue. Top left: the blue segment is *not* a left lane since it does not coincide with the terminal subsegment of the arc. Top right: the blue segment contains both left and right lanes. There is a left lane via case (2) of Definition 7.3 and a right lane via case (1) of Definition 7.4. Bottom left: the blue segment contains two left lanes via case (2) of Definition 7.3. Bottom right: the blue segment contains two left lanes via case (1) of Definition 7.3 (notice the orientation on the arc has been changed).

We note that the reason that there are two possibilities for a left lane is that we want the definition to be independent of the orientation of δ .

Definition 7.4 Given an arc δ that has highways, a segment ρ of the code is called a *right lane* if one of the following holds:

- (1) $\rho = P_u P_o \gamma$, where γ does not contain C and γ is maximal with respect to the property that the code $P_s \gamma$ coincides with the initial $\ell_c(P_s \gamma)$ many characters in δ and $\bar{\gamma} P_s$ coincides with the terminal $\ell_c(P_s \gamma)$ many characters in δ ; or
- (2) $\rho = \bar{\gamma} P_o P_u$ is the reverse of the segment from (1).

If δ is a symmetric arc (see Definition 4.4), then in Definitions 7.3 and 7.4 it suffices to check the overlap on just the initial part of the code for δ . However, in the general case, checking the overlap with both the initial and terminal parts of the code for δ is necessary.

Definition 7.5 Let δ be an arc with highways. The *lane length* $L(\xi)$ of a left or right lane ξ of δ is defined to be

$$L(\xi) = \ell_c(\xi) - 1.$$

We denote the collection of all left lanes of δ by \mathcal{L} and similarly of all right lanes of δ by \mathcal{R} .




Figure 36: In blue is (a portion) of an arc δ . The shaded disk around the puncture is D_p . This example has two right lanes (in green), a portion of which are shown, and two left lanes (in pink), which are shown in their entirety. The $P_{o/u}P_{u/o}$ portion of these lanes are contained in D_p while all other strands of δ are disjoint from D_p .

Through the rest of the section, we fix a closed topological disc D_p of sufficiently small radius with center at the puncture p such that D_p has empty intersection with each B_i where $i \in \mathbb{Z}$, contains B_p , and has empty intersection with each separating curve $\{S_i\}$ from Section 3.2.1. Moreover, we will only work with homotopy representatives of δ which have reduced code and the property that any pair $P_{o/u}P_{u/o}$ lies inside of D_p , while any other segments remain outside, except for the two that come from the initial and terminal two characters of δ (see Figure 36). Throughout the section, when we further homotope δ we will only do so relative to D_p , so one can assume that the set $\delta \cap D_p$ is pointwise fixed.

A left or right lane ξ is called *innermost* if every oriented straight line segment with initial point at the puncture and terminal point on the boundary circle of D_p intersects the $P_{o/u}P_{u/o}$ at the beginning (resp. end) of ξ before it intersects any other lane of the same type (left or right). If the oriented line segment does not intersect any lane, then this condition is vacuously satisfied. See Figure 37. In particular, innermost left and right lanes are the lanes which are closest to an initial and terminal subsegments of δ .

We then have the following lemma.

Lemma 7.6 *Let δ be an arc that has highways. Let $\lambda \in \mathcal{L}$ and $\rho \in \mathcal{R}$ denote the innermost left and right lanes of δ , respectively. Then $L(\lambda) \geq L(\lambda')$ for all $\lambda' \in \mathcal{L}$ and $L(\rho) \geq L(\rho')$ for all $\rho' \in \mathcal{R}$. Moreover, writing $\lambda = P_u P_o \beta_l$ or its reverse and $\rho = P_o P_u \beta_r$ or its reverse, then if γ is an arc disjoint from δ , then one of the following holds:*

- (1) γ has initial code $P_s \beta_l$ and terminal code $\overline{P_s \beta_l}$,
- (2) γ has initial code $P_s \beta_r$ and terminal code $\overline{P_s \beta_r}$.




Figure 37: The blue and red segments are the initial portions of innermost left and right lanes, respectively. The dotted gray segments intersect these segments before intersecting the initial or terminal portion of any other lane.




Figure 38: In black is (a portion) of an arc δ . The red disk is D_p . Any arc γ which is disjoint from δ must have initial and terminal segment which follows one of the three purple strands emanating from the puncture and continues to fellow travel δ .

Intuitively, this lemma states that the arc γ must “stay in a lane” of the arc δ ; see Figure 38. Note that we do not require γ to be distinct from δ , so this statement applies to δ as well. Indeed, the arc δ can always be homotoped to be disjoint from itself while retaining the same code.

Proof The hypothesis that δ has highways implies that $\mathcal{L}, \mathcal{R} \neq \emptyset$, and hence an innermost left and right lane exist. We prove the first statement for the innermost left lane λ ; an identical proof works for ρ .

Since λ is innermost, if there exists a left lane $\lambda' \in \mathcal{L}$ such that $L(\lambda') > L(\lambda)$, then we get a contradiction to the simplicity of δ , as λ must intersect either λ' or an initial segment of δ . See Figure 39.

For the final statement, fix an arc γ that is disjoint from δ . It must be the case that the initial and terminal subsegments of γ are each contained in a single connected component of $D_p \setminus \delta$ (not necessarily the same component). Moreover these initial and terminal subsegments must begin at the puncture p and therefore must also fellow travel λ , ρ , and the initial/terminal parts of δ (see Figure 38). Therefore, by the same reasoning as in the previous paragraph, we see that γ satisfies conclusion (1) of the lemma if $\ell_c(\beta_l) \leq \ell_c(\beta_r)$ and satisfies conclusion (2) of the lemma if $\ell_c(\beta_r) \leq \ell_c(\beta_l)$. \square

7.2 Highways for the arcs $\alpha_i^{(n)}$

The main goal of this section is to apply the technology of the previous section to show that any arc disjoint from an arc which starts like $\alpha_i^{(n)}$ starts like $\alpha_{i-1}^{(n)}$, provided i is large enough (see Theorem 7.10).




Figure 39: A schematic of the two left lanes: λ , which is innermost, and λ' , which is not. If λ' fellow travels an initial segment of δ for longer than λ does, then λ must intersect either λ' (pictured) or the initial segment of δ .

To do this, we will show that $\alpha_i^{(n)}$ has highways for all $i \geq 2$ and all n , beginning by proving the result for $i = 2$. For fixed i and n , define $\chi_i^{(n)}$ to be all of $\alpha_i^{(n)}$ except the initial P_s , so that

$$(6) \quad \alpha_i^{(n)} = P_s \chi_i^{(n)}.$$

Proposition 7.7 *The segment $\alpha_2^{(n)}$ contains a subsegment of the form $\bar{\chi}_1^{(n)} P_u P_o \chi_1^{(n)}$ that appears as the final subsegment of $\alpha_2^{(n)}$. In particular, $\alpha_2^{(n)}$ has highways whose lanes contain $\chi_1^{(n)}$.*

Proof We prove this by carefully examining the code for $\alpha_1^{(n)}$, which we assume is in standard position relative to the disk D_p . As in Section 6, we will retain the 0_u which immediately follows $\alpha_1^{(n)}$ when we compute images. We will show that $\bar{\chi}_1^{(n)} P_u P_o \chi_1^{(n)} 0_u$ appears as the final subsegment of $\alpha_2^{(n)} 0_u$, which then implies the statement of the proposition.

By Proposition 6.2, $\alpha_2^{(n)} 0_u = g_n(\alpha_1^{(n)} 0_u)$. Recall from (4) that

$$(7) \quad \alpha_1^{(n)} 0_u = P_s 0_o 1_o \dots (n+1)_o (n+1)_u n_o (n-1)_o \dots 2_o 1_o 0_o 0_u.$$

Defining $\sigma := \emptyset$ if $n = 1$ or $\sigma := n_o (n-1)_o \dots 2_o$ if $n > 1$, a computation shows that

$$\begin{aligned} h_1(\alpha_1^{(n)} 0_u) &= P_s 0_o 1_o 2_o \dots (n+2)_o (n+2)_u (n+1)_o n_o \dots 2_o 1_o 0_o 0_u \\ &= P_s 0_o 1_o 2_o \dots (n+2)_o (n+2)_u (n+1)_o \sigma 1_o 0_o 0_u. \end{aligned}$$

See also Figure 40, top pair. We will show that the image of the $1_o 0_o 0_u$ at the end of this last equation under $h_3^{(n)} \circ h_2^{(n)}$ produces the requisite segment in $\alpha_2^{(n)}$.

A direct computation shows that

$$h_2^{(n)}((n+2)_o (n+2)_u (n+1)_o \sigma 1_o 0_o 0_u) = (n+1)_o (n+1)_u \sigma h_2^{(n)}(1_o 0_o 0_u),$$

since all of σ is disjoint from D_2 . In standard position, $1_o 0_o 0_u$ contains a full crossing, and we compute that

$$h_2^{(n)}(1_o 0_o 0_u) = 1_o \overline{\partial D_2|_{(-n-1,0]}} (-n-1)_o (-n-1)_u \partial D_2|_{(-n-1,0]} 0_o 0_u,$$

in an unreduced code. We thus have the decomposition

$$\begin{aligned} (8) \quad h_2^{(n)}((n+2)_o (n+2)_u \sigma 1_o 0_o 0_u) &= (n+1)_o (n+1)_u \sigma 1_o \overline{\partial D_2|_{(-n-1,0]}} (-n-1)_o (-n-1)_u \partial D_2|_{(-n-1,0]} 0_o 0_u \\ &= \tau_1 P_o (-1)_u \tau_2 (-n-1)_o (-n-1)_u \tau_3 (-1)_u P_o 0_u, \end{aligned}$$

where each τ_i is defined by the second equality. See Figure 40, middle pair. None of τ_1, τ_2, τ_3 fully cross D_3 in standard position. As $(-1)_u P_o 0_u$ fully crosses D_3 twice, we compute that

$$\begin{aligned} (9) \quad h_3^{(n)}((-1)_u P_o 0_u) &= (-1)_u P_u 0_o \partial D_3|_{(0,n]} (n+1)_u (n+1)_o \overline{\partial D_3|_{(0,n]}} 0_o P_u P_o 0_o \partial D_3|_{(0,n]} \\ &\quad (n+1)_o (n+1)_u \overline{\partial D_3|_{(0,n]}} 0_o 0_u \\ &= (-1)_u P_u \bar{\chi}_1^{(n)} P_u P_o \chi_1^{(n)} 0_u. \end{aligned}$$

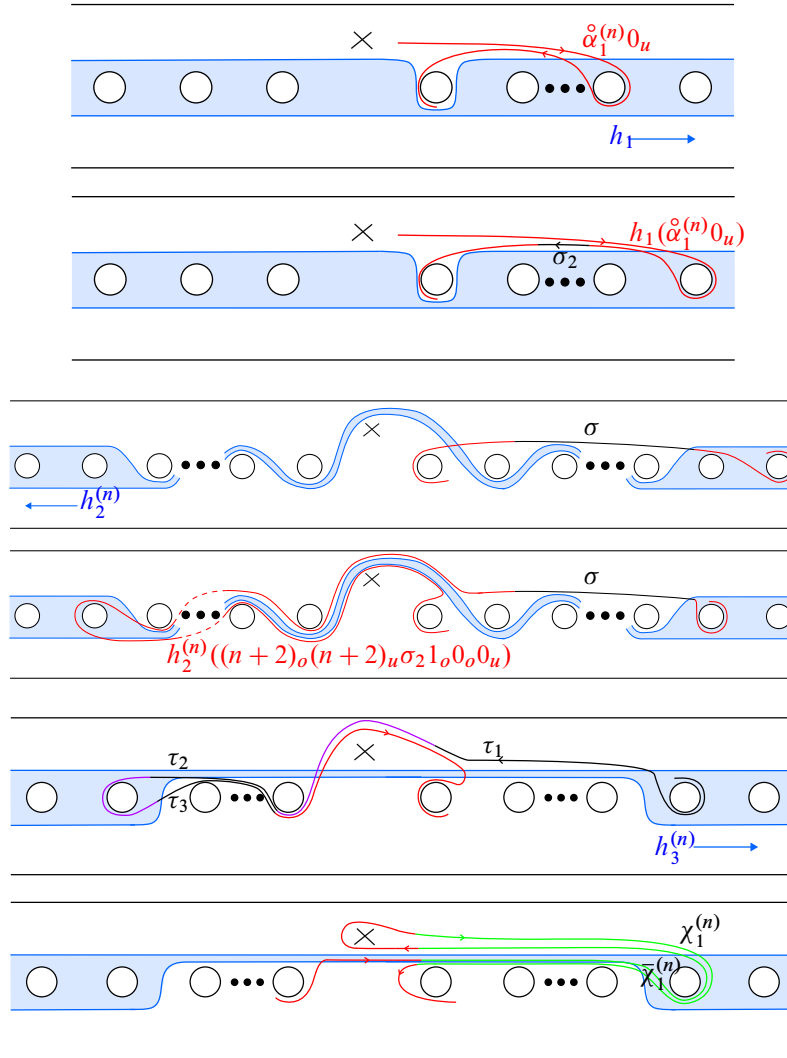


Figure 40: The direct computations done in the proof of Proposition 7.7. Top pair: the image of $\alpha_1^{(n)} 0_u$ under h_1 . Middle pair: the image under $h_2^{(n)}$ of the segment σ , shown in red and black. Bottom pair: the two full crossings in red give rise to $\chi_1^{(n)}$ and $\bar{\chi}_1^{(n)}$, which are shown in green, in the proof of Proposition 7.7 upon applying $h_3^{(n)}$.

