Algebraic and geometric properties of big mapping class groups

Anschel Schaffer-Cohen

A DISSERTATION

in

Mathematics

Presented to the Faculties of the University of Pennsylvania in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy

2022

Supervisor of Dissertation

David Futer, Professor of Mathematics

Graduate Group Chairperson

Ron Donagi, Graduate Chair

Dissertation Committee:
Julia Hartmann, Professor of Mathematics
David Futer, Professor of Mathematics
Wolfgang Ziller, Professor of Mathematics Emeritus

Acknowledgments

> The acknowledgments is the only part of a book no one edits and no one reviews. It's rather strange, actually.
>
> Dylan Schaffer, *Life, Death & Bialys: A Father/Son Baking Story*

For the research in Chapters 2 and 3 I am endebted chiefly to conversations with Priyam Patel, Nicholas G. Vlamis, Federica Fanoni, and of course my advisor Dave Futer. Kathryn Mann and Kasra Rafi, in addition to producing the seminal results that make the work in Chapter 3 possible, have also been patient and helpful in responding to my follow-up questions about their paper. Edgar Bering provided an insightful email during my work on Chapter 2, and Rylee Lyman suggested the $\natural$ notation used in Proposition 3.3.5. Some of this work was done at the Young Geometric Group Theory IX conference, which was partially funded by

NSF grant 1941077.

That said, it feels a bit remiss to credit only the people who helped me with this specific research. I wouldn't have gotten to the "compiling a PhD thesis" stage in the first place without a lot of help along the way. At the very least that support includes my parents, Nancy Cohen and Cullen Schaffer, who provided me with the love, inspiration, stability, and support that got me to the adulthood I have today. Support has also been provided, in abundance, by my partners and closest friends. To Adam Clearwater, Alice Sukhina, Allison Schinagle, Cori Croot, Ellie Brooks, Janet Nguyen, Man Cheung Tsui, Marcus Michelin, Marti Cohen, and Sara-Rivka Bass: none of this would have been possible without you. Thank you. My life would also have been incomplete without some important nonhuman companions. Yelka, Darran, and Carlos, your existence has brightened my life, and I will never forget you.

At this juncture a land acknowledgement would also be appropriate, but here we come upon a unique problem. Unlike nearly all land currently under the control of the US government, the city of Philadelphia was acquired through entirely peaceful means, but historical evidence of what land exactly was ceded to William Penn by the local Lenape people has been somewhat mysteriously lost. It is entirely possible, as has been seriously suggested by Jennings [JDD17], that the treaty was deliberately destroyed by Penn's unscrupulous heirs; the idea of its being lost through carelessness does not exactly suggest an attitude of respect toward Indian land. In

either case, it is impossible to say for certain whether the land currently occupied by the University of Pennsylvania was acquired ethically.

What can be said with certainty is that this university proudly traces its routes back to the "Academy of Philadelphia" founded around 1750 by Benjamin Franklin, and that one of Mr. Franklin's primary political projects in the 1750s was to end the Pennsylvania Colony's decades-long policy of strict pacifism in favor of active military support for white "settlers" stealing land from the Lenape. In this respect, if not in others, the University of Pennsylvania has fully lived up to its founder's legacy.

That isn't the brightest note to end on, so I'll turn back to math. It seems clear to me that every math class I've taken has in some way contributed to my reaching this point, going back at least to the time my father taught me symbolic logic before I could read. To that end, I've attempted to list, in chronological order, every teacher who taught me something that might be described as a "math class". Thank you to Ward Bronson, Claire Nestor, Jack Coggins, Philip Grossi, Nancy Couturie, Zbigniew Nitecki, Lenore Cowen, Christoph Borgers, Anselm Blumer, Genevieve Walsh, Loring Tu, Moon Duchin, George McNinch, Alberto López Martín, Fulton Gonzalez, Greta Panova, Herman Gluck, Márton Hablicsek, Robert Strain, Ted Chinburg, Scott Weinstein, Jonathan Block, Samuel J. Taylor, Matthew Stover, and Henry Towsner. Every one of you taught me something I'm glad to have learned.

ABSTRACT

Algebraic and geometric properties of big mapping class groups

Anschel Schaffer-Cohen

David Futer, Advisor

This thesis investigates mapping class groups of infinite-type surfaces, also called big mapping class groups, by studying their actions on certain graphs whose vertices are arcs and curves on the underlying surface. In particular, we show that the extended mapping class group of any surface with a finite, positive number of punctures is isomorphic to the relative arc graph of that surface; that the mapping class group of any translatable surface is quasi-isometric to that surface's translatable curve graph; and that the mapping class group of a sphere minus a Cantor set is quasi-isometric to that surface's loop graph.

Contents

Chapter 1

Introduction

> You'll be sort of surprised what
> there is to be found
> Once you go beyond Z and start
> poking around!
>
> Dr. Seuss, *On Beyond Zebra*

1.1 Surfaces big and small

A *surface* is a 2-manifold, i.e. a topological space locally homeomorphic to the Euclidean plane. For the purpose of this work, we also assume that a surface is Hausdorff, separable, and orientable. The first two conditions exclude examples generally thought of as pathological—the plane with two origins, or an interval crossed with the long line. Non-orientable surfaces are somewhat more socially acceptable, but we follow the tradition in low-dimensional topology of limiting our

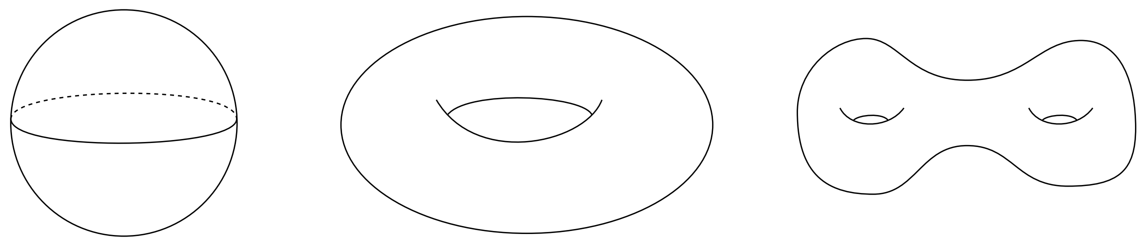

Figure 1.1: Compact surfaces of genus 0, 1, and 2.

study to orientable manifolds, with the non-orientable case set aside for further investigation.

Compact surfaces admit an extremely simple classification.[1] Namely, a compact surface is completely classified by its genus; see Figure 1.1. We might also want to puncture our surface, in the following sense. If we remove finitely many points from a compact surface, as in Figure 1.2, the resulting surfaces are "almost compact", in that we get our compact manifold back after adding finitely many points. This motivates the following definition.

Definition 1.1.1. A surface Σ has *finite type* if there is a compact surface Σ' and a (possibly empty) finite set $P \subseteq \Sigma'$ such that Σ is homeomorphic to $\Sigma' \setminus P$.

Note that we can classify such a surface by giving two non-negative integers: the genus of the surface Σ' and the size of the set P. Two equivalent definitions of finite

[1] It's not immediately obvious who to cite here. Gallier and Xu [GX13] credit Möbius [Mö61] with the initial classification, and Brahana [Bra21] with the first "rigorous proof"—I have not independently verified this claim. David Futer, meanwhile, recommends I refer to Conway's ZIP proof, as described by Francis and Weeks [FW99]. In any case, the classification is well-known and straightforward.

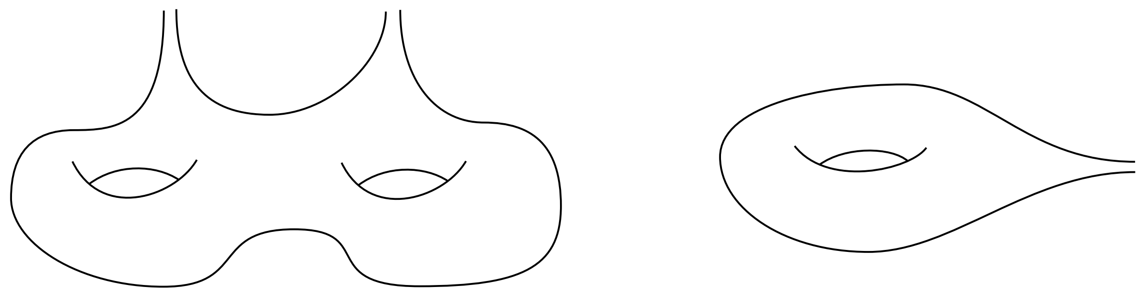

Figure 1.2: Some punctured surfaces.

type, which may help the reader's intuition, are that

- a surface has finite type if and only if it admits a complete, finite-volume, constant-curvature Riemannian metric; and
- a surface has finite type if and only if its fundamental group is finitely generated.

Surfaces that do not have finite type are said to have *infinite type*—some examples of infinite-type surfaces are given in Figure 1.3. Note however that when trying to give "examples" of infinite-type surfaces, we run into the same problem that plagues discussions of any uncountable collection: there are too many of them out there. The collection of infinite-type surfaces that *can be described* is of course countable. In particular, the surfaces we explicitly draw, name, or even comprehend will inevitably be more regular in structure than those we do not. Thus our theorems need to account for a level of weirdness that is present in "nearly all" infinite-type surfaces, but that will never appear in our examples.

The classification of these surfaces has a bit more to it than just counting genus and punctures. We start by defining the *space of ends* of a surface.

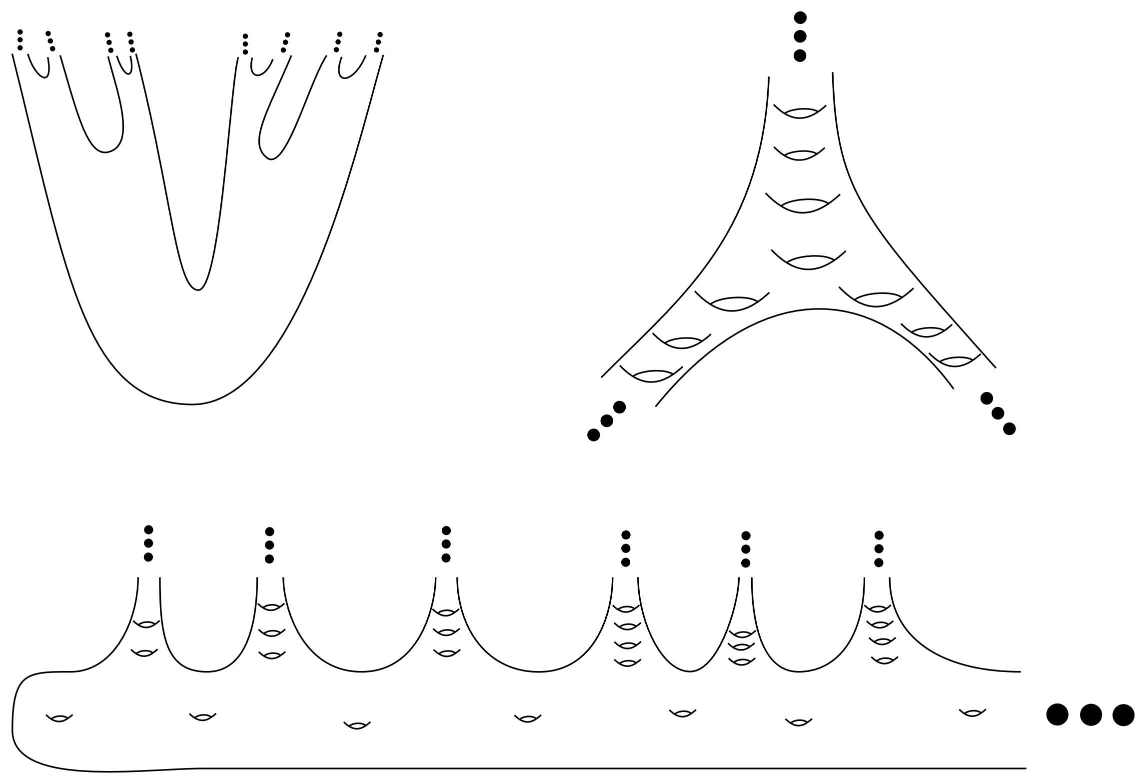

Figure 1.3: Some infinite-type surfaces.

Definition 1.1.2. Given a surface Σ, an *end* of Σ is the equivalence class of a nested sequence of connected subsurfaces $S_1 \supseteq S_2 \supseteq \cdots$ of Σ, each with compact boundary and with the property that for any compact subsurface $K \subseteq \Sigma$, $K \cap S_n = \emptyset$ for high enough n. Two such sequences $S_1 \supseteq S_2 \supseteq \cdots$ and $T_1 \supseteq T_2 \supseteq \cdots$ are equivalent if for every $n \in \mathbb{N}$ there exists $m \in \mathbb{N}$ such that $T_m \subseteq S_n$ and vice versa.

An end x given by a sequence $S_1 \supseteq S_2 \supseteq \cdots$ is said to be *accumulated by genus* if every S_n has positive genus. Otherwise x is said to be *planar*.

The *space of ends* of Σ, written $E(\Sigma)$, is a topological space whose points are the ends of Σ and whose basic open sets correspond to subsurfaces $S \subseteq \Sigma$ with compact boundary. An end $S_1 \supseteq S_2 \supseteq \cdots$ of Σ is contained in the basic open set corresponding to S if $S_n \subseteq S$ for high enough n.