Hence we see that $\bar{\chi}_1^{(n)} P_u P_o \chi_1^{(n)} 0_u$ is contained in an *unreduced* code for $g_n(\alpha_1^{(n)} 0_u)$. It remains to show that this segment persists in a *reduced* code for $g_n(\alpha_1^{(n)} 0_u) = \alpha_2^{(n)} 0_u$.

Proposition 6.2 shows that if $\bar{\chi}_1^{(n)} P_u P_o \chi_1^{(n)} 0_u$ persists in a reduced code for $\alpha_2^{(n)} 0_u$ then it persists in a reduced code for α_2 . To check if $\bar{\chi}_1^{(n)} P_u P_o \chi_1^{(n)} 0_u$ persists in a reduced code for $\alpha_2^{(n)} 0_u$, we need only consider a reduced code for the image under g_n of the characters following the $(n+1)_o(n+1)_u$ in (7), because Lemma 6.3 shows that the pair $(n+1)_o(n+1)_u$ blocks cascading cancellation between characters on either side. This is precisely the segment whose image under $h_2^{(n)}$ is given by (8).

We will use direct computation to show that the segment $\bar{\chi}_1^{(n)} P_u P_o \chi_1^{(n)} 0_u$ is the final segment of a reduced code for

$$(10) \quad h_3^{(n)}(\tau_1 P_o(-1)_u \tau_2(-n-1)_o(-n-1)_u \tau_3(-1)_u P_o 0_u).$$

Before making this computation, we point out the relevant pieces. Each $(-1)_u P_o$ (or its reverse) in (10) fully crosses D_3 and has image given by the first half of (9), up to reversing the orientation. The terminal $P_o 0_u$ also fully crosses D_3 , and so we may similarly compute its image. The segments τ_2, τ_3 are invariant under $h_3^{(n)}$. The image of τ_1 is $h_3^{(n)}(\tau_1) = (n+2)_o(n+2)_u(n+1)_o \sigma 1_o 0_o$. Moreover, the four characters $\tau_2^t(-n-1)_o(-n-1)_u \tau_3^i$ form a loop which fully crosses D_3 exactly once. (Note that $\tau_2^t = \tau_3^i = (-n)_{o/u}$, where the choice of o/u depends on the parity of n .) Combining all of these remarks, we find the following *reduced* code (see Figure 40, bottom pair):

$$\begin{aligned} & h_3^{(n)}(\tau_1 P_o(-1)_u \tau_2(-n-1)_o(-n-1)_u \tau_3(-1)_u P_o 0_u) \\ &= (n+2)_o(n+2)_u(n+1)_o \sigma 1_o 0_o P_o P_u 0_o \dots (n+1)_o(n+1)_u n_o \dots 0_o \\ & \quad P_u \overline{\partial D_2|_{(-n,-1]}}(-n)_u(-n)_o \partial D_3|_{(-n,n+1)}(n+1)_o(n+1)_u \overline{\partial D_2|_{(-n,n]}}(-n)_o(-n)_u \partial D_2|_{(-n,-1]} P_u \\ & \quad \partial D_3|_{[0,n+1)}(n+1)_u(n+1)_o \overline{\partial D_3|_{(P,n+1)}} P_u P_o \partial D_3|_{[0,n+1)}(n+1)_o(n+1)_u \overline{\partial D_3|_{[1,n+1)}} 0_o 0_u. \end{aligned}$$

In particular, we see the requisite segment as the last line of this string. Precisely, we have

$$\chi_1^{(n)} = \partial D_3|_{[0,n+1)}(n+1)_o(n+1)_u \overline{\partial D_3|_{[1,n+1)}} 0_o.$$

This completes the proof of Proposition 7.7. \square

The next corollary shows that $\alpha_i^{(n)}$ has highways whose lanes contain $\chi_{i-1}^{(n)}$. Notice that $g_n(\chi_{i-1}^{(n)}) = \chi_i^{(n)}$ by Proposition 6.2.

Corollary 7.8 *When $i \geq 2$, $\alpha_i^{(n)}$ contains a subsegment of the form $\bar{\chi}_{i-1}^{(n)} P_u P_o \chi_{i-1}^{(n)}$ that appears as the final subsegment of $\alpha_i^{(n)}$. In particular, $\alpha_i^{(n)}$ has highways for all $i \geq 2$ with lanes containing $\chi_{i-1}^{(n)}$.*

Proof As above, we will retain the 0_u which immediately follows $\alpha_i^{(n)}$ when we compute images. We will show by induction that $g_n(\alpha_{i-1}^{(n)} 0_u)$ ends with $\bar{\chi}_{i-1}^{(n)} P_u P_o \chi_{i-1}^{(n)} 0_u$ in a reduced code, which will show that $\bar{\chi}_{i-1}^{(n)} P_u P_o \chi_{i-1}^{(n)}$ appears as the final subsegment of $\alpha_i^{(n)}$. Since $g_n(\alpha_{i-1}^{(n)} 0_u) = \alpha_i^{(n)} 0_u$ by Proposition 6.2, this will imply that such a segment persists in $\alpha_i^{(n)}$.

The base case $i = 2$ was shown in Proposition 7.7, so we proceed to the induction step. Using the induction hypothesis, write

$$\alpha_{i-1}^{(n)} 0_u = P_s \eta_1 \bar{\chi}_{i-2}^{(n)} P_u P_o \chi_{i-2}^{(n)} 0_u,$$

for some subsegment η_1 . A calculation shows that in an unreduced code

$$\begin{aligned} g_n(\bar{\chi}_{i-2}^{(n)} P_u P_o \chi_{i-2}^{(n)} 0_u) &= g_n(\bar{\chi}_{i-2}^{(n)} + 0_o P_u P_o 0_o + \chi_{i-2}^{(n)} 0_u) \\ &= g_n(\bar{\chi}_{i-2}^{(n)}) + g_n(0_o P_u P_o 0_o) + g_n(\chi_{i-2}^{(n)} 0_u) \\ &= \bar{\chi}_{i-1}^{(n)} P_u P_o \chi_{i-1}^{(n)} 0_u, \end{aligned}$$

where one verifies that the last line is in fact reduced. We therefore reduce to showing that $\bar{\chi}_{i-1}^{(n)} P_u P_o \chi_{i-1}^{(n)} 0_u$ persists as the terminal substring of $\alpha_i^{(n)} 0_u$. For this, we must prove that no characters in an unreduced code for $g_n(P_s \eta_1 0_o)$ cancel with this terminal 0_o , which is $(\bar{\chi}_{i-1}^{(n)})^i$.

Recall from (6) that $\chi_{i-1}^{(n)}$ is by definition all of $\alpha_{i-1}^{(n)}$ except the initial P_s . In particular, its terminal character $0_o = (\chi_{i-1}^{(n)})^t$ is oriented to the left. Thus the initial 0_o of the reverse segment $\bar{\chi}_{i-1}^{(n)}$ is oriented to the right. Since we are considering the segment $P_s \eta_1 0_o = P_s \eta_1 (\bar{\chi}_{i-1}^{(n)})^i$, this implies that $(\eta_1)^t$ has numerical value at most 0. Moreover, simplicity of $\alpha_{i-1}^{(n)}$ shows that if η_1 ends by looping around p , then it must have $1_o 0_o P_o P_u$ as its final segment. Theorem 6.4 combined with Remark 6.5 then shows that there is no cancellation with $(\bar{\chi}_{i-1}^{(n)})^i = 0_o$ in a reduced code for $g_n(P_s \eta_1 0_o)$. \square

The following corollary follows immediately from Corollaries 6.9 and 7.8, our definition of $\chi_i^{(n)}$ in (6), and the definition of lane length.

Corollary 7.9 *Fix $i \geq 2$ and let \mathcal{L} and \mathcal{R} denote the collections of left and right lanes in $\alpha_i^{(n)}$, respectively. Then*

$$n_L = \max_{\lambda \in \mathcal{L}} L(\lambda) \geq \ell_c(\alpha_{i-1}^{(n)}) \quad \text{and} \quad n_R = \max_{\rho \in \mathcal{R}} L(\rho) \geq \ell_c(\alpha_{i-1}^{(n)}).$$

We are now in a position to prove the main result of the section.

Theorem 7.10 *If δ is an arc which starts like $\alpha_i^{(n)}$ for some $i \geq 2$ and γ is any arc disjoint from δ , then γ starts like $\alpha_{i-1}^{(n)}$.*

Proof Proposition 7.7 and Corollary 7.8 show that $\alpha_i^{(n)}$ contains the segment $\bar{\chi}_{i-1}^{(n)} P_u P_o \chi_{i-1}^{(n)}$ and hence $\alpha_i^{(n)}$ has highways for all $i \geq 2$.

As δ starts like $\alpha_i^{(n)}$, the first part of its code is identical to that of $\alpha_i^{(n)}$ and therefore δ also has highways and also contains the segment $\bar{\chi}_{i-1}^{(n)} P_u P_o \chi_{i-1}^{(n)}$. By Lemma 7.6, the innermost left lane $P_u P_o \alpha_l$ (or its reverse) and the innermost right lane $P_o P_u \alpha_r$ (or its reverse) in δ each have lane length at least $\ell_c(\alpha_{i-1}^{(n)})$. Hence the first $\ell_c(\alpha_{i-1}^{(n)}) - 1$ characters of these lanes immediately following $P_o/u P_u/o$ must agree with $\chi_{i-1}^{(n)}$ or its reverse. As γ is disjoint from δ , the moreover statement of Lemma 7.6 gives that the code for γ must begin with $P_s \alpha_l$ or $P_s \alpha_r$. Consequently it must begin with $\alpha_{i-1}^{(n)} = P_s \chi_{i-1}^{(n)}$, as required. \square

8 Loxodromic elements for $\mathcal{A}(\Sigma, p)$

In this section, we will first conclude the proof of Theorem 1.1 and then go on to explore some properties of our loxodromic elements g_n , including identifying an explicit geodesic axis for g_n , as well as describing the limit points of g_n on the boundary of the modified arc complex $\mathcal{A}(\Sigma, p)$ for an admissible surface Σ .

8.1 The proof of Theorem 1.1

Let S be the biinfinite flute surface we defined in Section 2 with isolated puncture p . Recall that $\mathcal{A}(S, p)$ denotes the modified arc graph of S , as in Definition 2.3. Let $\alpha_0 = P_s 0_o 0_u P_s \in \mathcal{A}(S, p)$ be the arc from Definition 5.2.

For each n , consider the map

$$(11) \quad \varphi_n: \mathcal{A}(S, p) \rightarrow \mathbb{Z}_{\geq 0},$$

defined by setting $\varphi_n(\delta)$ equal to the largest $i \geq 0$ such that δ starts like $\alpha_i^{(n)}$. If there is no such i then set $\varphi_n(\delta) = 0$.

Lemma 8.1 *For any $\gamma, \delta \in \mathcal{A}(S, p)$, we have $d_{\mathcal{A}(S, p)}(\gamma, \delta) \geq |\varphi_n(\gamma) - \varphi_n(\delta)|$.*

Proof It follows from Theorem 7.10 that if $d_{\mathcal{A}(S, p)}(\gamma, \delta) = 1$, then $|\varphi_n(\gamma) - \varphi_n(\delta)| \leq 1$ for any n . In particular, if $\varphi_n(\delta) = \varphi_n(\gamma)$, then we are done. If not, then without loss of generality, assume that $\varphi_n(\delta) < \varphi_n(\gamma)$ and $\varphi_n(\gamma) = j$. Then by Theorem 7.10, $\varphi_n(\delta) = j - 1$ so that $|\varphi_n(\gamma) - \varphi_n(\delta)| \leq 1$. The result then follows from the subadditivity of the absolute value. \square

Let $\{g_n\}_{n \in \mathbb{N}}$ be as in Definition 5.1, and let $\{\alpha_i^{(n)}\}_{i \in \mathbb{Z}}$ be as in Definition 5.2. We first show that the elements $g_n \in \text{Map}(S, p)$ are loxodromic with respect to the action on $\mathcal{A}(S, p)$.

Proposition 8.2 *For each $n \in \mathbb{N}$, the homeomorphism $g_n \in \text{Map}(S, p)$ is a loxodromic isometry of $\mathcal{A}(S, p)$ with a $(2, 0)$ -quasigeodesic axis $\{\alpha_i^{(n)}\}_{i \in \mathbb{Z}}$.*

Proof We first show that the map $\mathbb{Z}_{\geq 0} \rightarrow \langle g_n \rangle \alpha_0 \subset \mathcal{A}(S, p)$ defined by $i \mapsto g_n^i(\alpha_0)$ is a $(2, 0)$ -quasi-isometry. In other words, we will show that $\{\alpha_i^{(n)}\}_{i \geq 0}$ is a quasigeodesic half-axis for g_n along which g_n acts as translation.

Let φ_n be the map defined in (11). By Lemma 8.1, we have that $d_{\mathcal{A}(S, p)}(\gamma, \delta) \geq |\varphi_n(\gamma) - \varphi_n(\delta)|$ for any $\gamma, \delta \in \mathcal{A}(S, p)$. Since $\varphi_n(\alpha_i^{(n)}) = i$ for all $i \geq 0$, this implies that

$$(12) \quad d_{\mathcal{A}(S, p)}(\alpha_0, \alpha_i^{(n)}) \geq i.$$

Consider the arc $\beta = P_s(-1)_o(-1)_u P_s$. Then β is disjoint from both $\alpha_0^{(n)}$ and $\alpha_1^{(n)}$. Since g_n is a homeomorphism, $g_n^i(\beta)$ is disjoint from both $\alpha_i^{(n)}$ and $\alpha_{i+1}^{(n)}$, and thus $d_{\mathcal{A}(S, p)}(\alpha_i^{(n)}, \alpha_{i+1}^{(n)}) \leq 2$ for all i . Therefore, for all i ,

$$(13) \quad d_{\mathcal{A}(S, p)}(\alpha_0, \alpha_i^{(n)}) \leq 2i.$$

Together, (12) and (13) show that $\{\alpha_i^{(n)}\}_{i \geq 0}$ is a $(2, 0)$ -quasigeodesic half-axis for g_n .

Since $\{\alpha_i^{(n)}\}_{i \geq 0}$ is an unbounded orbit of g_n , we can see that g_n is not elliptic, and since g_n acts as translation along this quasigeodesic half-axis, g_n cannot be parabolic. Thus we may conclude that g_n is a loxodromic isometry of $\mathcal{A}(S, p)$. \square

Recall from Definition 2.6 that a surface Σ with an isolated puncture p is admissible if there exists a proper embedding $S \hookrightarrow \Sigma$ where S contains p , the two nonisolated ends of S correspond to distinct ends of Σ , and with the property that either there are countably many connected components of $\Sigma \setminus S$ of the same (nontrivial) topological type or countably many isolated punctures of S remain isolated punctures when embedded in Σ . We denote this special class of connected components by \mathcal{U} , so that the elements of \mathcal{U} are all homeomorphic to a fixed surface Σ_0 . Recall that this definition ensures that the shift maps $h_1, h_2^{(n)}$, and $h_3^{(n)}$ we are interested in on S extend to shift maps on Σ . In particular, $g_n = h_3^{(n)} \circ h_2^{(n)} \circ h_1$ extends to a homeomorphism of Σ . When we reference g_n below, we will try to be specific about when we are considering g_n as an element of $\text{Map}(S, p)$ versus an element of $\text{Map}(\Sigma, p)$. Recall from Definition 2.2 that a homeomorphism $f \in \text{Map}(\Sigma, p)$ is intrinsically infinite-type if $f \notin \overline{\text{Map}_c(\Sigma)}$.

Theorem 1.1 is then a direct consequence of the following theorem.

Theorem 8.3 *Let Σ be an admissible surface. For each $n \in \mathbb{N}$, the homeomorphism $g_n \in \text{Map}(\Sigma, p)$ is a loxodromic isometry of $\mathcal{A}(\Sigma, p)$ with a $(4, 0)$ -quasigeodesic axis $\{\alpha_i^{(n)}\}_{i \in \mathbb{Z}}$. Moreover, g_n is intrinsically infinite-type.*

Proof Fix $n \in \mathbb{N}$. By Proposition 8.2, g_n is loxodromic with respect to the action of $\text{Map}(S, p)$ on $\mathcal{A}(S, p)$. Moreover, g_n extends to an element of $\text{Map}(\Sigma, p)$ and so g_n acts on by isometries $\mathcal{A}(\Sigma, p)$ as well. By Lemma 2.10, there is a $(2, 0)$ -quasi-isometric embedding $\mathcal{A}(S, p) \hookrightarrow \mathcal{A}(\Sigma, p)$. Therefore the image of the $(2, 0)$ -quasigeodesic half-axis for g_n constructed in Proposition 8.2 is a $(4, 0)$ -quasigeodesic half-axis in $\mathcal{A}(\Sigma, p)$. It is clear that g_n stabilizes this $(4, 0)$ -quasigeodesic half-axis and so the arcs $\{\alpha_i^{(n)}\}_{i \in \mathbb{Z}_{\geq 0}}$ form a quasigeodesic half-axis for g_n in $\mathcal{A}(\Sigma, p)$. Therefore, g_n is loxodromic with respect to the action of $\text{Map}(\Sigma, p)$ on $\mathcal{A}(\Sigma, p)$.

We now show that $g_n \notin \overline{\text{Map}_c(\Sigma)}$. If Σ_0 has a nontrivial end space, then $g_n \notin \text{PMap}(\Sigma)$ since g_n translates the elements of \mathcal{U} . Note that $\overline{\text{Map}_c(\Sigma)} < \text{PMap}(\Sigma)$ so that g_n is of intrinsically infinite-type in this case.

Now suppose that the end space of Σ_0 is trivial, so that Σ_0 is a finite-genus surface with one boundary component. Therefore, $S \cup \mathcal{U} = \Sigma'$ is homeomorphic to an infinite-genus surface with two nonplanar ends and a countable number of planar ends. Note that $\Sigma \setminus \Sigma'$ consists of all of the additional topology of Σ that is irrelevant to our construction of shift maps. In this way, the planar ends of Σ' cut away the irrelevant topology of Σ .

By [3, Corollary 6], $\text{PMap}(\Sigma') = \overline{\text{Map}_c(\Sigma')} \rtimes \mathbb{Z}(h')$, where h' is the standard handleshift on Σ' , which also corresponds to a handleshift on Σ by extending h' via the identity on $\Sigma \setminus \Sigma'$. For simplicity of notation, we drop the subscript on g_n in what follows since the argument does not depend on n . Theorem 8.5 below tells us that when g is considered as an element of $\text{Map}(S, p)$, there is a compactly supported mapping class φ such that $g = \varphi h$, where h is the right shift on S shifting the punctures corresponding to the elements of \mathcal{U} . As an element of $\text{PMap}(\Sigma', p)$, and therefore of $\text{Map}(\Sigma, p)$, g is thus equal to $\varphi \cdot (h')^m$, where $\varphi \in \text{Map}_c(\Sigma') < \text{Map}_c(\Sigma)$ and m is the genus of Σ_0 . By the proof of [22, Proposition 6.3], $(h')^m \notin \overline{\text{Map}_c(\Sigma)}$, so that $g \notin \overline{\text{Map}_c(\Sigma)}$ as well. \square

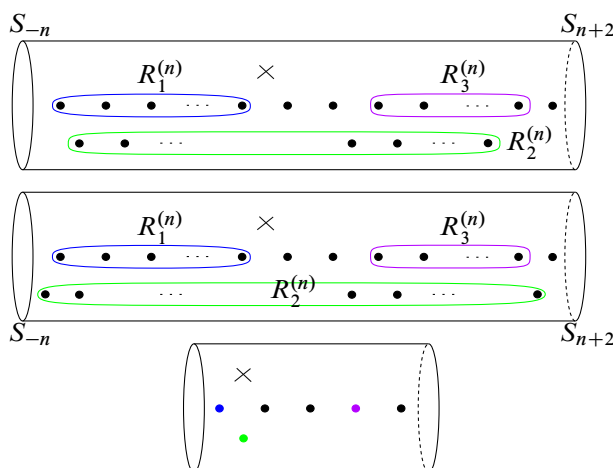


Figure 41: The curves $R_i^{(n)}$ when n is odd and even (top and middle, respectively), and the surface $\Pi = \Pi_n$ when $n \geq 3$ (bottom). In Π , the blue, green, and purple punctures correspond to where Π'_n was cut along the curves $R_1^{(n)}$, $R_2^{(n)}$, $R_3^{(n)}$, respectively.

8.2 Alternative description of g_n

In this section, we show that our mapping classes g_n can be written as the composition of a pseudo-Anosov on a finite-type subsurface and a standard shift. Moreover, we will show that g_n is the composition of *the same* pseudo-Anosov on a *fixed* finite-type subsurface and a standard shift whenever $n > 3$; however, this subsurface is embedded in Σ in different ways for different n , yielding distinct elements of $\text{Map}(\Sigma, p)$. Though the g_n can be expressed as a pseudo-Anosov composed with a shift map, they are not end-periodic maps as defined in the work of Cantwell, Conlon, and Fenley [13]. In particular, there are isolated planar ends, or punctures, that are fixed by our homeomorphisms (for example, the puncture p). However, it is possible that some of the methods used to study end periodic maps could also apply to the homeomorphisms we have constructed.

Let h be the shift map on our subsurface S that translates the punctures that correspond to elements of \mathcal{U} , that is, the right shift whose domain contains the simple closed curves B_i for all $i \in \mathbb{Z}$ and maps B_i to B_{i+1} .

Without loss of generality, for each n we may modify our separating curves S_{-n} and S_{n+1} so that a connected component Π'_n of $S \setminus (S_{-n} \cup S_{n+2})$ is a sphere with 2 boundary components and $n + 4$ punctures, one of which is p . This amounts to pushing any extra topology on the back of S outside of this subsurface. We now further modify Π'_n to form a subsurface Π_n for each n as follows. Let $R_1^{(n)}, R_2^{(n)}, R_3^{(n)}$ be the simple closed curves as shown in Figure 41. In particular,

- $R_1^{(n)}$ encloses all B_i with $-n \leq i \leq -1$ and i odd;
- $R_2^{(n)}$ encloses all B_i with $-n \leq i \leq n$ and i even; and
- $R_3^{(n)}$ encloses all B_i with $3 \leq i \leq n$ and i odd.




Figure 42: The train tracks τ_n on the finite-type subsurface Π_n . Left: $n = 1, 2$. The pink puncture only appears when $n = 2$. Right: $n \geq 3$.

Note that $R_i^{(1)}$ is empty for all $i = 1, 2, 3$, $R_i^{(2)}$ is empty when $i = 1, 3$, and $R_i^{(3)}$ is empty when $i = 3$. For all $n \geq 4$, $R_i^{(n)}$ is nonempty for each $i = 1, 2, 3$.

Definition 8.4 Let Π_n be the component of $\Pi'_n \setminus (R_1^{(n)} \cup R_2^{(n)} \cup R_3^{(n)})$ which contains the puncture p .

The surface Π_n is a sphere with 2 boundary components and some number of punctures: five punctures if $n = 1$, six punctures if $n = 2$, and seven punctures if $n \geq 3$. Notice that the Π_n are homeomorphic for all $n \geq 3$. However, the embedding $\iota_n: \Pi_n \rightarrow \Sigma$ are different for distinct n . For any $f \in \text{PMap}(\Pi_n, p)$, let $\bar{f} := \iota_n \circ f$.

Theorem 8.5 For each $n \geq 1$, there is a pseudo-Anosov $\varphi_n \in \text{PMap}(\Pi_n)$ so that $g_n = \bar{\varphi}_n h$. Moreover, for all $n, n' \geq 3$, $\Pi_n = \Pi_{n'}$ and $\varphi_n = \varphi_{n'}$ as elements of $\text{PMap}(\Pi_n)$. However, $\bar{\varphi}_n$ and $\bar{\varphi}_{n'}$ are distinct elements of $\text{Map}(\Sigma, p)$ since the embeddings ι_n are distinct.

Proof We define $\bar{\varphi}_n := g_n h^{-1}$ for all $n \geq 1$. It is clear that $\bar{\varphi}_n$ stabilizes the subsurface Π_n and is the identity on $\Sigma \setminus \Pi_n$. Let φ_n be the restriction of $\bar{\varphi}_n$ to $\text{PMap}(\Pi_n)$, so that $\iota \circ \varphi_n = \bar{\varphi}_n$. We will show that φ_n is pseudo-Anosov. To do this, we will apply [28, Lemma 3.1], which states that a mapping class is pseudo-Anosov if it preserves a large, generic, birecurrent train track whose associated transition matrix is Perron–Frobenius. We will construct such a train track τ_n for each n .

The cases $n = 1, 2$ are slightly different and we will deal with them separately. For all $n \geq 3$, the surfaces Π_n are homeomorphic and we will build a single train track which will satisfy the above conditions. Notice that τ_n is a train track on the surface Π_n , so to show that τ_n is large, generic, and birecurrent it suffices to consider Π_n . However, since φ_n is defined as a restriction of $\bar{\varphi}_n$, which is a product of elements of $\text{Map}(\Sigma, p)$ which are not supported on Π_n , we must consider the different embeddings of Π_n into Σ in order to show that τ_n is preserved by φ_n and to calculate the transition matrix of τ_n .