By construction, $E(\Sigma)$ is compact, separable, and totally disconnected—in other words, it is homeomorphic to a closed subspace of a Cantor space. The set of ends accumulated by genus is written $E_G(\Sigma)$ and is always a closed subspace of $E(\Sigma)$. We normally think of (E, E_G) as a pair, and call two pairs (E, E_G) and (E', E'_G) homeomorphic if there exists a homeomorphism $\varphi\colon E \to E'$ such that $\varphi(E_G) = E'_G$.[2]

With this data in hand we can move on to the classification. This claim was first stated by Kerékjártó [Ker23] and was proved by Richards [Ric63].

Theorem 1.1.3 (The principal result of [Ric63])**.** *An infinite-type surface is classified*

[2]Note that this is a stricter condition than "E is homeomorphic to E', and E_G is homeomorphic to E'_G".

by its genus (if finite), its space of ends E, and the subset $E_G \subseteq E$ of ends accumulated by genus. What's more, any homeomorphism of the pair (E, E_G) extends to a homeomorphism of the underlying surface.

Note that the genus of a surface is finite if and only if E_G is empty.

1.2 Mapping class groups

One way that mathematicians—and especially group theorists—study an object is via its symmetries. For a topological object like a surface, that means looking at the group of homeomorphisms, usually written $\mathrm{Homeo}(\Sigma)$. In our case, when the surface is orientable, we might concentrate on the group of orientation-preserving homeomorphisms, $\mathrm{Homeo}^+(\Sigma)$, which is a subgroup of $\mathrm{Homeo}(\Sigma)$ with index 2. Either group can be endowed with the compact-open topology, giving us an unreasonably large[3] topological group with many path-components.

A path in $\mathrm{Homeo}^+(\Sigma)$ is a continuous one-parameter family of homeomorphisms—that is, an isotopy. Thus two homeomorphisms are in the same path-component if and only if they are isotopic,[4] and $\mathrm{Homeo}_0(\Sigma)$, the path-component of the identity, is the subgroup of homeomorphisms isotopic to the identity. This subgroup is closed and normal, so we can quotient by it to get the *mapping class group*

[3]For instance, $\mathrm{Homeo}^+(\Sigma)$ his infinite covering dimension.

[4]Thanks to the work of Epstein [Eps66], every homotopy on a surface can be promoted to an isotopy. Thus two homeomorphisms are also in the same path-component if and only if they are homotopic, and we can use the concepts of homotopy and isotopy largely interchangeably.

$\mathrm{MCG}(\Sigma) := \mathrm{Homeo}^+(\Sigma)/\,\mathrm{Homeo}_0(\Sigma)$,[5] which we can think of as the group of orientation-preserving homeomorphisms of Σ, considered up to isotopy. If we do the same thing with $\mathrm{Homeo}(\Sigma)$ instead of $\mathrm{Homeo}^+(\Sigma)$, we get the *extended* mapping class group $\mathrm{MCG}^*(\Sigma)$, of which $\mathrm{MCG}(\Sigma)$ is an index-2 subgroup. We call the mapping class group of an infinite-type surface a *big* mapping class group.

The mapping class group of a surface thus tells us what kinds of topological symmetries the surface has, ignoring those that arise from isotopy. In addition to the intrinsic value of learning more about the surface, there are also applications elsewhere, for instance in Teichmüller theory, dynamics, and the construction of 3-manifolds.

It can be surprisingly difficult, when first encountering the idea of mapping class groups, to actually picture some interesting mapping classes concretely. The standard example of a mapping class is a *Dehn twist*, which we construct as follows. Pick a simple closed curve α on the surface Σ. Since Σ is orientable, α has a tubular neighborhood homeomorphic to an annulus. Let's call this neighborhood A and identify it with the annulus $\{re^{i\theta} \in \mathbb{C} \mid 1 \leq r \leq 2\}$.

Notice that the map $re^{i\theta} \mapsto re^{i(\theta+2\pi(r-1))}$, illustrated in Figure 1.4, restricts to the identity on the boundary of A but is not homotopic to the identity on the full annulus. Thus we can define a map T_α, called the *Dehn twist about the curve* α, which is the identity on $\Sigma \setminus A$ and takes $re^{i\theta}$ to $re^{i(\theta+2\pi(r-1))}$ on A; see Figure 1.5.

[5]Some algebraic topologists might recognize this quotient construction as $\pi_0(\mathrm{Homeo}^+(\Sigma))$.

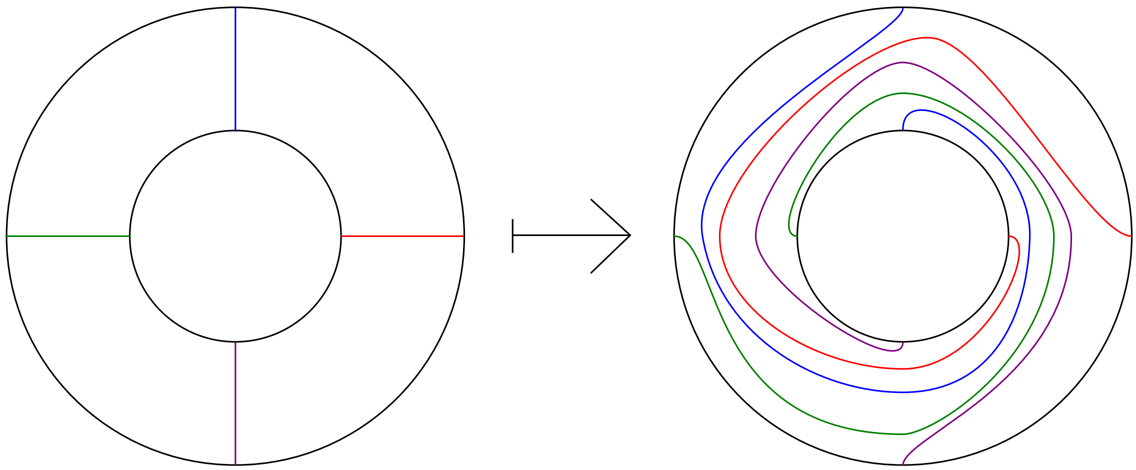

Figure 1.4: A Dehn twist on an annulus.

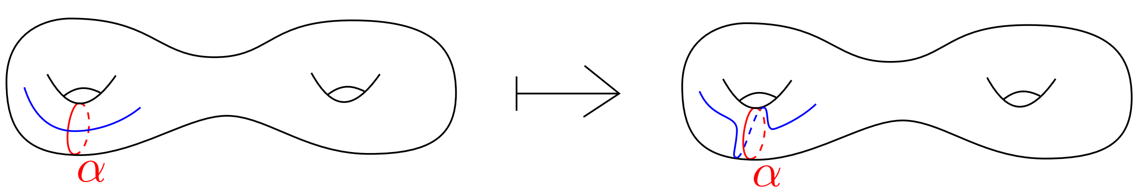

Figure 1.5: A Dehn twist about the curve α on a surface.

As long as α does not bound a disk or a once-punctured disk on Σ, this Dehn twist will not be homotopic to the identity, so T_α represents a nontrivial mapping class.

When the surface Σ has finite type, $\mathrm{Homeo}_0(\Sigma)$ is both closed and open in $\mathrm{Homeo}^+(\Sigma)$, and so the quotient $\mathrm{MCG}(\Sigma)$ is discrete and thus is normally not thought of as a topological group. When Σ has infinite type, on the other hand, $\mathrm{Homeo}_0(\Sigma)$ is closed but not open: fix a compact exhaustion $K_1 \subseteq K_2 \subseteq K_3 \subseteq \cdots$ of Σ, and for each $i \in \mathbb{N}$ let φ_i be a homeomorphism that is the identity on K_i but is not isotopic to the identity on Σ—say, a Dehn twist about a curve disjoint from K_i. Then the sequence φ_i converges to the identity in $\mathrm{Homeo}^+(\Sigma)$ but no φ_i is in $\mathrm{Homeo}_0(\Sigma)$, demonstrating that $\mathrm{Homeo}_0(\Sigma)$ is not open. In particular, this means that the quotient $\mathrm{MCG}(\Sigma)$ is topologically interesting, which will turn out to be important.

1.3 Graphs of curves and arcs

Geometric group theorists generally study groups via their actions by isometries on metric spaces. One large class of metric spaces that mapping class groups can act on is that of graphs of curves or arcs. For our purposes, a *curve* is the isotopy class of an embedded circle in Σ, ignoring orientation. Likewise, an *arc* is the isotopy class of an embedded line in Σ, again ignoring orientation. If both ends of the line limit to the same end of Σ, the arc may be called a *loop*. We generally do not consider homotopically trivial or peripheral cases, i.e. curves that bound a disk or a once-punctured disk, or arcs that bound a monogon. The phrases "all curves" and "all arcs" implicitly exclude these cases.

A *graph of curves or arcs* is a simplicial graph whose vertex set is some subset of the set of curves or arcs on the surface. Usually, the vertex set and the edge relation are chosen so that the mapping class group acts simplicially on the graph.[6] The most commonly encountered graphs of curves or arcs are

- the *curve graph* $\mathcal{C}(\Sigma)$, where the vertex is the set of all curves, and two vertices are connected by an edge if their corresponding curves have disjoint representatives; and
- the *arc graph* $\mathcal{A}(\Sigma)$, where the vertex set is the set of all arcs whose ends limit to punctures, and again two vertices are connected by an edge if their corresponding arcs have disjoint representatives. Note that this is only defined if Σ has at least one puncture.

Both of these are well-studied for finite-type surfaces, and have come under increasing scrutiny in the infinite-type case. Broadly speaking, we can study the action of the mapping class group on such a graph from two angles. Algebraically, we can ask about the relationship between $\mathrm{MCG}(\Sigma)$ and the automorphism group of a graph it acts on; and geometrically, we can ask what the action of $\mathrm{MCG}(\Sigma)$ on the graph as a metric space tells us about the geometry of the group itself.

[6] In the case of infinite-type surfaces, we may also want to ensure that the action is continuous.

1.4 Algebraic properties

Since $\mathrm{MCG}(\Sigma)$ is acting on a graph, it is natural to ask about the induced homomorphism from the mapping class group to the automorphism group of the graph. It is slightly too much to hope that these groups be isomorphic, because an orientation-reversing homeomorphism of Σ, though excluded from $\mathrm{MCG}(\Sigma)$, will induce an automorphism of any reasonable graph of curves or arcs. Thus we focus here on the extended mapping class group $\mathrm{MCG}^*(\Sigma)$, which includes those orientation-reversing maps. In the finite-type case, Ivanov [Iva97], Korkmaz [Kor99], and Luo [Luo00] showed that $\mathrm{MCG}^*(\Sigma)$ is isomorphic to $\mathrm{Aut}(\mathcal{C}(\Sigma))$ for all but a few surfaces, and Irmak and McCarthy [IM10] showed that $\mathrm{MCG}^*(\Sigma)$ is isomorphic to $\mathrm{Aut}(\mathcal{A}(\Sigma))$ whenever that graph is defined.

When Σ has infinite type, Hernández Hernández, Morales, and Valdez [HHMV18] and Bavard, Dowdall, and Rafi [BDR18] showed that $\mathrm{MCG}^*(\Sigma)$ is still isomorphic to $\mathrm{Aut}(\mathcal{C}(\Sigma))$. The analogous result for the arc graph remained open.

Following Bavard [Bav16] and Aramayona, Fossas, and Parlier [AFP17], we focus on the case of an infinite-type surface Σ with a finite, non-empty set of punctures P, and consider the *relative arc graph* $\mathcal{A}(\Sigma, P)$, where vertices are arcs that limit to the punctures in P. If P is the set of all punctures on Σ, the extended mapping class group $\mathrm{MCG}^*(\Sigma)$ has a well-defined action on $\mathcal{A}(\Sigma, P)$. The goal of Chapter 2 is to show that this action does in fact induce an isomorphism as in the finite-type case. Theorem 2.1.2 consists of this result. On the way, we prove Theorem 2.1.1, which focuses on the special case of a surface with a single puncture, or equivalently a single marked point.

1.5 Geometry of groups

It may seem strange, from an outsider's perspective, to ask about the geometric properties of groups. After all, groups do not seem to be inherently geometric objects. But groups can certainly *act on* geometric objects; and if some property holds for every action of a particular group on every geometric object, it makes sense to think of that as a property of the group itself.

Consider, as a first example, a group G, with finite generating set S, acting on a metric space X with distance function d. Pick a point $x_0 \in X$, and let $k = \max_{s \in S} d(x_0, s \cdot x_0)$. Then given $g \in G$, written as $g = s_1 s_2 \cdots s_n$ with each $s_i \in S$, we can immediately see that $d(x_0, g \cdot x_0) \leq kn$. Put more formally, and using the notation $|g|_S$ for the minimal word-length of the group element g with respect to the generating set S, we can say that for any $x_0 \in X$, there exists a $k \geq 1$ such that for all $g \in G$,

$$d(x_0, g \cdot x_0) \leq k|g|_S.$$

This strongly suggests the use of $|\cdot|_S$ as the basis for a metric on the group G itself. Remembering that, for a group action, $d(g \cdot x_0, h \cdot x_0) = d(x_0, g^{-1}h \cdot x_0)$, we define $d_S(g, h) = |g^{-1}h|_S$, and now our previous statement takes on an even more geometric flavor. That is, for any $x_0 \in X$, there exists a $k \geq 1$ such that for all $g, h \in G$,

$$d(g \cdot x_0, h \cdot x_0) \leq k d_S(g, h).$$

Or in other words, the orbit-map $g \mapsto g \cdot x_0$ is k-Lipschitz.