The train tracks τ_n for $n = 1, 2$ and for $n \geq 3$ are shown on Π_n in Figure 42. For all n , the following hold. All switches are trivalent and so τ_n is generic. Each complementary region is a once-punctured disk or a polygon and so τ_n is large. The weights on each branch of τ_n are positive and so τ_n is recurrent. Moreover,




Figure 43: A collection of simple closed curves $\{\gamma_i^{(n)}\}$ such that each branch of τ_n is intersected transversely and efficiently by some $\gamma_i^{(n)}$. As in Figure 42, the pink puncture only appears when $n = 2$. Left: $n = 1, 2$. Right: $n \geq 3$.

the finite collections of simple closed curves $\{\gamma_i^{(n)}\}$ in Figure 43 is such that each branch of τ_n is intersected transversely and *efficiently* by some $\gamma_i^{(n)}$, ie $\gamma_i^{(n)} \cup \tau_n$ has no complementary bigon regions for any i . Therefore τ_n is transversely recurrent. Since τ_n is recurrent and transversely recurrent, it is birecurrent.

Figures 44 and 45 show that τ_n is preserved by φ_n for $n = 1, 2$ and $n \geq 3$, respectively. It is immediate from these figures that the matrix A_n associated to τ_n is

$$A_1 = A_2 = \begin{pmatrix} 5 & 6 & 0 & 2 \\ 6 & 9 & 0 & 2 \\ 10 & 10 & 2 & 3 \\ 6 & 6 & 1 & 2 \end{pmatrix} \quad \text{and} \quad A_n = \begin{pmatrix} 5 & 6 & 0 & 2 & 0 \\ 6 & 9 & 0 & 2 & 0 \\ 10 & 10 & 2 & 2 & 1 \\ 6 & 6 & 1 & 1 & 1 \\ 6 & 6 & 1 & 2 & 0 \end{pmatrix},$$

when $n \geq 3$.

For each n , $(A_n)^2$ has all positive entries, hence A_n is Perron–Frobenius. We conclude that φ_n is pseudo-Anosov for all n by [28, Lemma 3.1]. \square

While it is not necessary for this paper, it is interesting to note that for all $n \in \mathbb{N}$, the top eigenvalue of A_n is $\frac{9}{2} + \frac{1}{2}\sqrt{41} + \sqrt{\frac{1}{2}(59 + 9\sqrt{41})}$, which is associated to a unique irrational lamination on Π_n that is carried by τ_n and fixed by φ_n .

We say that a homeomorphism of $\text{Map}(\Sigma, p)$ is a *pseudo-Anosov shift* if it can be written as the composition of a pseudo-Anosov on a finite-type subsurface containing p and a standard shift. The results of this section inspire the following questions.

Question 8.6 When is the composition of shift maps a pseudo-Anosov shift?

Question 8.7 Does every pseudo-Anosov shift act loxodromically on $\mathcal{A}(\Sigma, p)$?

8.3 Geodesic axes

The proof of Theorem 1.1 shows that, for each n , the sequence $(\alpha_i^{(n)})_{i \in \mathbb{Z}}$ is a $(4, 0)$ -quasigeodesic axis for g_n in $\mathcal{A}(\Sigma, p)$. In this section, we find a geodesic axis for g_n in $\mathcal{A}(\Sigma, p)$.

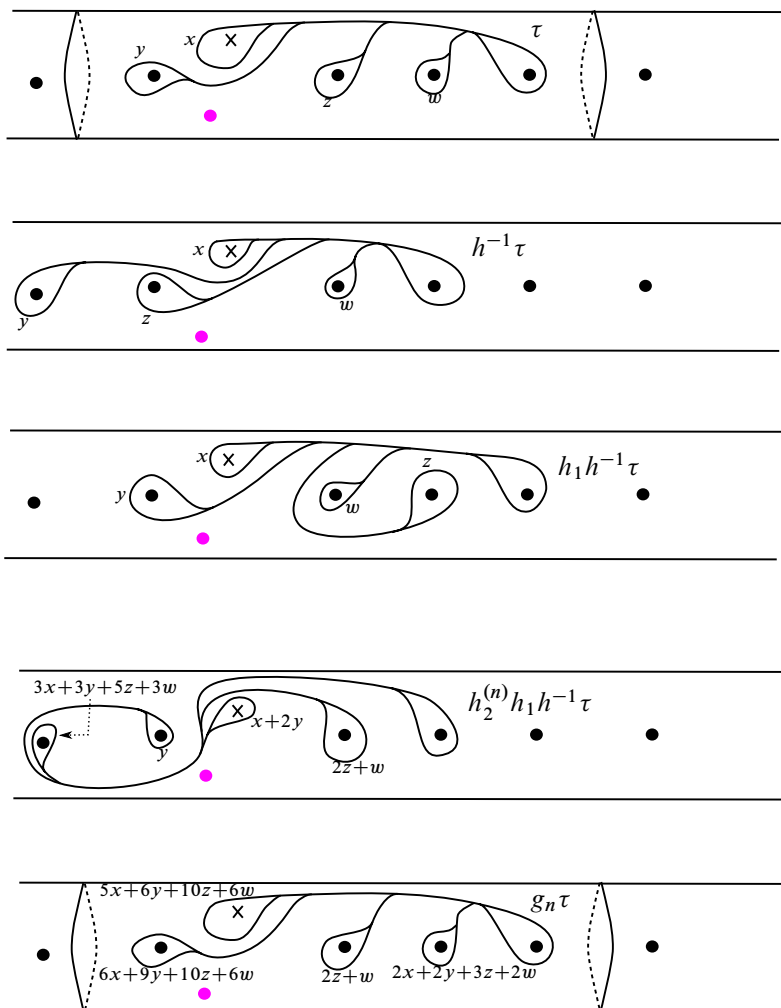


Figure 44: The train track τ_n is preserved by $\varphi_n = g_n h^{-1}$ for $n = 1, 2$; the pink puncture only appears when $n = 2$. The weights in the picture are used to calculate the matrix A_n associated to τ_n .

Theorem 8.8 For each $n \in \mathbb{N}$, there is a geodesic axis for g_n in $\mathcal{A}(\Sigma, p)$. Furthermore, g_n has translation length 1 on this axis.

Proof As $d_{\mathcal{A}(\Sigma, p)}(\gamma, \delta) \leq d_{\mathcal{A}(S, p)}(\gamma, \delta)$ for any arcs $\gamma, \delta \in \mathcal{A}(S, p)$, the image of a geodesic under the inclusion $\mathcal{A}(S, p) \hookrightarrow \mathcal{A}(\Sigma, p)$ is still a geodesic. Thus it suffices to construct a geodesic axis for g_n in $\mathcal{A}(S, p)$. Toward this goal, define

$$\beta_0^{(n)} = P_s(-1)_o(-2)_o \dots (-n-1)_o(-n-1)_u(-n)_o \dots (-2)_o(-1)_o P_s \in \mathcal{A}(S, p),$$

and let

$$\beta_i^{(n)} = g_n^i(\beta_0^{(n)}).$$

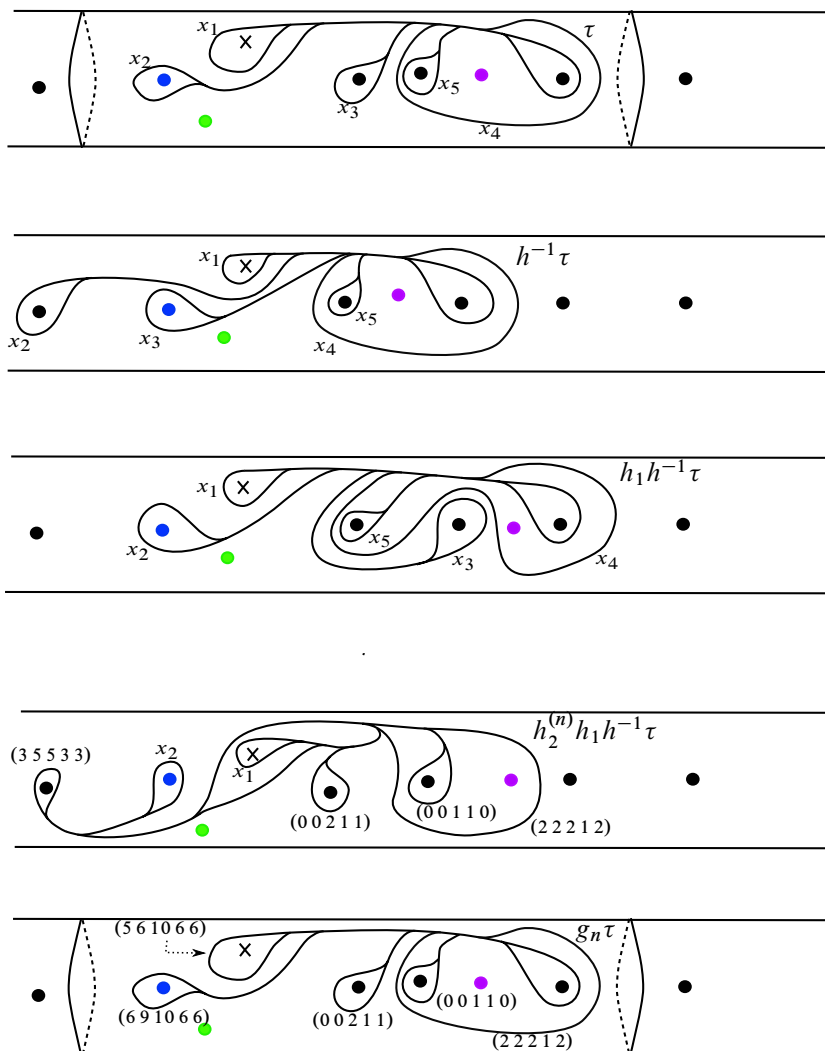


Figure 45: The train track τ_n is preserved by $\varphi_n = g_n h^{-1}$, when $n \geq 3$. The weights in the picture show how to calculate the matrix A_n associated to τ_n . For ease of notation, we often write the weights of each branch as a vector in the variables x_1, \dots, x_5 . For example, the label $(2\ 1\ 2\ 4\ 5)$ corresponds to the weight $2x_1 + x_2 + 2x_3 + 4x_4 + 5x_5$.

Since the arcs $\beta_i^{(n)}$ are the orbit of $\beta_0^{(n)}$ under $\langle g_n \rangle$ and g_n is a loxodromic isometry, it follows that they form a quasigeodesic axis in $\mathcal{A}(\Sigma, p)$ for g_n . We will show that $\beta_0^{(n)}$ is disjoint from $\beta_1^{(n)}$, which will prove that $(\beta_i^{(n)})_{i \in \mathbb{Z}}$ is a geodesic axis for g_n and that g_n has translation length one on this axis. In fact, it suffices to show that $\beta_1^{(n)}$ does not contain P_o or $k_{o/u}$ for any $k \leq -n - 1$, as it then follows that $\beta_0^{(n)}$ is disjoint from $\beta_1^{(n)}$; see Figure 46.

Applying h_1 to $\beta_0^{(n)}$ yields

$$(14) \quad h_1(\beta_0^{(n)}) = P_s(-1)_o(-2)_o \dots (-n)_o(-n)_u(-n+1)_o \dots (-2)_o(-1)_o P_s.$$




Figure 46: If an arc γ does not contain P_o or $k_{o/u}$ for any $k \leq -n-1$, then it must lie in the shaded region of S . In particular, γ must be disjoint from $\beta_0^{(n)}$.

Since all of $h_1(\beta_0^{(n)})$ is to the left of the puncture, the image under $h_2^{(n)}$ (which shifts to the left) will not contain P_o . Moreover, since P_o disagrees with the code for the domain of $h_3^{(n)}$, neither will $\beta_1^{(n)}$. Thus it remains to show that $\beta_1^{(n)}$ does not contain $k_{o/u}$ for any $k \leq -n-1$.

Recall that $h_3^{(n)}$ shifts to the right and has shift region $(-\infty, -n-1] \cup [n+1, \infty)$. Thus, any instance of $k_{o/u}$ with $k \leq -n-1$ in $\beta_1^{(n)}$ must be the image of $(k-1)_{o/u}$ in $h_2^{(n)}(h_1(\beta_0^{(n)}))$. Similarly, such a $(k-1)_{o/u}$ must be the image of $k_{o/u}$ in $h_1(\beta_0^{(n)})$, since $h_2^{(n)}$ shifts to the left. However, by (14), $h_1(\beta_0^{(n)})$ does not contain $j_{o/u}$ for $j \leq -n-1$, and we conclude that $\beta_1^{(n)}$ does not contain $k_{o/u}$ for any $k \leq -n-1$. \square

8.4 Limit points of the g_n

Since the relative arc graph $\mathcal{A}(\Sigma, p)$ is a hyperbolic metric space, it has a well-defined Gromov boundary. This boundary was described by Bavard and Walker [7; 8]. In this section, we describe the limit set of g_n on $\partial\mathcal{A}(\Sigma, p)$ in terms of Bavard and Walker's characterization of the boundary.

8.4.1 The Gromov boundary of $\mathcal{A}(\Sigma, p)$ We begin by recalling some definitions from [7; 8]. It is important to mention that in [7; 8], the word *loop* is used for what we call an arc in this paper. Any time we mention a result from one of these two papers, we will convert it to our terminology. Let $E(\Sigma)$ denote the space of ends of Σ , which necessarily contains our preferred puncture p .

Fix any hyperbolic metric (of the first kind) on Σ , as in [8, Theorem 3.0.1]. For a fixed lift of p to the universal cover \mathbb{H}^2 , which necessarily lies on $\partial\mathbb{H}^2$, there exists a parabolic subgroup $G < \pi_1(\Sigma)$ stabilizing this lift. Define $\hat{\Sigma} = \mathbb{H}^2/G$ to be the intermediate cover associated to this parabolic subgroup. The space $\hat{\Sigma}$ is called the *conical cover* of Σ . This cover has boundary \mathbb{S}^1 and contains a preferred lift \hat{p} of p which comes from our fixed choice in the universal cover. Let $\pi: \hat{\Sigma} \rightarrow \Sigma$ be the natural quotient map, let $\hat{\beta}$ be any geodesic ray from \hat{p} to $\partial\hat{\Sigma}$, and let $\beta = \pi(\hat{\beta})$. Thus $\hat{\beta}$ has one endpoint on \hat{p} and the other endpoint somewhere in $\partial\hat{\Sigma} \simeq \mathbb{S}^1$. The other endpoint may be a lift of p that is not our chosen \hat{p} , (a lift of) a point in $E(\Sigma) \setminus \{p\}$, or a point which is neither. If β is simple, then in the first case β is an arc,² in the second case β is a *short ray*, and in the last case β is a *long ray*. Equivalently, an embedding

²As noted above, Bavard and Walker call this a *loop*.

$\beta: (0, 1) \rightarrow \Sigma$ is a short ray if it can be continuously extended to a map $\hat{\beta}: [0, 1] \rightarrow \Sigma \cup E(\Sigma)$ such that $\hat{\beta}|_{(0,1)} = \beta$, $\hat{\beta}(0) = p$, and $\hat{\beta}(1) \in E(\Sigma) \setminus \{p\}$ (see [8, Section 5.1.1]), and β is a long ray if it is neither an arc nor a short ray.

Bavard and Walker construct a graph involving all three kinds of rays, which they use to describe the Gromov boundary $\partial\mathcal{A}(\Sigma, p)$ of the relative arc graph.

Definition 8.9 The *completed ray graph* $\mathcal{R}(\Sigma, p)$ is the graph whose vertices are isotopy classes of simple arcs, short rays, and long rays and whose edges correspond to homotopically disjoint isotopy classes.

By definition, $\mathcal{A}(\Sigma, p)$ embeds in $\mathcal{R}(\Sigma, p)$, but the following theorem shows that something stronger is true. Recall that a clique is a complete graph.

Theorem 8.10 [8, Theorem 5.7.1] *The connected component of $\mathcal{R}(\Sigma, p)$ containing all arcs is quasi-isometric to $\mathcal{A}(\Sigma, p)$. All other connected components are cliques.*

A particular type of long ray will be important in the description of the Gromov boundary of $\mathcal{A}(\Sigma, p)$. A long ray β is *k-filling* if k is the minimum natural number such that there exists an arc β_0 and long rays $\beta_1, \dots, \beta_k = \beta$ such that $\beta_i \cap \beta_{i+1} = \emptyset$ for all $i \geq 0$. In other words, a long ray is *k-filling* if it is distance exactly k from an arc in $\mathcal{R}(\Sigma, p)$.

Definition 8.11 A long ray β is said to be *high-filling* if both of the following hold:

- (1) β intersects every short ray; and
- (2) β is not *k-filling* for any $k \in \mathbb{N}$.

All of the vertices of the connected components that form cliques are high-filling rays; accordingly, such cliques are called *high-filling cliques*.

Theorem 8.12 [8, Theorem 6.3.1] *The set of high-filling cliques is in bijection with the Gromov boundary $\partial\mathcal{A}(\Sigma, p)$ of the relative arc graph.*

8.4.2 The limit set of g_n In [8, Section 7.1], Bavard and Walker prove that to associated to any $f \in \text{Map}(\Sigma, p)$ acting loxodromically on $\mathcal{A}(\Sigma, p)$, there exists an attracting clique of high-filling rays $C^+(f)$ and a repelling clique of high-filling rays $C^-(f)$ in $\mathcal{R}(\Sigma, p)$ that correspond to the attracting and repelling limit points of f in $\partial\mathcal{A}(\Sigma, p)$, respectively. The cliques $C^+(f)$ and $C^-(f)$ have the same (finite) number of vertices, called the *weight of f* , denoted by $w(f)$ (see [8, Definition 7.1.3]).

Following [7, Example 2.7.1], we have the following lemma.

Lemma 8.13 *For any $n \geq 1$, the homeomorphism g_n constructed in Theorem 1.1 satisfies $w(g_n) = 1$.*

Proof For notational simplicity, fix n and define $g = g_n$ and $\alpha_i = \alpha_i^{(n)}$. It suffices to prove that the attracting clique $C^+(g)$ consists of a single high-filling ray.

For each i , let $\hat{\alpha}_i$ be a lift of α_i to the conical cover $\hat{\Sigma}$. Then $\hat{\alpha}_i$ is a simple geodesic arc with one endpoint on \hat{p} and the other on some $x_i \in \partial\hat{\Sigma}$. By [8, Lemma 5.2.2], the set of endpoints of arcs, short rays, and long rays is compact in $\partial\hat{\Sigma}$, and hence there exists a subsequence (x_{i_k}) of (x_i) which limits to a point $x_\infty \in \partial\hat{\Sigma}$. Let $\hat{\alpha}_\infty$ be the geodesic ray from \hat{p} to x_∞ and $\alpha_\infty = \pi(\hat{\alpha}_\infty)$, where $\pi: \hat{\Sigma} \rightarrow \Sigma$ is the covering map. The construction of the conical cover $\hat{\Sigma}$ shows that α_∞ has an infinite-length code with initial segment $\hat{\alpha}_{i_k}$ for any k . Since $\ell_c(\hat{\alpha}_{i_k}) \rightarrow \infty$ as $k \rightarrow \infty$ and $\hat{\alpha}_{i_k}$ has an initial segment equal to $\hat{\alpha}_{i_{k-1}}$, this uniquely determines the entire infinite code. We claim that α_∞ is a high-filling ray and, moreover, that α_∞ is the sole vertex in $C^+(g)$.

That α_∞ is a ray follows from the fact that the set of endpoints of arcs in $\partial\hat{\Sigma}$ are isolated [8, Lemma 5.2.2]. To see that α_∞ is high-filling, it therefore suffices to show that it intersects every other ray (short or long). It follows from the same proof as Theorem 7.10 that any ray β which is disjoint from α_{i_k} must begin like $\alpha_{i_{k-1}}$. In particular, we see that the first $\ell_c(\hat{\alpha}_{i_{k-1}})$ characters in a code for β must agree with $\hat{\alpha}_{i_{k-1}}$. Since $\ell_c(\hat{\alpha}_{i_{k-1}}) < \ell_c(\hat{\alpha}_{i_k})$ for all k , taking $k \rightarrow \infty$ shows that any ray β which is disjoint from α_∞ must have identical code and hence indeed must be exactly α_∞ .

Finally, since α_∞ intersects every other ray, it must, in particular, intersect any other high-filling ray. By [7, Lemma 2.7.8], the connected component of any high-filling ray is a clique of high-filling rays, and thus α_∞ is its own connected component in $\mathcal{R}(\Sigma, p)$. Hence $w(g) = 1$, completing the proof. \square

We close this section with the remark that if the weight of the limit points of g_n were not all the same, then Bavard and Walker give a method for constructing nontrivial quasimorphisms [8, Theorem 9.1.1]. However, since this is not the case, we must use a different method for showing that the space of quasimorphisms is infinite-dimensional, which is related to Bavard's original proof for the arc graph from [5]. We do so in Section 9.

9 Infinite-type quasimorphisms

A *quasimorphism* of a group G is a map $f: G \rightarrow \mathbb{R}$ such that there exists a real constant C for which $|f(ab) - f(a) - f(b)| \leq C$ for all $a, b \in G$. The set of quasimorphisms forms a vector space $V(G)$ over \mathbb{R} , and, moreover, bounded functions and homomorphisms both form subspaces of $V(G)$. Let $\widetilde{\text{QH}}(G)$ denote the quotient of $V(G)$ by the subspace spanned by bounded functions and homomorphisms. We call $\widetilde{\text{QH}}(G)$ the *space of quasimorphisms* of G . The goal of this section is to use the elements constructed in Theorem 1.1 to prove Theorem 1.4, which we restate for the convenience of the reader.

Theorem 1.4 *Let Σ be an admissible surface. The space of nontrivial quasimorphisms on $\text{Map}(\Sigma, p)$ is infinite-dimensional.*




Figure 47: A schematic of D_p and the two arcs a^- and a^+ .

Several of the ideas in this section closely follow the strategy and ideas of Bavard [5], though the production of the elements which give rise to our quasimorphisms differs. We begin by studying a specific subclass of arcs with the goal of describing a particular homotopy operation and intersection pairing on them. We then use this intersection pairing and a theorem of Bestvina and Fujiwara to prove the theorem.

9.1 An intersection pairing on symmetric arcs and first properties

Recall that our surface Σ has an embedding of the biinfinite flute surface S such that $\Sigma \setminus S$ has infinitely many connected components, a countable collection of which are homeomorphic to a fixed surface Σ_0 . Moreover, the complement of each arc α_i separates Σ into two components, one of which is homeomorphic to $\text{int}(\Sigma_0)$, the interior of the fixed surface Σ_0 . Using Σ_0 , we define \mathcal{SA} to be the set of simple, symmetric arcs δ (see Definition 4.4) such that $\Sigma \setminus \delta$ has two connected components, one of which is homeomorphic to $\text{int}(\Sigma_0)$. Notice that $\alpha_i^{(n)} \in \mathcal{SA}$ for all $i \in \mathbb{Z}$, $n \in \mathbb{N}$ and that \mathcal{SA} is a $\text{Map}(\Sigma, p)$ -invariant subset of the set of all arcs.

Since p is isolated, we again fix the small once-punctured disk D_p containing p as in Section 7. This disk is homeomorphic to the closed unit disk minus an interior point. As in that section, given any element δ in \mathcal{SA} , we put δ in standard position so that its intersections with $\partial D_p \approx S^1$ are all transverse. Let x_s and x_t be the initial and terminal point where δ intersects ∂D_p . Let λ_0 and ρ_0 be the subsegments of δ which connect x_t and x_s to p , respectively.

We will modify δ to form a particular simple closed curve as follows. We can replace $\rho_0 \cup \lambda_0$ with either a^+ or a^- , as shown in Figure 47, forming two distinct simple closed curves, δ^+ and δ^- , respectively. One of these two curves bounds a surface homeomorphic to $\text{int}(\Sigma_0)$; in Figure 47, this curve is δ^- . Fixing a hyperbolic metric on the surface, we let B_δ be the geodesic representative of this curve.

We now choose the homotopy representative of δ that will allow us to define an intersection pairing. As δ is symmetric, there exists an arc

$$\delta' = r_\delta B_\delta r_\delta^{-1}$$




Figure 48: Two examples of arcs in \mathcal{SA} (in the case that Σ_0 is a one-holed torus) and their corresponding zippings. The purple and red arcs are zipped to the green and blue arcs, respectively.

in the homotopy class of δ , where r_δ is a simple ray from the puncture to B_δ that intersects B_δ only at its endpoint. Intuitively, one can think of the arc δ' as being constructed from δ by “zipping” the initial and terminal portions of the arc together for as long as they fellow travel to form r_δ . Particular examples of r_δ and δ' are given in Figure 48.

We are now in a position to define the intersection pairings.

Definition 9.1 We define a map $I^\pm: \mathcal{SA} \times \mathcal{SA} \rightarrow \mathbb{Z}_{\geq 0}$ as follows. Let $I^+(\delta, \gamma)$ be the number of positively oriented intersections between minimal position representatives for the homotopy classes of r_δ and r_γ that do not occur on B_δ or B_γ . Here we require that the homotopy fixes the puncture and fixes each of B_δ and B_γ setwise (though not necessarily pointwise). We define $I^-(\delta, \gamma)$ similarly using negative intersections.

Notice that $I^+(-, -)$ and $I^-(-, -)$ are not necessarily symmetric in their arguments. However, it is straightforward to verify that the relationship

$$I^+(\delta, \gamma) = I^-(\gamma, \delta)$$

holds for any $\delta, \gamma \in \mathcal{SA}$.

For the remainder of the section we will fix an $n \in \mathbb{N}$ and use the notation that $g = g_n$, $\alpha_i = \alpha_i^{(n)}$, and $\varphi = \varphi_n$ is the “starts like” function from Section 8.

Example 9.2 One can readily compute from Figure 49 that we have the following relations:

$$5 = I^-(\alpha_0, \alpha_2) = I^+(\alpha_2, \alpha_0) \quad \text{and} \quad 6 = I^+(\alpha_0, \alpha_2) = I^-(\alpha_2, \alpha_0).$$

These calculations will be relevant later in the section.

We now collect some useful properties of the intersection pairing and its interaction with $\text{Map}(\Sigma, p)$. These preliminary facts are inspired by Bavard [5], where similar statements are shown in Bavard’s context.




Figure 49: The oriented arc $r_{\alpha_0^{(n)}}$ is shown in blue and the oriented arc $r_{\alpha_2^{(n)}}$ is shown in pink. Each line of $r_{\alpha_0^{(n)}}$ and $r_{\alpha_2^{(n)}}$ represents two strands of $\alpha_0^{(n)}$ and $\alpha_2^{(n)}$, respectively.