To put a lower bound on $d(g \cdot x_0, h \cdot x_0)$ we need to impose some restrictions on the action, for instance to ensure that it is not trivial. To that end, we call the action of a group G on a metric space X *properly discontinuous* if, for every bounded[7] set $A \subseteq X$, the set $\{g \in G \mid g \cdot A \cup A \neq \emptyset\}$ is finite. In this case, we can chose constants $k \geq 1$ and $C \geq 0$ so that

$$d(g \cdot x_0, h \cdot x_0) \geq \frac{1}{k} d_S(g, h) - C.$$

Putting this all together gives us the famous lemma of Schwarz[8] [Š55] and Milnor [Mil68].

Lemma 1.5.1 (Schwarz-Milnor)**.** *Let G be a group generated by a finite set S, acting properly discontinuously on a metric space X. Then for any $x_0 \in X$, there exist*

[7] Most definitions use compactness here instead of boundedness. For proper metric spaces these conditions are equivalent, but many metric spaces that we care about are not proper—curve graphs, for instance, have infinite valence. For that reason we focus on boundedness, which is a stronger condition than compactness in a non-proper metric space.

[8] After [Š55] was published in the Soviet Union in 1955, English citations used a number of different transliterations of the author's Russian-transcribed Yiddish name, leading to a profusion of spellings and pronunciations in English-language references. Since that time, however, Albert Schwarz has immigrated to the United States, and he consistently spells his name in English as used here.

constants $k \geq 1$, $C \geq 0$ such that for any $g, h \in G$,

$$\frac{1}{k} d_S(g, h) - C \leq d(g \cdot x_0, h \cdot x_0) \leq k d_S(g, h) + C.$$

The inequality in Lemma 1.5.1 feels like an equivalence relation, but it is not symmetric: the orbit $G \cdot x_0$ might cover only one small region of X, making an inverse relationship impossible. To close this gap, we can insist that the orbit $G \cdot x_0$ get close to every point of X. In other words, we say that the action is *co-bounded*[9] if there is some bounded set $A \subseteq X$ such that $\bigcup_{g \in G} g \cdot A = X$. Finally, we tie this up by introducing a name for this equivalence relation.

Definition 1.5.2. Given two metric spaces X and Y, a map $f \colon X \to Y$ is a *quasi-isometric embedding* if there exist constants $k \geq 1$, $C \geq 0$ such that for any $x_1, x_2 \in X$,

$$\frac{1}{k} d_X(x_1, x_2) - C \leq d_Y(f(x_1), f(x_2)) \leq k d_X(x_1, x_2) + C.$$

If in addition every point of Y is within bounded distance of the image $f(X)$, then f is a *quasi-isometry* and the spaces X and Y are *quasi-isometric*.

Note that quasi-isometry is an equivalence relation. We can now write a shorter and slightly stronger version of Lemma 1.5.1.

Lemma 1.5.3 (Schwarz-Milnor, version 2)**.** *Let G be a group generated by a finite set S and equipped with the metric d_S, acting properly discontinuously and co-*

[9]Again, other definitions usually use co-compactness.

boundedly on a metric space X. Then for any $x_0 \in X$, the orbit map $g \mapsto g \cdot x_0$ is a quasi-isometry.

Crucially, this lets us avoid a problem that was glossed over before. The metric d_S obviously depends on the choice of a generating set S, and thus is not a property of the group itself per se. But given two finite generating sets S and T, the conditions of Lemma 1.5.3 apply to the action of G on itself by left-multiplication. Taking G as the group, S as the generating set, G with d_T as the metric space, and the identity as x_0, we see that the identity map from G to itself (but changing generating sets) is a quasi-isometry. Thus we can say that a finitely generated group has a well-defined metric *up to quasi-isometry*.

1.5.1 Topological groups and coarse boundedness

The previous section was written from the perspective of finitely generated groups, and this is the angle from which geometric properties of this kind have traditionally been studied. Indeed, the mapping class group of a finite-type surface was shown by Dehn [Deh38] to be generated by a finite collection of Dehn twists, so it makes sense to talk about the geometry of a classical mapping class group in this sense.

Big mapping class groups, on the other hand, are uncountable and therefore definitely not finitely generated. On the other hand, they have a topological structure that their finite-type analogues lack. To what extent can we replicate the classical results in this new environment of uncountable, topological groups?

The key insight, due to Rosendal [Ros18], is that the finiteness of the generating set in Lemma 1.5.3 and its antecedents is largely a means to an end. The only property of finiteness actually used is the fact that $k = \max_{s \in S} d(x_0, s \cdot x_0)$ exists. Indeed, for any generating set S, if we know that $d(x_0, s \cdot x_0) \leq k$ for all $s \in S$, it is still true that

$$d(g \cdot x_0, h \cdot x_0) \leq k d_S(g, h)$$

for all $g, h \in G$. This inspires a new definition, which can replace finiteness in many theorems of geometric group theory.

Definition 1.5.4. A subset A of a topological group G is *coarsely bounded* if for every continuous action of G on a metric space X and every $x_0 \in X$, the set $A \cdot x_0 = \{a \cdot x_0\}_{a \in A}$ is bounded in X.

Note that any finite set is coarsely bounded, as indeed is any compact set, so this framework forms an expansion of the existing literature about finitely or compactly generated groups. In particular, by arguments analogous to those above we can see that two different coarsely bounded generating sets for the same group give rise to quasi-isometric word metrics, so we can talk rigorously about the "geometric properties" of such groups.

Definition 1.5.4 may at first glance seem very hard to apply due to the broadness of the quantification over "every continuous action". How do we actually prove that a (non-compact) set is coarsely bounded? The most useful tool here is Lemma 3.2.1, which gives a more concrete group-theoretic condition for coarse boundedness assuming the group is Polish. Thankfully, big mapping class groups are in fact Polish, and we use this lemma extensively.

1.6 Geometric properties

Sadly, not all big mapping class groups admit coarsely bounded generating sets. Mann and Rafi [MR19] classified those that do, but did not provide any further insight into the geometric properties of these groups. Indeed, mapping class groups—and especially big mapping class groups—can be difficult to picture. The traditional solution to this problem has been to study not the mapping class group itself, but its action on a metric space, often a graph of curves or arcs. Masur and Minsky [MM00] showed that the word length in the mapping class group of a finite-type surface can be estimated using the action of that group on a collection of curve graphs, thus making it possible to study the geometry of the mapping class group by focusing entirely on the geometry of these graphs.

In the infinite-type case, we might hope to do even better. Ideally, if the action of the mapping class group on some graph of curves or arcs satisfied conditions similar to those of Lemma 1.5.3, then we would have a true quasi-isometry, and all the geometric information about the mapping class group would be contained in the geometry of that one graph. Crucially, all graphs of curves and most[10] graphs of

[10]Things can start to get uncountable if, for instance, we let our arcs limit to ends that are not punctures, or wrap around forever without limiting to an end at all. The graphs of arcs described

arcs are countable, unlike the groups that are acting on them, which makes them much simpler to work with.

To this end, we introduce Lemma 3.2.2, which is a simplified version of Lemma 1.5.3 for the case of groups acting vertex-transitively on graphs in the setting of coarse boundedness. With this tool in hand, we prove two new quasi-isometries.

Section 3.3 introduces a class of *translatable* surfaces, and Section 3.4 defines a graph of curves, the *translatable curve graph*, that can be built for any translatable surface. Finally Theorem 3.5.3 proves two important relationships between the mapping class group and the translatable curve graph. First, that as we had hoped the mapping class group of a translatable surface is in fact quasi-isometric to that surface's translatable curve graph. And second, that this is the only way a mapping class group can be quasi-isometric to a graph of curves. That is, if the mapping class group of a surface is quasi-isometric to *any* graph of curves, that surface must be translatable.

This result is in some sense a classification, but it still leaves us to wonder about the geometry of translatable curve graphs themselves. A less general but more immediately fruitful result comes in Section 3.6, which studies the plane minus a Cantor set, also called the *punctured Cantor tree*. As a surface with a single puncture, this is a case where the loop graph is well-defined, and Theorem 3.6.5 shows that the mapping class group of this surface is in fact quasi-isometric to its

here are all countable, however, as their arcs do limit to punctures.

loop graph.

This result is not a full classification—it is not clear at present whether there are other surfaces whose mapping class groups are quasi-isometric to loop or arc graphs—but the loop graph, unlike the translatable curve graph, is already somewhat well-studied. In particular, Bavard [Bav16] showed that it is δ-hyperbolic, and Bavard and Walker [BW18] described its Gromov boundary. Since δ-hyperbolicity and the homeomorphism type of the Gromov boundary are both quasi-isometry invariants, these results combine with Theorem 3.6.5 to tell us that this mapping class group is itself δ-hyperbolic, and to describe its Gromov boundary. What's more, we can also use Theorem 3.6.5, along with work of Cornulier and de la Harpe [CdlH16], to show that this mapping class group has a *coarsely bounded presentation*; that is, a group presentation in which the generating set is coarsely bounded and the relators have bounded length.

Chapter 2

Automorphisms

Alguien observará que la conclusión precedió sin duda a las "pruebas". ¿Quién se resigna a buscar pruebas de algo no creído por él o cuya prédica no le importa?

Jorge Luis Borges, "Tres versiones de Judas"

2.1 Introduction and Main Result

The goal of this chapter is to prove the isomorphisms stated in Section 1.4 between the extended mapping class group $\mathrm{MCG}^*(\Sigma)$ of a surface Σ and the loop or arc graph on Σ. The material here is lightly adapted from [SC22].

Fix a basepoint $p \in \Sigma$. For the purposes of this chapter, we consider a *loop* in Σ based at p to be an unoriented simple closed curve in Σ starting an ending at p, considered up to isotopy relative to p. The *loop graph* $\mathcal{L}(\Sigma, p)$ of Σ is the graph whose vertex set is the set of all such loops, with vertices connected by an edge if they have representatives intersecting only at p. The group that acts on this graph is the *extended based mapping class group* $\mathrm{MCG}^*(\Sigma, p)$, which is defined analogously to the extended mapping class group introduced in Section 1.2 but with all homeomorphisms and isotopies fixing the basepoint p. By construction, $\mathrm{MCG}^*(\Sigma, p)$ acts on $\mathcal{L}(\Sigma, p)$, giving a homomorphism $\mathrm{MCG}^*(\Sigma, p) \to \mathrm{Aut}(\mathcal{L}(\Sigma, p))$. Theorem 2.1.1 will show that this is in fact an isomorphism.

Theorem 2.1.1. *Given an infinite-type surface Σ and a basepoint $p \in \Sigma$, the map $\mathrm{MCG}^*(\Sigma, p) \to \mathrm{Aut}(\mathcal{L}(\Sigma, p))$ induced by the action is an isomorphism.*

Irmak and McCarthy [IM10] provide a successful program for proving this theorem in the finite-type case:

1. Fix a maximal set of disjoint loops, called a *triangulation*, and observe that its complementary regions are all triangles.

2. Show that some important local properties—three loops bounding a triangle, two triangles being adjacent, etc.—can be defined in terms of the loop graph and are therefore preserved by automorphisms of that graph. We use some of these results as Facts 2.4.1 and 2.4.2.

3. Use these local properties to construct a homeomorphism inducing a given transformation of our fixed triangulation.

4. Show that a homeomorphism which fixes a triangulation actually fixes the entire loop graph. We use this result as Fact 2.4.3.

In our extension to the infinite-type case we have an advantage, a disadvantage, and a trick. The advantage is of course that we can depend on the existing result for finite-type surfaces. The disadvantage is that "triangulations" in the infinite-type setting can be much more exotic, as seen in Section 2.2. They will in fact have some complementary regions that are not actually triangles. The trick is to notice that Irmak and McCarthy's proof is more general than the result requires: they start by fixing an arbitrary triangulation, but for their proof (and ours) it is sufficient to follow this program with any single triangulation. Thus we can construct a particularly useful triangulation for the specific purpose of building our homeomorphism.

In Section 2.5, we generalize Theorem 2.1.1 to the relative arc graph $\mathcal{A}(\Sigma, P)$ as defined in Section 1.4 with respect to a finite set of punctures P. The *extended mapping class group* $\mathrm{MCG}^*(\Sigma, P)$ of (possibly orientation-reversing) homeomorphisms stabilizing (but not necessarily fixing) the set P acts on $\mathcal{A}(\Sigma, P)$. Then we can prove the following:

Theorem 2.1.2. *Given an infinite-type surface Σ and a finite set P of punctures of Σ, the map $\mathrm{MCG}^*(\Sigma, P) \to \mathrm{Aut}(\mathcal{A}(\Sigma, P))$ induced by the action is an isomorphism.*

There are two special cases to keep in mind: first, if P consists of a single puncture p, then $\mathcal{A}(\Sigma, P) = \mathcal{L}(\Sigma \cup \{p\}, p)$, so Theorem 2.1.2 really is a generalization of Theorem 2.1.1. Second, if Σ has finitely many punctures, then P can contain all of them and so $\mathrm{MCG}^*(\Sigma, P) = \mathrm{MCG}^*(\Sigma)$ and we have a bona fide action of the mapping class group.

The idea of the proof of Theorem 2.1.2 will be to reduce it to that of Theorem 2.1.1 by picking a puncture $p \in P$ and considering $\mathcal{L}(\Sigma \cup \{p\}, p)$ as an induced subgraph of $\mathcal{A}(\Sigma, P)$. The main hurdle will therefore be to prove that an automorphism of the arc graph preserves properties like "this arc is actually a loop" and "these two loops are based at the same puncture".