Lemma 9.3 *The intersection pairing is mapping class group invariant. That is, for any $\delta, \gamma \in \mathcal{SA}$ and any $f \in \text{Map}(\Sigma, p)$*

$$I^\pm(f(\delta), f(\gamma)) = I^\pm(\delta, \gamma).$$

Proof This is immediate from the fact that $\text{Map}(\Sigma, p)$ is orientation preserving and preserves \mathcal{SA} . \square

Recall that in Section 8, we defined the “starts like” function

$$\varphi: \mathcal{A}(S, p) \rightarrow \mathbb{Z}_{\geq 0},$$

by setting $\varphi(\delta)$ equal to the largest $i \geq 0$ such that δ starts like α_i . We now extend φ to all of $\mathcal{A}(\Sigma, p)$ by setting $\varphi(\delta) = 0$ if δ does not have a homotopy representative that is contained in S . We continue to call this extension φ .

Lemma 9.4 *If $\delta, \gamma \in \mathcal{SA}$ are arcs such that $2 + \varphi(\delta) \leq \varphi(\gamma)$, then*

$$6 \leq I^-(\gamma, \delta).$$

Proof By the mapping class group invariance of Lemma 9.3, the quantities $I^\pm(\alpha_i, \alpha_j)$ depend only on $j - i$. As α_{i+1} starts like α_i , the pairing $I^-(\alpha_i, \alpha_j)$ must be monotonically increasing in the difference $j - i$. In particular, if $2 + i \leq j$, then

$$6 = I^-(\alpha_2, \alpha_0) \leq I^-(\alpha_j, \alpha_i).$$

If $\varphi(\delta) = i$ and $\varphi(\gamma) = j$, then δ starts like α_i and γ starts like α_j , so the arcs γ and δ must have at least as many negatively oriented intersections as α_i and α_j . Thus we have

$$6 \leq I^-(\alpha_j, \alpha_i) \leq I^-(\gamma, \delta),$$

and the result follows. \square

9.2 Production of “infinite-type” quasimorphisms

We now use the intersection pairing from the previous subsection to show that the elements g_n give rise to nontrivial quasimorphisms. For this, we need the following theorem of Bestvina and Fujiwara. We explain the two conditions on h_1, h_2 after the statement.

Theorem 9.5 [10, Theorem 1] *Suppose that a group G has a nonelementary action by isometries on a δ -hyperbolic graph X . If there exist independent loxodromic elements $h_1, h_2 \in G$ such that $h_1 \sim h_2$, then the space of quasimorphisms is infinite-dimensional.*

Two loxodromic isometries h_1 and h_2 are *independent* if their limit sets in the boundary ∂X of X are disjoint. For the second condition, fix constants $K \geq 1$ and $K' \geq 0$ so that h_i has a (K, K') -quasigeodesic axis ℓ_i in X for $i \in \{1, 2\}$. A fundamental property of δ -hyperbolic spaces is that there exists $B = B(K, K', \delta)$ such that any two finite (K, K') -quasigeodesics with common endpoints are within distance B of each other. Define an equivalence relation on elements $h_1, h_2 \in G$ so that $h_1 \sim h_2$ if the following holds: for any arbitrarily long segment L of ℓ_1 , there exists an $f \in G$ such that $f(L)$ is contained in the B -neighborhood of ℓ_2 and the map $f: L \rightarrow f(L)$ is orientation-preserving with respect to the h_i -orientation on ℓ_i for $i \in \{1, 2\}$. For the definition of the h_i -orientation on L and $f(L)$, see [10, page 72].

We now recall some arguments from [5, Section 4.3] which, when adapted into our language, show that $g \sim g^{-1}$ for our loxodromic isometries $g = g_n$. Fix $B \geq 1$ to be the constant defined above for all $(4, 0)$ -quasigeodesics in $\mathcal{A}(\Sigma, p)$. Let $\ell = \{g^i(\alpha_0)\}_{i \in \mathbb{Z}}$, so that ℓ is a $(4, 0)$ -quasigeodesic axis of g by Theorem 8.3. We then have the following statements that are similar to [5, Lemmas 4.6, 4.7]. We supply the proofs for the reader’s convenience.

Lemma 9.6 *Let L be a subpath of ℓ from α_i to α_j for $0 < i < j$. Let $f \in \text{Map}(\Sigma, p)$ be such that $d_{\mathcal{A}(\Sigma, p)}(\alpha_i, f(\alpha_j)) \leq B$ and such that $f(L) \subset N_B(\ell)$ with the opposite orientation. If $j - i > 8B + 3$, then there exists some k such that $i \leq k < j$ and $\varphi(f(\alpha_{k+2})) \leq \varphi(f(\alpha_k)) - 2$.*

Proof Since $d_{\mathcal{A}(\Sigma, p)}(f(\alpha_j), \alpha_i) \leq B$, we conclude by Lemma 8.1 that $\varphi(f(\alpha_j)) \leq i + B$. Since $f(L) \subset N_B(\ell)$ with the opposite orientation, we may apply [5, Lemma 4.4] to L and the reverse of $f(L)$ to conclude that $d_{\mathcal{A}(\Sigma, p)}(\alpha_j, f(\alpha_i)) \leq 3B$. Note that [5, Lemma 4.4] is stated for geodesics, but the exact same proof goes through for quasigeodesics. Again applying Lemma 8.1, we see that $\varphi(f(\alpha_i)) \geq j - 3B$.

Suppose towards a contradiction that $\varphi(f(\alpha_{k+2})) > \varphi(f(\alpha_k)) - 2$ for all $1 \leq k < j$. Equivalently, $\varphi(f(\alpha_{k+2})) \geq \varphi(f(\alpha_k)) - 1$ for every k . Then, if $j - i$ is even,

$$i + B \geq \varphi(f(\alpha_j)) \geq \varphi(f(\alpha_i)) - \frac{1}{2}(j - i) \geq j - 3B - \frac{1}{2}(j - i),$$

where the second inequality is obtained by applying $\varphi(f(\alpha_{k+2})) \geq \varphi(f(\alpha_k)) - 1$ repeatedly starting with $j = k + 2$. If $j - i$ is odd, then by the same reasoning we have

$$i + B \geq \varphi(f(\alpha_j)) \geq \varphi(f(\alpha_{i+1})) - \frac{1}{2}(j - i - 1).$$

Since

$$|\varphi(f(\alpha_i)) - \varphi(f(\alpha_{i+1}))| \leq d(f(\alpha_i), f(\alpha_{i+1})) = d(\alpha_i, \alpha_{i+1}) = 2,$$

it follows that

$$i + B \geq \varphi(f(\alpha_{i+1})) - \frac{1}{2}(j - i - 1) \geq \varphi(f(\alpha_i)) - \frac{3}{2} - \frac{1}{2}(j - i) \geq j - 3B - \frac{3}{2} - \frac{1}{2}(j - i).$$

Hence we conclude that, in either case,

$$4B + \frac{3}{2} \geq \frac{1}{2}(j - i),$$

which contradicts that $j - i > 8B + 3$. Thus there must be some k for which $\varphi(f(\alpha_{k+2})) \leq \varphi(f(\alpha_k)) - 2$, as required. \square

Proposition 9.7 *For any segment L of ℓ whose length is greater than $32B + 12$ and any $f \in \text{Map}(\Sigma, p)$, if $f(L) \subset N_B(\ell)$ then $f(L)$ has the same orientation as ℓ .*

Proof After possibly increasing the length of L , we may assume that L is a subpath of ℓ from α_i to α_j for some $i < j$. Since the length of L is greater than $32B + 12$ and ℓ is a $(4, 0)$ -quasigeodesic edge path by Theorem 8.3, we have that $j - i > 8B + 3$. By precomposing f with a suitable power of g , we can and do assume that $i, j > 0$.

Assume for contradiction that $f(L)$ has the opposite orientation as ℓ . By Example 9.2 and Lemma 9.3 we have that

$$(15) \quad I^-(f(\alpha_k), f(\alpha_{k+2})) = I^-(\alpha_k, \alpha_{k+2}) = I^-(\alpha_0, \alpha_2) = 5$$

for all $k \in \mathbb{Z}$. On the other hand, by Lemma 9.6 there is some fixed index $i \leq k < j$ for which

$$(16) \quad 2 + \varphi(f(\alpha_{k+2})) \leq \varphi(f(\alpha_k)).$$

Applying Lemma 9.4 to $f(\alpha_k)$ and $f(\alpha_{k+2})$ for this index k shows that

$$6 \leq I^-(f(\alpha_k), f(\alpha_{k+2})) = I^-(\alpha_k, \alpha_{k+2}) = I^-(\alpha_0, \alpha_2).$$

However, this contradicts (15), and so we conclude that $f(L)$ has the same orientation as L . \square

Proposition 9.7 implies that $g \sim g^{-1}$ since the axis ℓ has opposite orientations for g and g^{-1} . Additionally, since conjugate elements are equivalent under the relationship \sim , we have the following immediate corollary of Proposition 9.7.

Corollary 9.8 *For fixed $n \in \mathbb{N}$, the loxodromic elements $g = g_n$ have the property that $g \sim hg^{-1}h^{-1}$ for any $h \in \text{Map}(\Sigma, p)$.*

With this in hand, we can prove the main result of this section.

Proof of Theorem 1.4 Fix $n \in \mathbb{N}$, and continue to use the notation that $g = g_n$. By Theorem 9.5 and Corollary 9.8 it suffices to show that there exists an $h \in \text{Map}(\Sigma, p)$ such that g and $hg^{-1}h^{-1}$ are independent loxodromic elements. For this we can use any $h \in \text{Map}(\Sigma, p)$ which does not fix the limit set of g (which is the same as the limit set of g^{-1}). For example, fix any finite-type subsurface Π_0 with boundary of sufficient complexity which contains the puncture p , and take any pseudo-Anosov $h \in \text{Map}(\Pi_0, p)$. Extending h by the identity outside of Π_0 , we may consider h as an element of $\text{Map}(\Sigma, p)$. By Lemma 2.10, there is a $(2, 0)$ -quasi-isometric embedding $\iota: \mathcal{A}(\Pi_0, P) \hookrightarrow \mathcal{A}(\Sigma, p)$. Since h is loxodromic with respect to the action of $\text{Map}(\Pi_0, p)$ on $\mathcal{A}(\Pi_0, p)$, it is therefore also loxodromic with respect to the action of $\text{Map}(\Sigma, p)$ on $\mathcal{A}(\Sigma, p)$. In addition, out of all such pseudo-Anosovs, there is a choice of h whose limit points are different from the limit points ξ_{\pm} of g . (Note that the quasi-isometric embedding ι ensures that two pseudo-Anosovs with distinct limit points in the boundary of $\mathcal{A}(\Pi_0, p)$ also have distinct limit points in the boundary of $\mathcal{A}(\Sigma, p)$.) Therefore, the limit points $h(\xi_{\pm})$ of $hg^{-1}h^{-1}$ are distinct from those of g . In particular, g and $hg^{-1}h^{-1}$ are independent loxodromic elements, as required. \square

10 Convergence to a geodesic lamination

The goal of this section is to prove Theorem 1.2, which we restate for the convenience of the reader. As in [26], we equip Σ with its unique conformal hyperbolic metric. Additionally, we require that with this metric, Σ is equal to its convex core, which is equivalent to eliminating hyperbolic funnels and half-planes in Σ . This is necessary in order to consider geodesic laminations on an infinite-type surface; see [26].

Theorem 1.2 *If Σ is an admissible surface equipped with its conformal hyperbolic metric that is equal to its convex core, then there exists a simple closed curve c_0 on Σ such that the sequence $(g_n^i(c_0))_{i \in \mathbb{N}}$ converges to a geodesic lamination on Σ .*

10.1 Geodesic laminations

We begin by reviewing some facts about geodesic laminations on infinite-type surfaces. For a complete treatment of the subject, we refer the reader to [26].

Definition 10.1 A *geodesic lamination* λ on Σ is a foliation of a closed subset of Σ by complete geodesics.

Fix a locally finite geodesic pants decomposition $\{P_n\}$ of Σ and a train track Θ on Σ constructed as in [26, Section 4]. Denote by $\tilde{\Theta}$ the lift of Θ to $\tilde{\Sigma}$, the universal cover of Σ . An *edge path* of $\tilde{\Theta}$ is a finite, infinite, or biinfinite sequence of edges of $\tilde{\Theta}$ such that consecutive edges meet smoothly at each vertex. Every biinfinite edge path has two distinct accumulation points on $\partial_{\infty} \tilde{\Sigma}$ by [26, Proposition 4.5].

Given a biinfinite edge path $\tilde{\gamma}$ of $\tilde{\Theta}$, let $G(\tilde{\gamma})$ be the geodesic of $\tilde{\Sigma}$ whose endpoints on $\partial_{\infty} \tilde{\Sigma}$ are the two distinct endpoints of $\tilde{\gamma}$. A geodesic g of $\tilde{\Sigma}$ is *weakly carried* by $\tilde{\Theta}$ if there exists a biinfinite path $\tilde{\gamma}$ in $\tilde{\Theta}$ such that $G(\tilde{\gamma}) = g$.

We now gather various results from [26] which will be useful in what follows. These are all analogous to the situation for finite-type surfaces (see, for example, [11; 23; 27]). The first gives a correspondence between geodesics weakly carried by $\tilde{\Theta}$ and biinfinite edge paths.

Proposition 10.2 [26, Proposition 4.7] *There is a one-to-one correspondence between biinfinite edge paths of $\tilde{\Theta}$ and geodesics weakly carried by $\tilde{\Theta}$.*

Let γ be an edge path in Θ and $\tilde{\gamma}$ a single component of the lift of γ to $\tilde{\Sigma}$. Then, as above, we construct a geodesic $G(\tilde{\gamma})$ with the same endpoints as $\tilde{\gamma}$ and denote by $G(\gamma)$ its projection to Σ . We then say that the geodesic $G(\gamma)$ is *weakly carried by Θ* if $G(\tilde{\gamma})$ is weakly carried by $\tilde{\Theta}$. A geodesic lamination λ on Σ is *weakly carried by Θ* if every geodesic of λ is weakly carried by Θ . The next result gives a correspondence between geodesic laminations and certain families of biinfinite edge paths.

Proposition 10.3 [26, Proposition 4.11] *The set of geodesic laminations on Σ that are weakly carried by Θ is in one-to-one correspondence with the families Γ of biinfinite edge paths of $\tilde{\Theta}$ that satisfy*

- any two biinfinite edge paths γ and γ' in Γ do not cross; and
- if γ is a biinfinite edge path such that for any finite edge subpath there is a biinfinite edge path in Γ that contains it, then $\gamma \in \Gamma$.

The following proposition describes when a sequence of geodesics carried by Θ converge. Let $G(\tilde{\Sigma})$ be the set of unoriented geodesics on $\tilde{\Sigma}$.

Proposition 10.4 [26, Proposition 4.9] *Let $f_n, f \in G(\tilde{\Sigma})$ be weakly carried by a train track $\tilde{\Theta}$, and denote by $\tilde{\gamma}_n, \tilde{\gamma}$ the corresponding biinfinite edge paths in $\tilde{\Theta}$. Then f_n converges to f as $n \rightarrow \infty$ if and only if for each finite subpath $\tilde{\gamma}'$ of $\tilde{\gamma}$ there exists $n_0 \geq 0$ such that $\tilde{\gamma}'$ is contained in the path $\tilde{\gamma}_n$ for all $n \geq n_0$.*

The final proposition we will need shows that every geodesic lamination is weakly carried by a train track.

Proposition 10.5 [26, Proposition 4.12] *Every geodesic lamination λ on a hyperbolic surface X is weakly carried by a train track Θ that is constructed as above starting from a fixed locally finite geodesic pants decomposition.*

10.2 Construction of the train track Θ

We now construct a train track on our surface Σ . In the next section, we will define our simple closed curve c_0 and show that $g_n^i(c_0)$ converges as $i \rightarrow \infty$ to a geodesic lamination that is weakly carried by this train track.

First, fix a simple closed curve γ'_p on S so that the bounded component of $S \setminus \gamma_p$ is the pair of pants with cuffs γ'_p , p_{-1} , and p (recall that punctures are allowed to be cuffs). Let γ_p be the image of γ'_p under the embedding $S \hookrightarrow \Sigma$. We construct our first pair of pants Q_p to have cuffs γ_p , B_{-1} , and p . Next




Figure 50: The cuffs of the pants decomposition are in blue and black, and the pairs of pants are labeled in red. Anything dotted occurs on the back of the embedded copy of S in Σ .

choose two simple closed separating curves $\gamma_{-1,l}$ and $\gamma_{-1,r}$ on the embedded copy of S in Σ so that one component of $\Sigma \setminus (\gamma_{-1,l} \cup \gamma_{-1,r})$ contains two pairs of pants: Q_p and a pair of pants Q_{-1} with cuffs $\gamma_{-1,l}$, $\gamma_{-1,r}$, and γ_p . Let $\gamma_{-1,l}$ be the curve which bounds the connected component of $\Sigma \setminus (\gamma_{-1,l} \cup \gamma_{-1,r})$ containing B_{-2} . For each $i \in \mathbb{Z} \setminus \{-1\}$, choose two simple closed separating curves $\gamma_{i,l}$ and $\gamma_{i,r}$ on the embedded copy of S in Σ so that one component of $\Sigma \setminus (\gamma_{i,l} \cup \gamma_{i,r})$ contains a pair of pants, Q_i , with cuffs $\gamma_{i,l}$, $\gamma_{i,r}$, and B_i . Similarly, let $\gamma_{i,l}$ be the curve which bounds the connected component of $\Sigma \setminus (\gamma_{-1,l} \cup \gamma_{-1,r})$ containing B_{i-1} .

Next, for all $i \in \mathbb{Z}$, consider the component C_i of $\Sigma \setminus (\gamma_{i,r} \cup \gamma_{i+1,l})$ which does not contain $\gamma_{i,l}$ for any i . If C_i is a cylinder, that is, if $\gamma_{i,r}$ is homotopic to $\gamma_{i+1,l}$, then modify $\gamma_{i,r}$ and Q_i so that $\gamma_{i,r} = \gamma_{i+1,l}$. If C_i is a pair of pants, let $Q_{i,r} = C_i$. If C_i is neither a pair of pants nor a cylinder, fix a simple closed curve δ_i so that one component of $C_i \setminus \delta_i$ is a pair of pants $Q_{i,r}$ with cuffs $\gamma_{i,r}$, $\gamma_{i+1,l}$, and δ_i . See Figure 50.

We replace each cuff with a geodesic representative of the same homotopy class; by a slight abuse of notation, we continue to call the resulting geodesic pairs of pants Q_i , Q_p , and $Q_{i,r}$. This is a locally finite geodesic pants decomposition of a subsurface of the embedded copy of S in Σ , and we extend it to a locally finite geodesic pants decomposition of Σ , which we denote \mathcal{Q} .

We next construct the train track Θ as follows. The specific connectors on the front of the surface for all Q_i , Q_p , and $Q_{i,r}$ are as in Figure 51. For all other pairs of pants, we choose any connectors that satisfy the conditions of [26].




Figure 51: A portion of the train track Θ . The connectors are in purple, and the labels are the codes for certain connectors and cuffs.

We code a subset of the connectors and cuff segments of Θ which lie on the front of S in the following way (see Figure 51):

- In each pair of pants Q_i with $i \neq -1$, denote the connectors from $\gamma_{i,l}$ and $\gamma_{i,r}$ to B_i by i_L and i_R , respectively. The two segments of the cuff B_0 are denoted $0_o, 0_u$, as in our standard code.
- In each pair of pants $Q_{i,r}$, denote the single connector on the front of S by i_{RR} .
- In the pair of pants Q_p , denote the connector which has both endpoints on γ_p by P_u . This connector divides the cuff γ_p into two segments, denote these by P_o and $(-1)_u$, as in Figure 51. We will not code the second connector or the segments on the cuff B_{-1} in this pair of pants.
- In the pair of pants Q_{-1} , denote the connectors from $\gamma_{-1,l}$ and $\gamma_{-1,r}$ to γ_p by $(-1)_L$ and $(-1)_R$, respectively.

If a simple closed curve is carried by this train track and does not intersect any subpaths without a code, then we call this a Θ -code for the given simple closed curve. We say a Θ -code is *reduced* if no two adjacent characters are the same.

10.3 The simple closed curves c_i

For the remainder of this section, we fix $n \in \mathbb{N}$ and use the notation $g = g_n$ and $\alpha_i = \alpha_i^{(n)}$.

In [26], it is assumed that the relevant geodesic laminations do not contain geodesics that run out a cusp at one (or both) ends (see the discussion before [26, Proposition 4.12]). It is impossible for our sequence α_i to converge to a geodesic lamination of this type. In this section, we describe how to associate simple closed curves c_i to each α_i so that $g(c_{i-1}) = c_i$. In the following subsection, we prove that they converge to a geodesic lamination as in [26].

Let D_p be a small disk around the puncture p which is invariant under our homeomorphism g , as in Section 7.1. As before, we choose D_p small enough so that, for each i , $\alpha_i \cap D_p$ consists of two segments, one starting at p and one ending at p , with endpoints $z_{i,1}, z_{i,2}$ on ∂D_p . Up to homotopy, we may assume without loss of generality that $z_{i,1} = z_{j,1}$ and $z_{i,2} = z_{j,2}$ for all i, j . To form the simple closed curves c_i , we start with α_i and remove $\alpha_i \cap \text{int}(D_p)$. We then add an arc of ∂D_p from $z_{i,1}$ to $z_{i,2}$; there two possible choices of arc, and we choose the one so that one connected component of $S \setminus c_0$ contains p and p_0 . Note that this is the opposite choice than the one made in Section 9.1. See Figure 52.

For example, given the train track in Figure 51, a Θ -code of c_0 is

$$(17) \quad P_o(-1)_R(-1)_{RR}0_L0_o0_u0_L(-1)_{RR}(-1)_R P_u.$$

Note that since curves do not have a well-defined starting point, any cyclic permutation of this Θ -code for c_0 is also a Θ -code for c_0 . For the rest of the section, we fix the starting point P_o for the Θ -code for c_0 as in (17).




Figure 52: Forming c_i (right) from α_i (left). The initial and terminal segments of α_i are in purple (left), and the corresponding segment of c_i is in purple (right).

Our curves c_i and the arcs α_i agree outside of D_p , and since D_p is invariant under g , it follows that $g(c_i) = c_{i+1}$. For each i , we fix the starting point P_o for a Θ -code for c_i in such a way that applying g to the Θ -code for c_i yields the Θ -code for c_{i+1} .

10.4 Proof of Theorem 1.2

We will show that the simple closed curves c_i defined in the previous section converge to a geodesic lamination T on Σ . Our strategy is as follows. We will first fix a lift $\tilde{\gamma}_i$ of each curve c_i in $\tilde{\Sigma}$ and show that the sequence $\tilde{\gamma}_i$ converges to some $\tilde{\gamma}$. We then show that if $f_i = G(\tilde{\gamma}_i)$ and $f = G(\tilde{\gamma})$ are the corresponding geodesics that are weakly carried by $\tilde{\Theta}$, then $\lim_{i \rightarrow \infty} f_i = f$. Finally, we take the geodesic lamination T to be the image of f in Σ .

For each i , let $\ell_{\Theta}(c_i)$ be the length of any reduced Θ -code for c_i . Recall that $\ell_c(\alpha_i)$ is the code length of α_i , as in Definition 3.6.

Lemma 10.6 For each i , $\ell_c(\alpha_i) \leq \ell_{\Theta}(c_i) \leq 5\ell_c(\alpha_i)$.

Proof It is clear that the Θ -code for c_i is at least as long as the code for α_i . For the second inequality, notice that each instance of $k_{o/u}$ with $k \neq -1$, P in α_i is replaced with one of the following strings, depending on what precedes/follows the character $k_{o/u}$ and on the chosen pants decomposition,

$$k_L k_{o/u} k_R, \quad k_L k_{o/u}, \quad k_R k_{o/u}, \quad (k-1)_{RR} k_L k_{o/u} k_R, \quad (k-1)_{RR} k_L k_{o/u} k_R k_{RR}, \quad \text{or} \quad k_L k_{o/u} k_R k_{RR}.$$

If $k = -1$, then $(-1)_{o/u}$ is replaced with

$$\begin{aligned} &(-1)_L(-1)_{o/u}, \quad (-1)_{o/u}(-1)_R, \quad (-2)_{RR}(-1)_L(-1)_{o/u}, \\ &(-1)_{o/u}(-1)_R(-1)_{RR}, \quad \text{or} \quad (-2)_{RR}(-1)_L(-1)_{o/u}(-1)_{RR}. \end{aligned}$$

Finally, if $k = P$, then each $P_{o/u}$ remains the same and P_s is replaced with $P_{o/u}$. Therefore, each character in the code for α_i is replaced with at most 5 characters in the Θ -code for c_i , which gives the upper bound. \square

For each i , let $\ell_i = \ell_{\Theta}(c_i)$, and consider the Θ -code for c_i

$$(18) \quad c_i = c_1^i c_2^i \dots c_{\ell_i}^i.$$

For each i , fix the lift $\tilde{\gamma}_i$ of c_i that is the periodic biinfinite edge path

$$\tilde{\gamma}_i = (\dots, b_{-1}^i, b_0^i, b_1^i, b_2^i, \dots),$$

with period ℓ_i , where $b_j^i = c_j^i$ for each $1 \leq j \leq \ell_i$, and $b_j^i = b_{j-\ell_i}^i$ for all j . The codes for α_i and α_{i+1} agree on the first $\frac{1}{2}\ell_c(\alpha_i)$ characters. Thus it follows from Lemma 10.6 that the Θ -codes of c_i and c_{i+1} defined by (18) agree on at least the first $L_i = \lfloor \frac{1}{10}\ell_i \rfloor$ characters. Therefore, $b_j^i = b_j^{i+1}$ for all $1 \leq j \leq L_i$.