2.2 Triangulations of infinite-type surfaces

Irmak and McCarthy [IM10] define a triangulation as a maximal set of disjoint arcs, or equivalently as a set of arcs whose complementary regions are all triangles. Hatcher [Hat91], allowing for punctures not in P, observes that the complementary regions of a triangulation may also include punctured monogons. In the infinite-type case triangulations are considerably more exotic, as we shall see. So we simply define a *triangulation*[1] to be a maximal clique in $\mathcal{A}(\Sigma, P)$. The following facts are then immediate:

[1]This somewhat misleading name is (we hope) justified by the fact that it is a straightforward extention of the usage in [IM10] and [Hat91].

Lemma 2.2.1. *Given a set T of arcs and an automorphism f of $\mathcal{A}(\Sigma, P)$, T is a triangulation if and only if $f(T)$ is a triangulation.*

Lemma 2.2.2. *Any set of disjoint arcs can be extended to a triangulation.*

Proof. This follows from Zorn's Lemma. □

A finite-type surface admits countably many triangulations; each triangulation has the same number of loops depending only on the Euler characteristic of the surface; and any two triangulations are connected by a finite sequence of *elementary moves*, in which an arc α is removed from the triangulation and replaced with β, where the geometric intersection number $\iota(\alpha, \beta) = 1$ [Hat91]. An infinite-type surface, on the other hand, admits uncountably many triangulations by application of Lemma 2.2.2, each with countably many arcs, and so most pairs of triangulations will not be connected by elementary moves.

The topology of a triangulation is also interesting in the infinite-type setting: consider a finite union C of small circles centered at each $p \in P$ that intersects each arc finitely many times. By compactness, there must be points on C that are the limits of points on distinct arcs of the triangulation. A priori there is no reason to assume these limit points are contained in the triangulation at all, and so the triangulation may not be closed—in fact, we conjecture that it never is.

These and other questions about the nature of triangulations in general will not be studied further in this paper. However, they suggest possible future areas of research, and motivate the construction, in the next section, of a special triangulation to overcome the potential pitfalls of an arbitrary one.

2.3 Building a useful triangulation

For Sections 2.3 and 2.4 we fix an infinite-type surface Σ, a basepoint p, and an automorphism $f\colon \mathcal{L}(\Sigma, p) \to \mathcal{L}(\Sigma, p)$.

In Section 2.4, we will need a triangulation $\mathcal{T}$ of Σ and a corresponding exhaustion of Σ by compact subsurfaces $\{\Sigma_n\}_{n\in\mathbb{N}}$, with the property that for each n, the restriction of $\mathcal{T}$ to Σ_n is a triangulation of Σ_n. Such a triangulation may not a priori exist, which is why we construct it explicitly in this section. We will start by building an embedded tree in $\mathcal{S} \subseteq \Sigma$, called the *skeleton* of Σ, that has exactly one infinite branch for each end of Σ.

Fix a pants decomposition[2] $\{\Pi_0, \Pi_1, \ldots\}$ of Σ, with $p \in \Pi_0$ and such that $\bigcup_{i=0}^{n} \Pi_i$ is connected for each n. Note that, since Σ may have punctures, the holes in a pair of pants may all be boundary components connected to other pairs of pants, or one or two of them may be punctures. For the purpose of constructing the skeleton $\mathcal{S}$ we distinguish between four types of pairs of pants: namely, for each Π_i we can ask how many of its boundary components connect to some Π_j for $j < i$. In our inductive definition of $\mathcal{S}$, we will refer to such a boundary component as *already connected.*

[2]By "pants decomposition" I mean a collection of pairwise disjoint subsurfaces of Σ, each homeomorphic to a pair of pants, such that the closure of the union of these subsurfaces is Σ.

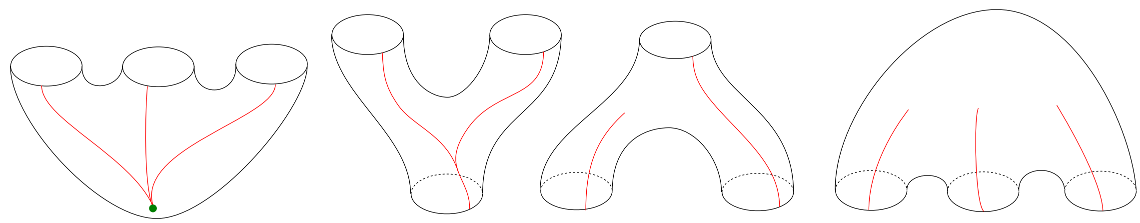

Figure 2.1: The construction of the skeleton on each type of pair of pants Π_i. The green dot is the basepoint p, and the red curves are pieces of the skeleton.

In Figure 2.1 we show how to draw the skeleton on each of the four types of pairs of pants; the first type, with no boundary components already connected, occurs only for Π_0. Note that by construction $\mathcal{S}$ is indeed a tree, and its infinite based rays are in bijection with the ends of Σ; there are also some paths in $\mathcal{S}$ that end after a finite distance, which do not therefore correspond to ends of Σ.

The loops of $\mathcal{T}$ will be defined with reference to the skeleton in the following way: when we draw an arc between two points a and b on the skeleton, such that the arc that does not otherwise intersect the skeleton, we are indicating the loop that starts at p, takes the unique nonbacktracking path to the point a, follows the arc to the point b, and takes the unique nonbacktracking path from b back to p; see Figure 2.2 for some examples. Note that two such arcs, if disjoint, indicate disjoint loops.

We now draw the loops of our triangulation on each pair of pants according to the arcs pictured in Figures 2.3–2.5. The skeleton in a pair of pants with zero boundary components already connected is identical to the skeleton in a pair of pants with one boundary component already connected, so both of these cases are covered in Figure

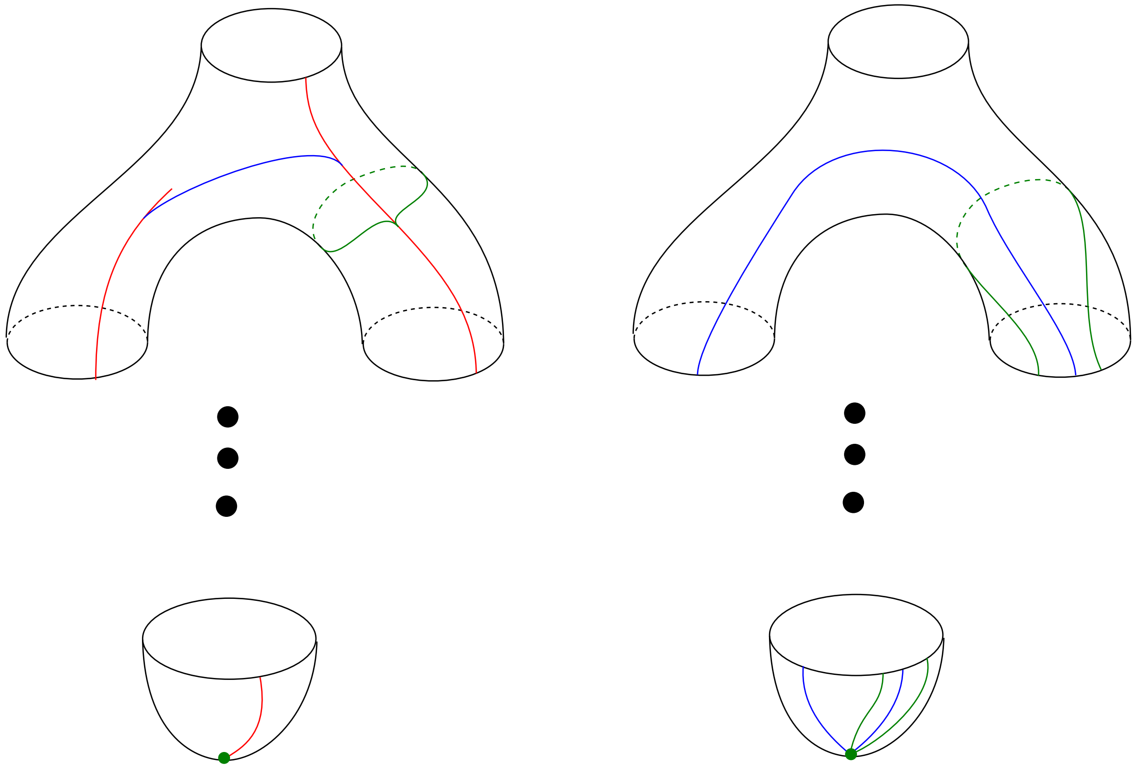

Figure 2.2: Each arc with endpoints on the skeleton represents a loop based at p.

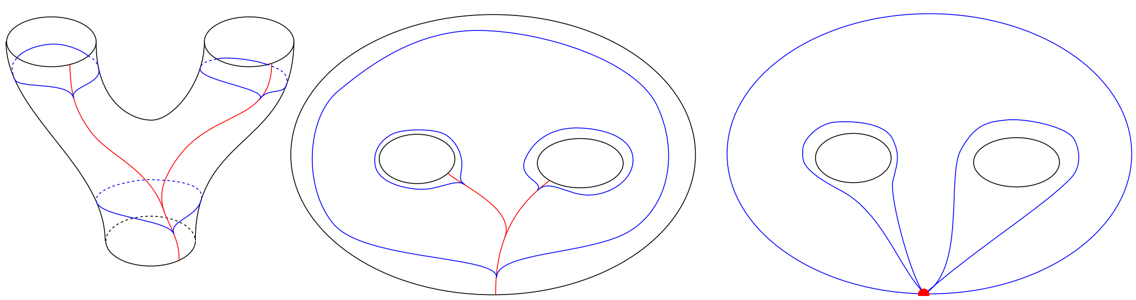

Figure 2.3: The skeleton and loops on a pair of pants with at most one boundary component already connected. The first and second figures are homeomorphic; in the third figure, the skeleton has been contracted to a point so that the triangle decomposition is more clearly visible.

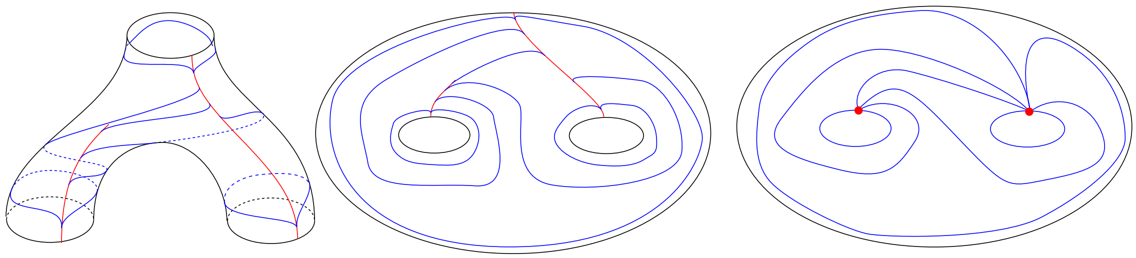

Figure 2.4: The skeleton and loops on a pair of pants with two boundary components already connected. As in Figure 2.3, the first and second figures are homeomorphic; in the third figure, the skeleton has been contracted to two points.

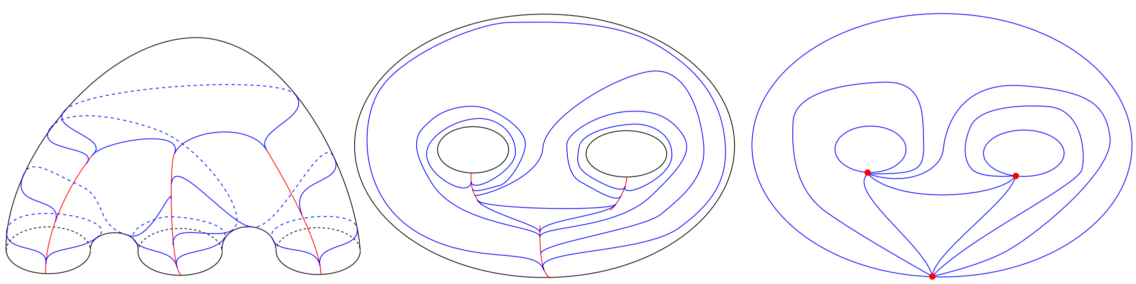

Figure 2.5: The skeleton and loops on a pair of pants with three boundary components already connected. As in Figure 2.3, the first and second figures are homeomorphic; in the third figure, the skeleton has been contracted to three points.

2.3, with the other two cases covered in Figures 2.4 and 2.5 respectively. Note that some of these loops are redundant, as representatives of the same homotopy class may be drawn on more than one pair of pants.

Although we have defined our loops in terms of arcs with endpoints on $\mathcal{S}$, this is purely a notational convenience. After this section we will consider only the loops themselves, and ignore the arcs with which they were defined.

Remark 2.3.1. It turns out, somewhat surprisingly, that we will not need to know whether $\mathcal{T}$ is in fact a triangulation in the sense of Section 2.2. But the reader can easily verify that it is in fact a maximal set of loops; more importantly, $\mathcal{T}$ restricts to a triangulation on each finite-type subsurface $\bigcup_{i=0}^{n} \Pi_i$.

2.4 Constructing a homeomorphism

We first note a few facts due to Irmak and McCarthy; keep in mind that loops are simply a special case of arcs, and the loop graph a special case of the relative arc graph.