We define a biinfinite path

$$\tilde{\gamma} = (\dots, d_{-1}, d_0, d_1, d_2, \dots),$$

as follows. Intuitively, our goal is to define $\tilde{\gamma}$ so that it agrees with each $\tilde{\gamma}_i$ from d_{-L_i} to d_{L_i} . In the first step, we define the characters d_{-L_0} to d_{L_0} of $\tilde{\gamma}$ so that they agree with $\tilde{\gamma}_0$. In the second step, we define the characters d_{-L_1} to d_{L_1} of $\tilde{\gamma}$ so that they agree with $\tilde{\gamma}_1$. The key point here is that $\tilde{\gamma}_0$ and $\tilde{\gamma}_1$ agree on the characters of $\tilde{\gamma}$ that we have already defined in the first step. Thus we are not redefining d_i if $-L_0 \leq i \leq L_0$. Rather, these characters remain, and the additional information from the second step is the definition of d_i if $-L_1 \leq i < -L_0$ or $L_0 < i \leq L_1$. We then continue this process.

Formally, this is equivalent to the following definition. For each $i \geq 0$ and each $1 \leq j \leq L_i$, define

$$d_1 = b_1^i, \quad \dots, \quad d_j = b_j^i, \quad \dots, \quad d_{L_i} = b_{L_i}^i,$$

and define

$$d_0 = b_1^i, \quad \dots, \quad d_{-j+1} = b_j^i, \quad \dots, \quad d_{-L_i+1} = b_{L_i}^i.$$

For each i and all $1 \leq j \leq L_i$, since $b_j^i = b_j^{i+1}$, there is no conflict with the previous defined edges as i increases.

By construction, $\tilde{\gamma}$ is a biinfinite path in $\tilde{\Theta}$. Let $f_i = G(\tilde{\gamma}_i)$ and $f = G(\tilde{\gamma})$ be the corresponding geodesics which are weakly carried by $\tilde{\Theta}$, the existence of which is guaranteed by Proposition 10.2.

Lemma 10.7

$$\lim_{i \rightarrow \infty} f_i = f.$$

Proof This is almost immediate from the construction of $\tilde{\gamma}$. Fix any finite subpath $T \subset \tilde{\gamma}$. Then T is supported on $[-k, l]$ for some $k, l \geq 1$. Let $L = \max\{k, l\}$, and fix N such that $L_N \geq L$. Such an N exists since $\ell_c(\alpha_i) \rightarrow \infty$ implies that $\lim_{i \rightarrow \infty} L_i = \infty$. Then by construction T appears in all $\tilde{\gamma}_i$ with $i \geq N$. Convergence follows by Proposition 10.4. \square

Let T be the image of f in Σ . Figure 53 shows the train track which weakly carries the lamination, that is, the image of $\tilde{\gamma}$ in Σ .

Lemma 10.8 T is a geodesic lamination on Σ .

Proof This follows immediately from Proposition 10.3 applied to $\Gamma = \{\tilde{\gamma}\}$. \square

Finally, Theorem 1.2 follows from Lemma 10.8.




Figure 53: The train track on Σ which weakly carries the lamination T when $n = 1$. Note that the train track is, in fact, contained on the front of the embedded copy of S in Σ .

References

- [1] **American Institute of Mathematics**, *Surfaces of infinite type*, problem lists (2019) Available at <http://aimpl.org/genusinfinity/>
- [2] **J Aramayona, A Fossas, H Parlier**, *Arc and curve graphs for infinite-type surfaces*, Proc. Amer. Math. Soc. 145 (2017) 4995–5006 MR Zbl
- [3] **J Aramayona, P Patel, N G Vlamis**, *The first integral cohomology of pure mapping class groups*, Int. Math. Res. Not. 2020 (2020) 8973–8996 MR Zbl
- [4] **J Aramayona, N G Vlamis**, *Big mapping class groups: an overview*, from “In the tradition of Thurston: geometry and topology” (K Ohshika, A Papadopoulos, editors), Springer (2020) 459–496 MR Zbl
- [5] **J Bavard**, *Hyperbolicité du graphe des rayons et quasi-morphismes sur un gros groupe modulaire*, Geom. Topol. 20 (2016) 491–535 MR Zbl
- [6] **J Bavard, A Genevois**, *Big mapping class groups are not acylindrically hyperbolic*, Math. Slovaca 68 (2018) 71–76 MR Zbl
- [7] **J Bavard, A Walker**, *The Gromov boundary of the ray graph*, Trans. Amer. Math. Soc. 370 (2018) 7647–7678 MR Zbl
- [8] **J Bavard, A Walker**, *Two simultaneous actions of big mapping class groups*, Trans. Amer. Math. Soc. 376 (2023) 7603–7650 MR Zbl
- [9] **M Bestvina, K Bromberg, K Fujiwara**, *Constructing group actions on quasi-trees and applications to mapping class groups*, Publ. Math. Inst. Hautes Études Sci. 122 (2015) 1–64 MR Zbl
- [10] **M Bestvina, K Fujiwara**, *Bounded cohomology of subgroups of mapping class groups*, Geom. Topol. 6 (2002) 69–89 MR Zbl
- [11] **F Bonahon**, *Closed curves on surfaces*, preprint (2000) Available at <https://dornsife.usc.edu/francis-bonahon/wp-content/uploads/sites/205/2023/06/Bouquin.pdf>
- [12] **D Calegari**, *Big mapping class groups and dynamics*, blog post (2009) Zbl Available at <https://lamington.wordpress.com/2009/06/22/big-mapping-class-groups-and-dynamics>
- [13] **J Cantwell, L Conlon, S R Fenley**, *Endperiodic automorphisms of surfaces and foliations*, Ergodic Theory Dynam. Systems 41 (2021) 66–212 MR Zbl
- [14] **F Fournier-Facio, Y Lodha, M C B Zaremsky**, *Braided Thompson groups with and without quasimorphisms*, Algebr. Geom. Topol. 24 (2024) 1601–1622 MR Zbl

- [15] **S Hensel, P Przytycki, R C H Webb**, *1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs*, J. Eur. Math. Soc. 17 (2015) 755–762 MR Zbl
- [16] **W P Hooper**, *The invariant measures of some infinite interval exchange maps*, Geom. Topol. 19 (2015) 1895–2038 MR Zbl
- [17] **B von Kerékjártó**, *Vorlesungen über Topologie, I: Flächentopologie*, Springer (1923) Zbl
- [18] **J Lanier, M Loving**, *Centers of subgroups of big mapping class groups and the Tits alternative*, Glas. Mat. Ser. III 55(75) (2020) 85–91 MR Zbl
- [19] **W B R Lickorish**, *A representation of orientable combinatorial 3-manifolds*, Ann. of Math. 76 (1962) 531–540 MR Zbl
- [20] **I Morales, F Valdez**, *Loxodromic elements in big mapping class groups via the Hooper–Thurston–Veech construction*, Algebr. Geom. Topol. 22 (2022) 3809–3854 MR Zbl
- [21] **J Nielsen**, *Surface transformation classes of algebraically finite type*, Danske Vid. Selsk. Mat.-Fys. Medd. 2, Munksgaard, Copenhagen (1944) MR Zbl
- [22] **P Patel, N G Vlamis**, *Algebraic and topological properties of big mapping class groups*, Algebr. Geom. Topol. 18 (2018) 4109–4142 MR Zbl
- [23] **R C Penner, J L Harer**, *Combinatorics of train tracks*, Ann. of Math. Stud. 125, Princeton Univ. Press (1992) MR Zbl
- [24] **A J Rasmussen**, *WWPD elements of big mapping class groups*, Groups Geom. Dyn. 15 (2021) 825–848 MR Zbl
- [25] **I Richards**, *On the classification of noncompact surfaces*, Trans. Amer. Math. Soc. 106 (1963) 259–269 MR Zbl
- [26] **D Šarić**, *Train tracks and measured laminations on infinite surfaces*, Trans. Amer. Math. Soc. 374 (2021) 8903–8947 MR Zbl
- [27] **W P Thurston**, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. 19 (1988) 417–431 MR Zbl
- [28] **Y Verberne**, *A construction of pseudo-Anosov homeomorphisms using positive twists*, Algebr. Geom. Topol. 23 (2023) 1601–1639 MR Zbl

Department of Mathematics, Brandeis University
Waltham, MA, United States

Department of Mathematics, University of Oklahoma
Norman, OK, United States

Department of Mathematics, University of Utah
Salt Lake City, UT, United States

carolynabbott@brandeis.edu, nickmbmiller@ou.edu, patelp@math.utah.edu

Received: 9 August 2023

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

EDITORS

PRINCIPAL ACADEMIC EDITORS

John Etnyre
etnyre@math.gatech.edu
Georgia Institute of Technology

Kathryn Hess
kathryn.hess@epfl.ch
École Polytechnique Fédérale de Lausanne

BOARD OF EDITORS

Julie Bergner	University of Virginia jeb2md@eservices.virginia.edu	Christine Lescop	Université Joseph Fourier lescop@ujf-grenoble.fr
Steven Boyer	Université du Québec à Montréal cohf@math.rochester.edu	Robert Lipshitz	University of Oregon lipshitz@uoregon.edu
Tara E Brendle	University of Glasgow tara.brendle@glasgow.ac.uk	Norihiko Minami	Yamato University minami.norihiko@yamato-u.ac.jp
Indira Chatterji	CNRS & Univ. Côte d'Azur (Nice) indira.chatterji@math.cnrs.fr	Andrés Navas	Universidad de Santiago de Chile andres.navas@usach.cl
Alexander Dranishnikov	University of Florida dranish@math.ufl.edu	Robert Oliver	University Paris 13 bobol@math.univ-paris13.fr
Tobias Ekholm	Uppsala University, Sweden tobias.ekholm@math.uu.se	Jessica S Purcell	Monash University jessica.purcell@monash.edu
Mario Eudave-Muñoz	Univ. Nacional Autónoma de México mario@matem.unam.mx	Birgit Richter	Universität Hamburg birgit.richter@uni-hamburg.de
David Futер	Temple University dfuter@temple.edu	Jérôme Scherer	École Polytech. Féd. de Lausanne jerome.scherer@epfl.ch
John Greenlees	University of Warwick john.greenlees@warwick.ac.uk	Vesna Stojanoska	Univ. of Illinois at Urbana-Champaign vesna@illinois.edu
Ian Hambleton	McMaster University ian@math.mcmaster.ca	Zoltán Szabó	Princeton University szabo@math.princeton.edu
Matthew Hedden	Michigan State University mhedden@math.msu.edu	Maggy Tomova	University of Iowa maggy-tomova@uiowa.edu
Hans-Werner Henn	Université Louis Pasteur henn@math.u-strasbg.fr	Chris Wendl	Humboldt-Universität zu Berlin wendl@math.hu-berlin.de
Daniel Isaksen	Wayne State University isaksen@math.wayne.edu	Daniel T Wise	McGill University, Canada daniel.wise@mcgill.ca
Thomas Koberda	University of Virginia thomas.koberda@virginia.edu	Lior Yanovski	Hebrew University of Jerusalem lior.yanovski@gmail.com
Markus Land	LMU München markus.land@math.lmu.de		


See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2025 is US \$760/year for the electronic version, and \$1110/year (+\$75, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

AGT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing

<https://msp.org/>

© 2025 Mathematical Sciences Publishers

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 25 Issue 1 (pages 1–644) 2025

Cutting and pasting in the Torelli subgroup of $\text{Out}(F_n)$	1
JACOB LANDGRAF	
Hyperbolic groups with logarithmic separation profile	39
NIR LAZAROVICH and CORENTIN LE COZ	
Topology and geometry of flagness and beltiness of simple handlebodies	55
ZHI LÜ and LISU WU	
Property (QT) for 3-manifold groups	107
SUZHEN HAN, HOANG THANH NGUYEN and WENYUAN YANG	
On positive braids, monodromy groups and framings	161
LIVIO FERRETTI	
Highly twisted diagrams	207
NIR LAZAROVICH, YOAV MORIAH and TAL PINSKY	
Rational homology ribbon cobordism is a partial order	245
STEFAN FRIEDL, FILIP MISEV and RAPHAEL ZENTNER	
A cubulation with no factor system	255
SAM SHEPHERD	
Relative h -principle and contact geometry	267
JACOB TAYLOR	
Relations amongst twists along Montesinos twins in the 4-sphere	287
DAVID T GAY and DANIEL HARTMAN	
Complexity of 3-manifolds obtained by Dehn filling	301
WILLIAM JACO, JOACHIM HYAM RUBINSTEIN, JONATHAN SPREER and STEPHAN TILLMANN	
The enumeration and classification of prime 20-crossing knots	329
MORWEN B THISTLETHWAITE	
An exotic presentation of $\mathbb{Z} \times \mathbb{Z}$ and the Andrews–Curtis conjecture	345
JONATHAN ARIEL BARMAN	
Generalizing quasicategories via model structures on simplicial sets	357
MATT FELLER	
Quasiconvexity of virtual joins and separability of products in relatively hyperbolic groups	399
ASHOT MINASYAN and LAWK MINEH	
Mapping tori of \mathcal{A}_∞ -autoequivalences and Legendrian lifts of exact Lagrangians in circular contactizations	489
ADRIAN PETR	
Infinite-type loxodromic isometries of the relative arc graph	563
CAROLYN ABBOTT, NICHOLAS MILLER and PRIYAM PATEL	