Fact 2.4.1 (Propositions 3.3 and 3.4 of [IM10]). *The condition that three arcs bound a triangle, or that two arcs bound a degenerate triangle,*[3] *is preserved under automorphisms of the arc graph.*

[3]A *degenerate triangle* is one where two sides are formed by the same arc. Note that this is only possible when our arcs go between at least two points, and thus never appears in the loop graph.

Fact 2.4.2 (Propositions 3.5–3.7 of [IM10])**.** *When two triangles are adjacent (i.e. share one or two edges) their relative orientations are preserved under automorphisms of the arc graph.*

Fact 2.4.3 (Proposition 3.8 of [IM10])**.** *If a homeomorphism of a finite-type surface preserves some triangulation up to isotopy, then it induces the identity on the arc graph of that surface.*

The proofs of the first two facts do not depend on the surface being finite-type; they are based exclusively on local properties of the arc graph. Since our surface may have punctures, some loops may bound punctured monogons as well as triangles, and this property will also be preserved:

Lemma 2.4.4. *The condition that a loop bounds a punctured monogon is preserved under automorphisms of the arc graph.*

Proof. If a nontrivial loop λ does not bound a punctured monogon, then $\Sigma \setminus \lambda$ has either at least two ends or at least one handle in each component.[4] In either case, we can draw two disjoint triangles adjacent to λ, which we cannot do if λ does bound a punctured monogon. Since the property of three arcs bounding a triangle is preserved by Fact 2.4.1, so is the property of bounding a punctured monogon. □

In order to leverage Fact 2.4.3 we will need some finite-type surfaces to which it can be applied. For this reason, we construct an exhaustion of Σ by finite-type

[4]There may be one or two components, depending on whether λ is separating.

surfaces Σ_n, with corresponding homeomoprhisms φ_n from Σ_n to an appropriate subsurface of Σ. The natural choice is to let the subsurface $\Sigma_n \subseteq \Sigma$ be the union $\bigcup_{i=0}^{n} \Pi_n$ of the first n pairs of pants. Note that by our choice of pants decomposition, Σ_n is connected and contains p.

However, this definition is somewhat unhelpful for the purpose of constructing φ_n, because there is no obvious choice for the image $\varphi_n(\Sigma_n)$—after all, f will not in general preserve our pants decomposition. But Σ_n has a useful alternate definition. If we let $\mathcal{T}_n$ be the set of loops in $\mathcal{T}$ contained in Σ_n, then Σ_n is also the largest subsurface of Σ (up to isotopy) filled by the loops of $\mathcal{T}_n$.

We now have a natural choice for the image of φ_n: since Σ_n is the largest subsurface filled by $\mathcal{T}_n$, $\varphi_n(\Sigma_n)$ should be the largest subsurface filled by $\{f(\lambda) \mid \lambda \in \mathcal{T}_n\}$. We can in fact construct such a map:

Lemma 2.4.5. *For $n \in \mathbb{N}$, there exists a map $\varphi_n : \Sigma_n \to \Sigma$, a homeomorphism onto its image, such that $\varphi_n(\lambda) = f(\lambda)$ for each $\lambda \in \mathcal{L}(\Sigma, p)$ supported on Σ_n. In addition, these homeomorphisms are compatible: that is, $\varphi_{n+1}|_{\Sigma_n} = \varphi_n$ for all n.*

Proof. For each $n \in \mathbb{N}$, the image of φ_n will be the largest subsurface filled by $f(\mathcal{T}_n) = \{f(\lambda) \mid \lambda \in \mathcal{T}_n\}$; call this subsurface Ω_n. Whenever three loops $\lambda_1, \lambda_2, \lambda_3 \in \mathcal{T}_n$ bound a triangle in Σ, so do their images $f(\lambda_1), f(\lambda_2), f(\lambda_3)$ by Fact 2.4.1. The latter triangle is thus included in Ω_n and can be the image of the former. When two such triangles are adjacent, their orientations are preserved by Fact 2.4.2 and so the homeomorphisms on adjacent triangles can be stitched together after applying some isotopy.

When a loop $\lambda \in \mathcal{T}_n$ bounds a punctured monogon in Σ, so does $f(\lambda)$ by Lemma 2.4.4. Then by construction, the punctured monogon bounded by $f(\lambda)$ is contained in Ω_n, and the two punctured monogons are homeomorphic. This homeomorphism can be stitched to those above along λ and $f(\lambda)$.

When a loop $\lambda \in \mathcal{T}_n$ is parallel to the boundary of Σ_n—that is, λ together with one of the boundary curves of Σ_n bound an annulus—λ must bound a triangle in $\mathcal{T}$ with at least one side not in $\mathcal{T}_n$. By Fact 2.4.1, $f(\lambda)$ therefore bounds a triangle in $f(\mathcal{T})$ with at least one side not in $f(\mathcal{T}_n)$ and so this triangle is not included in Ω_n. Thus $f(\lambda)$ is parallel to the boundary of Ω_n, and we can extend our homeomorphism to a tubular neighborhood of λ.

Since Σ_n is made up of tubular neighborhoods of loops in $\mathcal{T}_n$, triangles bounded by loops in $\mathcal{T}_n$, and punctured monogons bounded by loops in $\mathcal{T}_n$—and likewise for Ω_n and loops in $f(\mathcal{T}_n)$—this algorithm gives a homeomorphism $\varphi_n : \Sigma_n \to \Omega_n$. Since $\varphi_n(\lambda) = f(\lambda)$ for each $\lambda \in \mathcal{T}_n$, which is a triangulation of Σ_n, it follows by Fact 2.4.3 that $\varphi_n(\lambda) = f(\lambda)$ for every $\lambda \in \mathcal{L}(\Sigma, p)$ contained in Σ_n.

By construction $\Omega_n \subseteq \Omega_{n+1}$, and the only choices we made in defining our homeomorphisms were isotopies on the interiors of triangles and punctured monogons. Thus after an isotopy, φ_n and φ_{n+1} agree on Σ_n. □

To prove Theorem 2.1.1 we need to combine these partial maps φ_n, and we also need to prove injectivity. The following lemma achieves the latter result.

Lemma 2.4.6. *If a homeomorphism $\varphi : \Sigma \to \Sigma$ induces the identity automorphism of* $\mathrm{Aut}(\mathcal{L}(\Sigma, p))$ *then φ is isotopic to the identity.*

Proof. The key insight here is that if φ preserves the isotopy class of each based loop it must also preserve the isotopy class of each free loop. Then it acts trivially on the curve graph of Σ, and so by Corollary 1.2 of [HHMV19], φ is isotopic to the identity. □

Corollary 2.4.7. *The homomorphism* $\mathrm{MCG}^*(\Sigma, p) \to \mathrm{Aut}(\mathcal{L}(\Sigma, p))$ *is bijective.*

Proof. Injectivity follows directly from Lemma 2.4.6.

To prove surjectivity, fix $f \in \mathrm{Aut}(\mathcal{L}(\Sigma, p))$, construct $\{\varphi_n\}_{n\in\mathbb{N}}$ as in Lemma 2.4.5. Define $\varphi : \Sigma \to \Sigma$ by letting $\varphi(x) = \varphi_n(x)$ for some n where $x \in \Sigma_n$. Since the Σ_n exhaust Σ and φ_n agrees with φ_m wherever both are defined, this map φ is well-defined. Since it is a homeomorphism on each Σ_n and the Σ_n exhaust Σ, it is a homeomorphism onto its image. And since the Ω_n also exhaust Σ, this image is in fact Σ, so $\varphi : \Sigma \to \Sigma$ is a homeomorphism. Thus $[\varphi] \in \mathrm{MCG}^*(\Sigma, p)$, and its image is f. □

2.5 The relative arc graph

The contents of Theorem 2.1.1 can be generalized to the case of the *relative arc graph* $\mathcal{A}(\Sigma, P)$. Recall the definition of the relative arc graph from Section 2.1, where we noted that that the loop graph $\mathcal{L}(\Sigma, p)$ is simply a special name for the relative

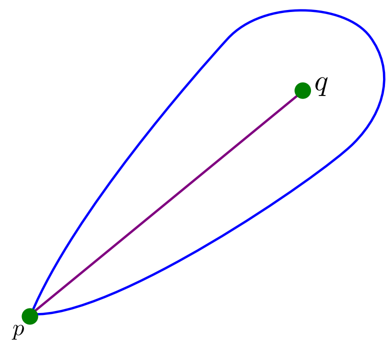

Figure 2.6: Adding a single arc (shown here in purple) turns a punctured monogon containing q into a degenerate triangle.

arc graph $\mathcal{A}(\Sigma \setminus \{p\}, \{p\})$; more generally, if $p \in P$ then $\mathcal{L}(\Sigma \cup \{p\}, p)$ is an induced subgraph of $\mathcal{A}(\Sigma, P)$. So for the remainder of this section we pick some $p \in P$ and $f \in \mathrm{Aut}(\mathcal{A}(\Sigma, P))$ and let $\Sigma' = \Sigma \cup \{p\}$. This gives an immediate result:

Lemma 2.5.1. *The homomorphism* $\mathrm{MCG}^*(\Sigma, P) \to \mathrm{Aut}(\mathcal{A}(\Sigma, P))$ *is injective.*

Proof. A homeomorphism $\varphi : \Sigma \to \Sigma$ that induces the identity automorphism on $\mathcal{A}(\Sigma, P)$ must also induce the identity automorphism on $\mathcal{L}(\Sigma', p)$ and so by Lemma 2.4.6 it is isotopic to the identity. □

We will now show surjectivity. In Section 2.3, we built a special triangualation $\mathcal{T}$. In this section, we will instead start with an arbitrary triangulation T' in $\mathcal{L}(\Sigma', p)$ as described in Section 2.2 and extend it to a triangulation T in $\mathcal{A}(\Sigma, P)$.

Lemma 2.5.2. *If T' is a triangulation in $\mathcal{L}(\Sigma', p)$, then there is a unique extension of T' to a triangulation T in $\mathcal{A}(\Sigma, P)$.*

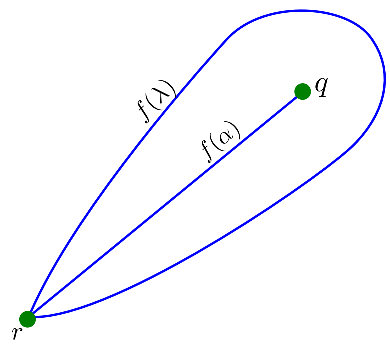

Figure 2.7: The image of a degenerate triangle under f

Proof. Consider a puncture $q \in P \setminus \{p\}$. Observe that T' must include a loop that bounds a punctured monogon around q—if not, such a loop could be added to T', contradicting maximality. The only way to extend T' with an arc having an endpoint at q is to turn this monogon into a degenerate triangle as in Figure 2.6. Repeating this process for each $q \in P \setminus \{p\}$ will give the only triangulation in $\mathcal{A}(\Sigma, P)$ containing T'. Call this triangulation T. □

Note that each arc of T is either contained in T'—in which case it is a loop based at p—or is one of these new arcs, each of which forms two sides of a single degenerate triangle in T'. The following lemma shows that this condition holds after applying the automorphism f.

Lemma 2.5.3. *Let T' be a triangulation in $\mathcal{L}(\Sigma', p)$ and T its extension to $\mathcal{A}(\Sigma, P)$ by Lemma 2.5.2. Then there is a puncture $\varphi(p) \in P$ such that for every loop $\lambda \in T'$ based at p, $f(\lambda)$ is a loop based at $\varphi(p)$, and for every arc $\alpha \in T$ with exactly one endpoint at p, $f(\alpha)$ has exactly one endpoint at $\varphi(p)$.*

Proof. Let $n = |P|$. By construction T has exactly $n - 1$ degenerate triangles, and Fact 2.4.1 ensures that $f(T)$ has the same number of degenerate triangles. Let us label one such degenerate triangle as in Figure 2.7. Observe that any arc in $\mathcal{A}(\Sigma, P)$ with an endpoint at q either is $f(\alpha)$ or intersects $f(\lambda)$. Thus the only arc in $f(T)$ with an endpoint at q is $f(\alpha)$.

By applying the above argument to each degenerate triangle of $f(T)$, we find $n-1$ punctures, each of which is the endpoint of exactly one arc, leaving one puncture left over. This puncture will be $\varphi(p)$, and note that in Figure 2.7 the puncture r is the endpoint of at least two arcs ($f(\lambda)$ and $f(\alpha)$) and so $r = \varphi(p)$.

If α is an arc in T with exactly one endpoint at p, then by construction it is the doubled edge in a degenerate triangle and so $f(\alpha)$ has exactly one endpoint at $\varphi(p)$. All other arcs in $f(T)$ must have both endpoints at $\varphi(p)$ because there are no other punctures in P available. It follows that all other arcs in $f(T)$—that is, $f(\lambda)$ for every $\lambda \in T'$—are loops based at $\varphi(p)$. □

We would like $\varphi(p)$ to be defined independently of our choice of a triangulation T' in the statement of Lemma 2.5.3, but this identification is not immediately obvious.

Lemma 2.5.4. *For any loop $\lambda \in \mathcal{L}(\Sigma', p)$, $f(\lambda)$ is a loop based at $\varphi(p)$.*

Proof. By extending λ to a triangulation in $\mathcal{L}(\Sigma', p)$ via Lemma 2.2.2, we see by Lemma 2.5.3 that $f(\lambda)$ is indeed a loop. If $\mu \in \mathcal{L}(\Sigma', p)$ is disjoint from λ, then the set $\{\lambda, \mu\}$ can also be extended to a triangulation in $\mathcal{L}(\Sigma', p)$, which means $f(\lambda)$ and $f(\mu)$ are based at the same puncture.

Even if λ and μ are not disjoint, the loop graph is connected (Theorem 1.1 of [AFP17]) and so there is a path $\lambda = \lambda_0, \ldots, \lambda_k = \mu$ so that each λ_i and λ_{i+1} are disjoint. Then $f(\lambda_i)$ and $f(\lambda_{i+1})$ are based at the same puncture, and so by induction are $f(\lambda)$ and $f(\mu)$.

Thus the image of every loop will be based at $\varphi(p)$, regardless of which triangulation was used to find $\varphi(p)$ in Lemma 2.5.3. $\square$

We now have very nearly all the ingredients necessary to apply the results of Section 2.4. If $\varphi(p) = p$, then Lemma 2.5.4 means $f \in \mathrm{Aut}(\mathcal{L}(\Sigma', p))$ and we can apply Corollary 2.4.7 directly. If not, we need only a bit more bookkeeping.

Lemma 2.5.5. *There exists a homeomorphism $\varphi \in \mathrm{MCG}^*(\Sigma, P)$ inducing the automorphism f on $\mathcal{A}(\Sigma, P)$.*

Proof. Let ψ be a homeomorphism of Σ that transposes p and $\varphi(p)$ while otherwise fixing the ends of Σ, and let g be the automorphism of $\mathcal{A}(\Sigma, P)$ induced by ψ. Then consider $f' = g \circ f$. By Lemma 2.5.4 and the construction of g, every loop $\lambda \in \mathcal{L}(\Sigma', p)$ is mapped to a loop $f'(\lambda)$ based at p. In other words, $f' \in \mathrm{Aut}(\mathcal{L}(\Sigma', p))$, and so by Corollary 2.4.7 it is induced by a homeomorphism $\varphi' \in \mathrm{MCG}^*(\Sigma', p) \subseteq \mathrm{MCG}^*(\Sigma, P)$. Let $\varphi = \psi^{-1} \circ \varphi'$.

By construction, φ agrees with f on every loop based at p; in particular they agree on the image of $\mathcal{T}'$ and thus of $\mathcal{T}$. Then as in the proof of Lemma 2.4.5, φ agrees with f on each finite-type subsurface Σ_n because it preserves a triangulation of that subsurface. Since every arc is contained in some Σ_n for high enough n, φ must agree with f on all arcs, and so φ induces f. $\square$

Chapter 3

Quasi-isometries

> This particularly rapid
> unintelligible patter isn't
> generally heard, and if it is it
> doesn't matter!
>
> W. S. Gilbert, *The Pirates of Penzance*

3.1 Introduction

The goal of this chapter is to prove Theorems 3.5.3 and 3.6.5, as outlined in Section 1.6. It is lightly adapted from [SC20].

This work was motivated by the resutls of Mann and Rafi [MR19], which provide a thorough application of the philosophy of coarsely bounded generating sets (see Section 1.5.1) to the area of big mapping class groups. In particular, their paper provides (under the technical condition of *tameness*—see Definition 3.2.8) a classification of which infinite-type surfaces admit coarsely bounded identity neighborhoods and generating sets. The natural question that then follows, as suggested in Section 1.6, is whether these big mapping class groups are quasi-isometric to some "nice" metric spaces.

A number of such spaces have been studied. Examples include Bavard's *loop graph* [Bav16], Rasmussen's *nonseparating curve graph* [Ras20], the graphs of separating curves defined by Durham, Fanoni, and Vlamis [DFV18], and the general curve and arc graphs defined by Aramayona, Fossas, and Parlier [AFP17].[1] The fact that all of these are graphs of curves or arcs motivates our focus on such graphs.

For graphs of curves we come to a very satisfying conclusion: we define a class of *translatable* surfaces—essentially, surfaces admitting a map that acts with north-south dynamics with respect to two distinct ends (see Definition 3.3.2)—and an associated *translatable curve graph*, and show that the mapping class group of a translatable surface is quasi-isometric to its translatable curve graph. What's more, we show that non-translatable surfaces do not admit such a graph of curves, except when that graph has finite diameter. More precisely, we prove:

Theorem 3.5.3. *Let Σ be an infinite-type surface with tame end space such that* $\mathrm{MCG}(\Sigma)$ *admits a coarsely bounded neighborhood of the identity and a coarsely*

[1]Fanoni, Ghaswala, and McLeay [FGM20] give an interesting action of a big mapping class group on the graph of *omnipresent arcs*, but this action is not continuous when the graph is given the topology of a simplicial complex.

bounded generating set—and thus has a well-defined quasi-isometry type—but is not itself coarsely bounded. Then the following are equivalent:

1. *There exists a graph* Γ *whose vertices are curves and such that the action of* $\mathrm{MCG}(\Sigma)$ *on* Γ *induces a quasi-isometry.*

2. Σ *is translatable.*

3. Σ *has no nondisplaceable finite-type surfaces, making it an* avenue surface *in the sense of Horbez, Qing, and Rafi [HQR20].*

We do not attempt in this paper to study the geometry of the translatable curve graph, although we hope this will be a fruitful avenue for further research. One property is however immediate: by the results of Horbez, Qing, and Rafi [HQR20] the translatable curve graph—and thus the mapping class group of a translatable surface—cannot be non-elementary δ-hyperbolic.

In the case of graphs of arcs we have not found such a general classification, but we exhibit one particularly striking quasi-isometry. Note that the surface in question is not translatable.

Theorem 3.6.5. *Let* $\Sigma = \mathbb{R}^2 \setminus C$ *be the plane minus a Cantor set. Then* $\mathrm{MCG}(\Sigma)$ *is quasi-isometric to the loop graph* $\mathcal{L}(\Sigma)$.

Though less general than the previous result, this quasi-isometry is of interest because the loop graph is already well-studied; for instance, the hyperbolicity of the loop graph was demonstrated by Bavard [Bav16], and thus $\mathrm{MCG}(\Sigma)$ is also hyperbolic; see Corollary 3.6.6. The Gromov boundary of this graph was also described by Bavard and Walker [BW18].

This is, to our knowledge, the first case of a big mapping class group being shown to be non-elementary δ-hyperbolic. By extension, it is also the first case in which two big mapping class groups have been shown to have distinct, non-trivial quasi-isometry types; see Corollary 3.6.7. Finally, it follows from Corollary 3.6.6 and work of Cornulier and de la Harpe that $\mathrm{MCG}(\Sigma)$ has a coarsely bounded presentation; see Corollary 3.6.10.

Before getting to the meat of the paper, we introduce an important motivating example and preview some of the techniques that will be used.

Durham, Fanoni, and Vlamis [DFV18], studying the Jacob's ladder surface (which has two ends, both accumulated by genus—see Figure 3.1) present the following subgraph of the curve graph of that surface: its vertices are curves separating the two ends, with an edge between two such curves if they cobound a genus-one subsurface. The main immediate application of this graph results from the fact that, unlike the full curve graph of an infinite-type surface, it has infinite diameter. In particular, a translation acts on this graph with unbounded orbits, which provides an easy proof that the mapping class group of this surface is not coarsely bounded.

An early version of Vlamis's notes on the topology of big mapping class groups [Vla19] claimed that this graph is quasi-isometric to the mapping class group of the Jacob's ladder surface. Vlamis's proof was incomplete—it showed only that the

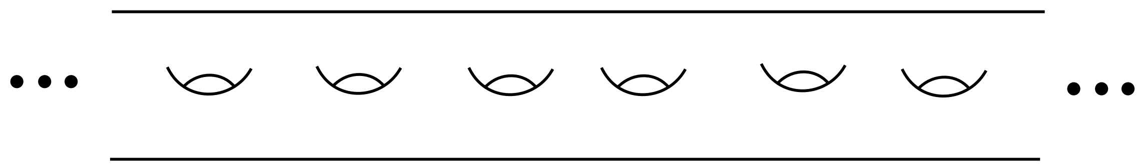

Figure 3.1: The Jacob's ladder surface.

vertex stabilizers are coarsely bounded, which is not sufficient to conclude quasi-isometry—but it provided significant inspiration for the results which eventually became Theorem 3.4.9.

First, this graph could in fact be shown to be quasi-isometric to the mapping class group, although it would take some additional effort. Second, the class of surfaces for which such a graph might be built could be expanded significantly beyond the Jacob's ladder surface. The properties of Durham, Fanoni, and Vlamis's graph depended largely on the translatable nature of the Jacob's ladder surface, rather than the details of the translation itself. Other surfaces admitting a similar kind of translation include the bi-infinite flute (see Figure 3.2) and more complicated surfaces that might be built by joining many copies of a single surface as in Figure 3.3 (on page 53). By modifying the construction of Durham, Fanoni, and Vlamis [DFV18] we are able to produce a general *translatable curve graph* $\mathcal{TC}(\Sigma)$ which is quasi-isometric to the mapping class group $\mathrm{MCG}(\Sigma)$; this is Theorem 3.4.9.

One obvious follow-up question, given this quasi-isometry, is whether other such graphs can be produced. For instance, are there other cases where a big mapping class group is quasi-isometric to a graph whose vertices are curves? What if the

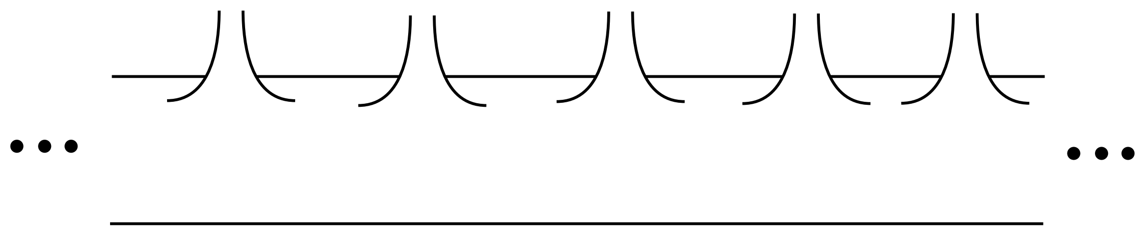

Figure 3.2: The bi-infinite flute.

vertices are arcs? Theorem 3.5.3 answers the first question; Theorem 3.6.5 is a partial answer to the second.

The structure of the rest of this chapter is as follows. In Section 3.2, we recall relevant results from previous papers ([Ric63], [Ros18], and [MR19]) that are used here, with the goal of making the chapter accessible to anyone with some knowledge of geometric group theory and low-dimensional topology, but who may not have worked previously with infinite-type surfaces or with the concept of coarse boundedness. Section 3.2 also presents and proves Lemma 3.2.2, which is a limited version of the Schwarz-Milnor lemma for the case of groups with coarsely bounded neighborhoods of the identity acting transitively on graphs.

In Section 3.3, we define translatable surfaces and prove some of their properties, most notably Proposition 3.3.5, which shows that every translatable surface can be written as an infinite connected sum of copies of some subsurface S as in Figure 3.3.

In Section 3.4 we define the translatable curve graph itself and prove the quasi-isometry to the mapping class group in Theorem 3.4.9. The main tools are a study of the maximal ends of the subsurface S found in Proposition 3.3.5, and Lemma

3.3.8, which allows us to embed the set of all mapping classes that fix half of our translatable surface in a conjugate of any neighborhood of the identity.

In Section 3.5 we show that translatable surfaces are in fact the only surfaces with non-coarsely-bounded mapping class groups quasi-isometric to a graph of curves, proving Theorem 3.5.3. The main tools here are, on one hand, a demonstration that under some reasonable conditions any surface with two equivalent maximal ends and zero or infinite genus is translatable; and on the other hand, that all other surfaces have mapping class groups that are either themselves coarsely bounded or have no coarsely bounded curve stabilizers, making such a quasi-isometry impossible.

Finally, Section 3.6 uses methods parallel to those in Sections 3.3 and 3.4 to prove that the mapping class group of the plane minus a Cantor set is quasi-isometric to the loop graph of that surface.

3.2 Preliminaries

Before we begin, we recall several results from the work of Rosendal [Ros18] and Mann and Rafi [MR19].

3.2.1 Coarse boundedness

Recall the definition of coarse boundedness from Section 1.5.1. We call a group *locally coarsely bounded* if it has a coarsely bounded neighborhood of the identity.

Recall also from Section 1.5.1 that if A and B are two coarsely bounded generating sets for a locally coarsely bounded group G, the word metrics with respect to the generating sets A and B are quasi-isometric.

We make heavy use of the following alternate characterization of coarse boundedness:

Lemma 3.2.1 (From [Ros18])**.** *Given G a Polish group, and a subset $A \subseteq G$. Then A is coarsely bounded if and only if for every identity neighborhood $V \subseteq G$ there is some $k \in \mathbb{N}$ and a finite set $F \subseteq G$ such that $A \subseteq (FV)^k$.*

In light of this definition, we will want to talk about specific identity neighborhoods in the mapping class group of a surface Σ: if S is a subsurface of Σ, let V_S be the set of mapping classes with representatives that restrict to the identity on S. Note that the set $\{V_S \mid S \subseteq \Sigma \text{ of finite type}\}$ forms a neighborhood basis of the identity in MCG(Σ).

We're looking to prove quasi-isometries between groups and graphs, so we want something reminiscent of the Schwarz-Milnor lemma.

Lemma 3.2.2. *Let G be a locally coarsely bounded group acting transitively by isometries on a connected graph Γ equipped with the edge metric. Suppose that for some vertex $v_0 \in \Gamma$, the set $A = \{g \in G \mid d(v_0, gv_0) \leq 1\}$ is coarsely bounded. Then the orbit map $g \mapsto gv_0$ is a quasi-isometry.*

Proof. Coarse surjectivity follows directly from the transitivity of the action.

Fix some $g \in G$. Since Γ is connected, there is a minimal-length path $v_0, v_1, \ldots, v_n = gv_0$ from v_0 to gv_0 with $d(v_i, v_{i+1}) = 1$. Since $d(v_0, v_1) = 1$ and the action of G

is transitive, there is some $g_0 \in A$ such that $g_0 v_0 = v_1$. Likewise, there is some $g_1' \in g_0 A g_0^{-1}$ such that $g_1' v_1 = v_2$. Writing $g_1' = g_0 g_1 g_0^{-1}$ with $g_1 \in A$ and $v_1 = g_0 v_0$, we see that $g_0 g_1 v_0 = v_2$. Continuing in this way, we can find $g_0, g_1, \ldots, g_{n-1} \in A$ such that $g_0 g_1 \cdots g_{n-1} v_0 = v_n$. Let $g_n = g^{-1} g_0 g_1 \cdots g_{n-1}$. Then $g_n v_0 = v_0$, so $g_n \in A$, and $g_0 g_1 \cdots g_{n-1} g_n = g$. Thus A is a generating set for G and the word-metric length of g in A is at most $n = d(v_0, gv_0) + 1$.

On the other hand, suppose $g_0 g_1 \cdots g_k = g$ with each $g_i \in A$ and k minimal. By the definition of A, $d(g_0 g_1 \cdots g_i v_0, g_0 g_1 \cdots g_{i+1} v_0) = d(v_0, g_{i+1} v_0) \leq 1$, so the distance $d(v_0, gv_0) \leq k + 1$. Thus the map $g \mapsto gv_0$ coarsely preserves the word metric with the generating set A, and thus the orbit map is a quasi-isometry for any choice of coarsely bounded generating set for G. □

Before we can apply this lemma, we need to know a bit more about infinite-type surfaces.

3.2.2 Infinite-type surfaces

Recall from Section 1.1 that an infinite-type surface is classified by its (possibly infinite) genus, its spacee of ends E, and the subset E_G of ends accumulated by genus. Following the example of Mann and Rafi, we will mostly avoid referencing E_G explicitly, and implicitly assume it is preserved. For instance, when we say that two subsets $U, V \subseteq E$ are homeomorphic, we mean that there is a homeomorphism $f\colon U \to V$ such that $f(U \cap E_G) = V \cap E_G$.

Mann and Rafi introduce the following partial pre-order on E, which provides a valuable for tool for studying its topology:

Definition 3.2.3. Given $x, y \in E$, we say $x \succcurlyeq y$ if for every clopen neighborhood U of x, there exists a clopen neighborhood V of y homeomorphic to a clopen subset of U.

As might be expected, we write $y \prec x$ when $y \preccurlyeq x$ but $x \not\preccurlyeq y$, and $y \sim x$ when $y \preccurlyeq x$ and $x \preccurlyeq y$. We use the notation $E(x) = \{y \in E \mid x \sim y\}$ for the equivalence class of $x \in E$ under the relation $\sim$.

Crucially, this order has maximal elements, which have a fairly rigid structure:

Fact 3.2.4 (Proposition 4.7 of [MR19])**.** *The partial pre-order $\preccurlyeq$ has maximal elements. Furthermore, for every maximal element $x \in E$, the equivalence class $E(x)$ is either finite or a Cantor set.*

Mann and Rafi also define the following self-similarity condition, and prove some useful properties of it:

Definition 3.2.5. A clopen neighborhood U of an end $x \in E$ is *stable* if for every clopen neighborhood U' of x contained in U, there is a clopen subset of U' homeomorphic to U.

Fact 3.2.6 (Lemma 4.17 of [MR19])**.** *If $x \sim y \in E$ and x has a stable neighborhood U, then all sufficiently small neighborhoods of y are homeomorphic to U via a homeomorphism taking x to y.*

Fact 3.2.7 (Lemma 4.18 of [MR19])**.** *Let $x, y \in E$ and assume x has a stable neighborhood V_x, and that x is an accumulation point of $E(y)$. Then for any sufficiently small clopen neighborhood U of y, $U \cup V_x$ is homeomorphic to V_x.*

Many of the results in later sections will assume the existence of certain stable neighborhoods, in the form of what Mann and Rafi call *tameness*:

Definition 3.2.8. An end space E is said to be *tame* if any $x \in E$ that is either maximal or an immediate predecessor to a maximal end has a stable neighborhood.

It is an open question (Problem 6.15 of [MR19]) whether there exists any surface with non-tame end space whose mapping class group is not coarsely bounded but has a well-defined quasi-isometry type. For this reason, we consider tameness to be an acceptable condition to impose in some of our results.

To achieve the negative results of subsection 3.5.3, we will need to consider the properties of mapping class groups that are locally coarsely bounded and admit a coarsely bounded generating set. Here $\mathcal{M}(X)$ denotes the set of maximal ends of some subspace $X \subseteq E$.

Fact 3.2.9 (Theorem 1.4 of [MR19])**.** $\mathrm{MCG}(\Sigma)$ *is locally coarsely bounded if and only if either* $\mathrm{MCG}(\Sigma)$ *is itself coarsely bounded or there is a finite-type surface $K \subseteq \Sigma$ such that the complementary regions of K each have infinite type and zero or infinite genus, and partition E into finitely many clopen sets*

$$E = \left(\bigsqcup_{A \in \mathcal{A}} A \right) \sqcup \left(\bigsqcup_{P \in \mathcal{P}} P \right)$$

such that:

1. *Each $A \in \mathcal{A}$ is self-similar, with $\mathcal{M}(A) \subseteq \mathcal{M}(E)$ and $\mathcal{M}(E) \subseteq \bigsqcup_{A \in \mathcal{A}} \mathcal{M}(A)$,*

2. *each $P \in \mathcal{P}$ homeomorphic to a clopen subset of some $A \in \mathcal{A}$, and*

3. *for any $x_A \in \mathcal{M}(A)$, and any neighborhood V of the end x_A in Σ, there is $f_V \in \mathrm{MCG}(\Sigma)$ so that $f_V(V)$ contains the complementary component to K with end space A.*

Moreover, in this case V_K—the set of mapping classes restricting to the identity on K—is a coarsely bounded neighborhood of the identity.

We will also make use of the following necessary condition for $\mathrm{MCG}(\Sigma)$ to have a coarsely bounded generating set:

Definition 3.2.10 (Definition 6.2 of [MR19])**.** We say that an end space E has *limit type* if there is a finite-index subgroup G of $\mathrm{MCG}(\Sigma)$, a G-invariant set $X \subseteq E$, points $z_n \in E$ indexed by $n \in \mathbb{N}$ which are pairwise inequivalent, and a nested family of clopen sets U_n with $\bigcap_{n \in \mathbb{N}} U_n = X$ such that

$$E(z_n) \cap U_n \neq \emptyset, \quad E(z_n) \cap U_0^c \neq \emptyset, \quad \text{and} \quad E(z_n) \subseteq (U_n \cup U_0^c)$$

where $U_0^c = E \setminus U_0$.

This definition is somewhat daunting, so we present the following example of a surface whose end space has limit type. Let z_1 be a puncture, and for each $n > 1$ let z_n be an end accumulated by countably many points locally homeomorphic to

z_{n-1}. Then let z_ω be an end accumulated by $\{z_n\}_{n\in\mathbb{N}}$. Let F be the set of ends just defined, and let Σ be a surface with zero genus and end space homeomorphic to the disjoint union of n copies of F for any $n > 1$. It can be verified that the end space of Σ has limit type.

Fact 3.2.11 (Lemma 6.4 of [MR19])**.** *If E has limit type, then* MCG(Σ) *does not admit a coarsely bounded generating set.*

3.3 Translations on surfaces

We first define a useful notion of convergence:

Definition 3.3.1. Given a surface Σ, an end e of Σ, and a sequence $\alpha_1, \alpha_2, \ldots$ of curves on Σ, we say that $\lim_{n\to\infty} \alpha_n = e$ if, for every neighborhood V of the end e in the surface Σ, all but finitely many of the α_i are (after some isotopy) contained in V.

A translation, then, will be a map that moves all curves toward one end and away from another. That is:

Definition 3.3.2. Given a surface Σ, a map $h \in \mathrm{MCG}(\Sigma)$ is called a *translation* if there are two distinct ends e_+ and e_- of Σ such that for any curve α on Σ, $\lim_{n\to\infty} h^n(\alpha) = e_+$ and $\lim_{n\to\infty} h^{-n}(\alpha) = e_-$. If such a translation exists, we call the surface Σ *translatable.*

Remark 3.3.3. This definition brings to mind several other classes of infinite-type surfaces with two special ends that have recently been defined for various reasons.

- The *telescoping* surfaces of Mann and Rafi [MR19] form a strict subset of the translatable surfaces: though every telescoping surface can be shown to be translatable, the Jacob's ladder surface is translatable but not telescoping.
- The *doubly pointed* surfaces of Aougab, Patel, and Vlamis [APV20] are a strict superset of the translatable surfaces: every translatable surface is doubly pointed, but the surface with zero genus and end space homeomorphic to $\omega^\omega 2 + 1$ is doubly pointed but not translatable.[2]
- The *avenue* surfaces of Horbez, Qing, and Rafi [HQR20] turn out to be precisely those translatable surfaces that have tame end space. This result is part of Theorem 3.5.3.

It follows directly from the definition that a translatable surface Σ cannot contain any finite-type nondisplaceable surfaces, and so in particular Σ cannot have finite type, finite positive genus, or any ends with a finite $MCG(\Sigma)$-orbit of size more than 2. This gives us lots of non-examples, but there are also plenty of translatable surfaces if we go looking for them.

[2]There are easier counterexamples, e.g. a surface with two inequivalent maximal ends (see Lemma 3.5.16). This example demonstrates however that a doubly pointed surface may not be translatable even if it has exactly two equivalent maximal ends.

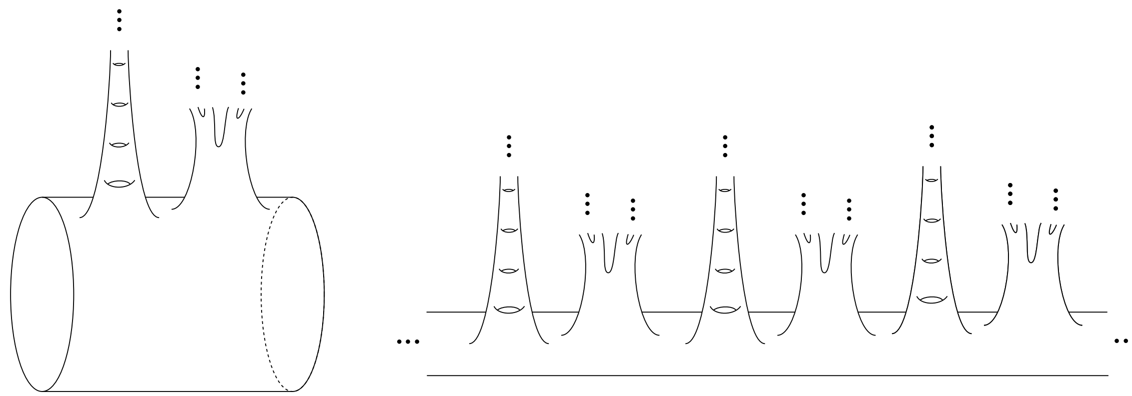

Figure 3.3: On the left, a surface S with two boundary components (in this case, the connected sum of an annulus, a Loch Ness Monster, and a Cantor tree). On the right, the translatable surface $S^{\natural\mathbb{Z}}$.

The Jacob's ladder surface and the bi-infinite flute, in Figures 3.1 and 3.2, are clearly translatable; we can think of the flute as being derived from the ladder by replacing each handle with a puncture. More generally, we might replace each handle in the ladder with some other surface, as follows. Let S be any surface, not necessarily of finite type, with two compact boundary components, and let $\Sigma = S^{\natural\mathbb{Z}}$ be the gluing along their boundaries[3] of countably many copies of S, arranged like $\mathbb{Z}$ as in Figure 3.3. Then the map that takes each copy of S to the next one over is a translation, and so Σ is translatable.

A natural question to ask is whether this last example includes all translatable surfaces, and in fact it does. However, we may have to choose a different translation. For this and future results, the following notation will be useful:

[3]The use of $\natural$ here is intended to invoke the standard use of $\#$ for connected sum, and was suggested to me on Facebook by Rylee Lyman.

Definition 3.3.4. Suppose Σ is a translatable surface and α a curve in Σ separating e_+ and e_-. Then we denote by α_+ (resp. α_-) the component of $\Sigma \setminus \alpha$ containing the end e_+ (resp. e_-). If β is a curve separating α and e_+, then we denote by (α, β) the subsurface $\alpha_+ \cap \beta_-$ of Σ bounded by α and β.

Proposition 3.3.5. *Let Σ be a translatable surface with translation h and α a curve separating the ends e_+ and e_-. Then there is a surface $S = (\alpha, h^N(\alpha))$ for some N such that Σ is homeomorphic to $S^{\natural\mathbb{Z}}$.*

Proof. By the definition of translation, there is some $N \in \mathbb{N}$ such that for all $n \geq N$, $h^n(\alpha) \subseteq \alpha_+$. Then h^N is also a translation, so without loss of generality we can replace h with h^N and assume that all $h^i(\alpha)$ are disjoint.

Let $S = (\alpha, h(\alpha))$. If $x \in \Sigma$ but not in any $h^i(\alpha)$, then x is either in α_+ or α_-. Supposing without loss of generality that $x \in \alpha_+$, there must be some least i such that $x \notin h^i(\alpha)_+$, otherwise $h^i(\alpha)$ could not converge to e_+. Then $x \in (h^{i-1}(\alpha), h^i(\alpha)) = h^{i-1}(S)$. On the other hand, $h^i(S) \cap S = \emptyset$ for all $i \neq 0$ by construction, and so every point of Σ is in exactly one $h^i(S)$ or $h^i(\alpha)$.

The resulting subsurface S has two boundary components, and the copies of S are glued together exactly as desired. $\square$

This decomposition depended on the choice of a curve separating e_+ and e_-. We might ask how important that choice was, and it turns out the answer is "not much".

Lemma 3.3.6. *Let Σ be a translatable surface, and α and β two curves separating the ends e_+ and e_-. Then there is some $f \in \mathrm{MCG}(\Sigma)$ which fixes the ends e_+ and e_- and such that $f(\alpha) = \beta$.*

Proof. First, replace β with some $h^n(\beta)$ so that $\beta \subseteq \alpha_+$, and then replace h with some power of h so that $\beta \subseteq (\alpha, h(\alpha))$. By Proposition 3.3.5, we can write $\Sigma = S^{\natural\mathbb{Z}} = T^{\natural\mathbb{Z}}$, where $S = (\alpha, h(\alpha))$ and $T = (\beta, h(\beta))$.

If we let $X = (\alpha, \beta)$, $Y = (\beta, h(\alpha))$, and $Z = (h(\alpha), h(\beta))$, then $(\alpha, h(\alpha)) = X \natural Y$ and $(\beta, h(\beta)) = Y \natural Z$. But X and Z are homeomorphic, and thus so are $S = (\alpha, h(\alpha))$ and $T = (\beta, h(\beta))$. It follows that we can map each copy of S to the appropriate copy of T, giving us a homeomorphism of Σ that takes α to β and fixes the ends e_+ and e_-. □

There is one more symmetry of a translatable surface worth discussing here:

Lemma 3.3.7. *Let Σ be a translatable surface, and α a curve separating the ends e_+ and e_-. Then there is some $r \in \mathrm{MCG}(\Sigma)$ that transposes e_+ and e_- and restricts to an orientation-reversing homeomorphism on α.*

Proof. In a tubular neighorhood of α, r is just a rotation by π about a diameter of the circle α—see Figure 3.4. By Proposition 3.3.5, α_+ and α_- are homeomorphic, so this r can be extended to a homeomorphism on all of Σ. □

The following lemma will be quite useful in light of Lemma 3.2.1.

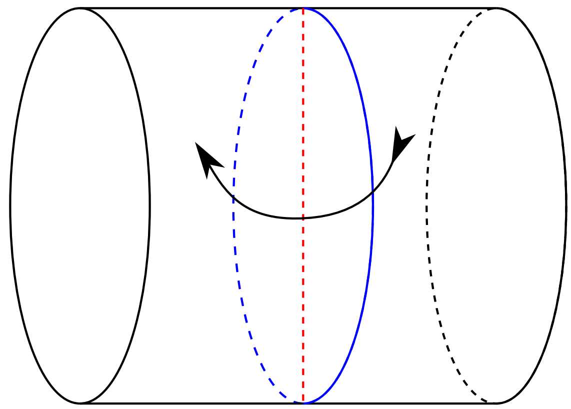

Figure 3.4: Within a tubular neighborhood of the blue curve α, the curve's orientation can be reversed via a rotation by π about its red diameter.

Lemma 3.3.8. *Let Σ be a translatable surface with translation h, and α a curve separating the ends e_+ and e_-. Then for any identity neighborhood V in* $\mathrm{MCG}(\Sigma)$*, there is some $n \in \mathbb{N}$ such that $V_{\alpha_-} \subseteq h^{-n}Vh^n$ and $V_{\alpha_+} \subseteq h^nVh^{-n}$.*

Proof. By the topology of $\mathrm{MCG}(\Sigma)$, there is some finite-type subsurface $T \subseteq \Sigma$ such that $V_T \subseteq V$. Let $n \in \mathbb{N}$ so that $T \subseteq h^n(\alpha)_-$. Then we have

$$\begin{aligned} T &\subseteq h^n(\alpha)_- \\ V_{h^n(\alpha)_-} &\subseteq V_T \\ V_{h^n(\alpha)_-} &\subseteq V \\ h^nV_{\alpha_-}h^{-n} &\subseteq V \\ V_{\alpha_-} &\subseteq h^{-n}Vh^n \end{aligned}$$

and likewise $V_{\alpha_+} \subseteq h^nVh^{-n}$. □

Corollary 3.3.9. *Let Σ be a translatable surface with translation h, and α a curve separating the ends e_+ and e_-. Then the set $H = \{f \in \mathrm{MCG}(\Sigma) \mid f(\alpha)$ is homotopic to $\alpha\}$ of mapping classes stabilizing α is coarsely bounded.*

Proof. Fix an identity neighborhood V, and using Lemma 3.3.8 find $n \in \mathbb{N}$ so that $V_{\alpha_-} \subseteq h^{-n}Vh^n$ and $V_{\alpha_+} \subseteq h^nVh^{-n}$. Let $F = \{r^{-1}, h^n, h^{-n}\}$ where r is the map defined in Lemma 3.3.7. We claim that $H \subseteq (FV)^5$, which gives the result by Lemma 3.2.1.

Pick $f \in H$. Up to homotopy, $f_{|\alpha}$ is either the identity or a reflection map; in the latter case replace f with rf so that $f_{|\alpha}$ is the identity. Then the action of f on α_+ and α_- do not interact, and so f can be decomposed as $f = f_-f_+$, where $f_- \in V_{\alpha_-}$ and $f_+ \in V_{\alpha_+}$. It follows by our choice of n that $f_- \in h^{-n}Vh^n$ and $f_+ \in h^nVh^{-n}$, so $f \in h^{-n}Vh^nh^nVh^{-n} \subseteq (FV)^4$. Since we may also have applied r^{-1}, this gives $f \in (FV)^5$. □

Corollary 3.3.10. *Let Σ be a translatable surface. Then* $\mathrm{MCG}(\Sigma)$ *is locally coarsely bounded.*

Proof. The stabilizer H of a curve separating the ends e_+ and e_- is an identity neighborhood, and by Corollary 3.3.9 it is coarsely bounded. □

Note that unlike in the following sections, here we have not assumed tameness.

3.4 The translatable curve graph

We are now ready to define a graph quasi-isometric to $\mathrm{MCG}(\Sigma)$ when Σ is a translatable surface.

Definition 3.4.1. Fix a collection $\mathcal{S}$ of subsurfaces of Σ. The *translatable curve graph* $\mathcal{TC}(\Sigma, \mathcal{S})$ of Σ with respect to the set of subsurfaces $\mathcal{S}$ is the graph whose vertices are curves separating e_+ and e_-, with an edge between two curves α and β if they have disjoint representatives and (α, β) or (β, α) is homeomorphic to some $S \in \mathcal{S}$.

We will eventually define a canonical and finite set $\mathcal{S}$ depending only on the surface Σ; once this has been defined, we will omit $\mathcal{S}$ and write simply $\mathcal{TC}(\Sigma)$.

Note that Corollary 3.3.9 implies that vertex stabilizers of $\mathcal{TC}(\Sigma, \mathcal{S})$ are coarsely bounded. Also, with Σ the Jacob's ladder surface and S a surface with genus 1 and two boundary components, the graph $\mathcal{TC}(\Sigma, \{S\})$ is precisely the motivating example from page 42.

For an arbitrary translatable surface Σ, we might try taking a subsurface S such that $\Sigma = S^{\natural\mathbb{Z}}$ as in Proposition 3.3.5. But $\mathcal{TC}(\Sigma, \{S\})$ will not in general be connected. Instead, we will use the topology of S to construct a collection of subsurfaces $\mathcal{S}$ such that $\mathcal{TC}(\Sigma, \mathcal{S})$ satisfies the conditions of Lemma 3.2.2.

Consider the space of ends $E(S)$ of S; we may ask how these relate to the order structure on $E(\Sigma)$.

Lemma 3.4.2. *The maximal ends of the subsurface S are either maximal ends of Σ or immediate predecessors to maximal ends of Σ.*

Proof. It follows from the decomposition given in Proposition 3.3.5 that e_+ and e_- are maximal ends of Σ, and that $e_+ \sim e_-$. In fact, they are global maxima of the partial preorder: for any end x of Σ, $x \preccurlyeq e_+$. If $E(e_+)$ contains some point y distinct from e_+ and e_-, then y is in the end space of some copy of S. In particular, since it is still true that for every end x of Σ, $x \preccurlyeq e_+ \sim y$, $E(e_+)$ contains *all* the maximal ends of S.[4]

If, on the other hand, $E(e_+) = \{e_+, e_-\}$, then let x be maximal in S. We know that $x \prec e_+$. Suppose we have an end y such that $x \preccurlyeq y \preccurlyeq e_+$. If y is an end of some copy of S, then by maximality $y \sim x$. But if y is not an end of any copy of S, the only other possibility is that $y = e_\pm$, in which case $y \sim e_+$. Thus x is an immediate predecessor to e_+. □

Corollary 3.4.3. *If the end space of Σ is tame, then every maximal end of S has a stable neighborhood.*

In general, a surface might have infinitely many equivalence classes of maximal ends. With tameness, however, the possibilities are much more limited.

[4]This implies, by Proposition 4.8 of [MR19], that the end space of Σ is self-similar and thus $\mathrm{MCG}(\Sigma)$ is coarsely bounded, and we will in fact see that $\mathcal{TC}(\Sigma, \mathcal{S})$ has finite diameter under these circumstances.

Lemma 3.4.4. *If every maximal end of $T \subseteq \Sigma$ has a stable neighborhood, then the end space $E(T)$ has finitely many equivalence classes of maximal ends.*

Proof. For each maximal end x of T, let V_x be a stable neighborhood of x. For every end y, pick a clopen neighborhood U_y of y such that U_y is homeomorphic to a clopen subset of V_x for some maximal end x.

Since the neighborhoods U_y cover the end space $E(T)$, and $E(T)$ is compact, there is a finite set $U_1, \ldots, U_n$ covering $E(T)$, where each U_i is homeomorphic to a clopen subset of V_{x_i}.

Now for each $y \in E(T)$, $y \in U_i$ for some i. Let V be an arbitrary clopen neighborhood of x_i; by stability, V contains a homeomorphic copy of V_{x_i}, which in turn contains a homeomorphic copy of U_i by construction. Thus $y \preccurlyeq x_i$. It follows that the set $\{x_1, \ldots, x_n\}$ contains a representative of every equivalence class of maximal ends, so there are at most n such equivalence classes. □

While there are finitely many equivalence classes of maximal ends, a priori the non-maximal ends might contribute meaningfully to the topology of a subsurface. However, this is not the case.

Lemma 3.4.5. *Given a subsurface $T \subseteq \Sigma$ with end space $E(T)$ and such that every maximal end of T has a stable neighborhood, $E(T)$ can be written as the disjoint union $\bigsqcup_{i=1}^{k} V_k$ where each V_i is a stable neighborhood of a maximal end of T.*

Proof. Let M be a set containing every maximal end x such that $E(x) \cap E(T)$ is finite, and a single representative from $E(x)$ for every equivalence class of maximal ends such that $E(x) \cap E(T)$ is infinite. By Lemma 3.4.4 M is finite and thus discrete, so we can pick disjoint stable neighborhoods $V_x \subseteq E(T)$ for each $x \in M$. Let $V = \bigsqcup_{x \in M} V_x$.

For each $y \in E(T) \setminus V$, there is by maximality some $x \in M$ such that $y \preccurlyeq x$. If $y \prec x$, then x is an accumulation point of $E(y)$; and if $y \sim x$, then since $y \notin V$ the set $E(x) \cap E(T)$ must be infinite, and thus a Cantor set by Fact 3.2.4, and so again x is an accumulation point of $E(y) = E(x)$. Thus in either case we can apply Fact 3.2.7 to find some clopen $U_y \ni y$ such that $U_y \cup V_x$ is homeomorphic to V_x. The set $E(T) \setminus V$ is clopen and thus compact, and is covered by the neighborhoods U_y, so there is a finite set of these neighborhoods covering $E(T) \setminus V$; since they are all clopen we can ensure they are disjoint.

For each U_y in this finite set, pick V_x so that $V_x \sqcup U_y$ is homeomorphic to V_x, and then replace V_x with $V_x \sqcup U_y$. After finitely many steps, the entire end space $E(T)$ is contained in the disjoint union $\bigsqcup_{x \in M} V_x$. See Figure 3.5 for an example. □

We are now ready to define the canonical collection $\mathcal{S}$ of subsurfaces that will be used in the definition of the graph $\mathcal{TC}(\Sigma, \mathcal{S})$. The following construction assumes Σ is tame. Let $\{f_1, \ldots, f_n, c_1, \ldots, c_m\}$ be representatives of the equivalence classes of maximal ends of S, with each $E(f_i)$ intersecting the end space of S finitely many times, and each $E(c_i)$ intersecting the end space of S in a Cantor set. Pick disjoint stable neighborhoods V_{f_i} and V_{c_i} of each representative.

For each $1 \leq i \leq n$, let T_i be a surface with two boundary components and

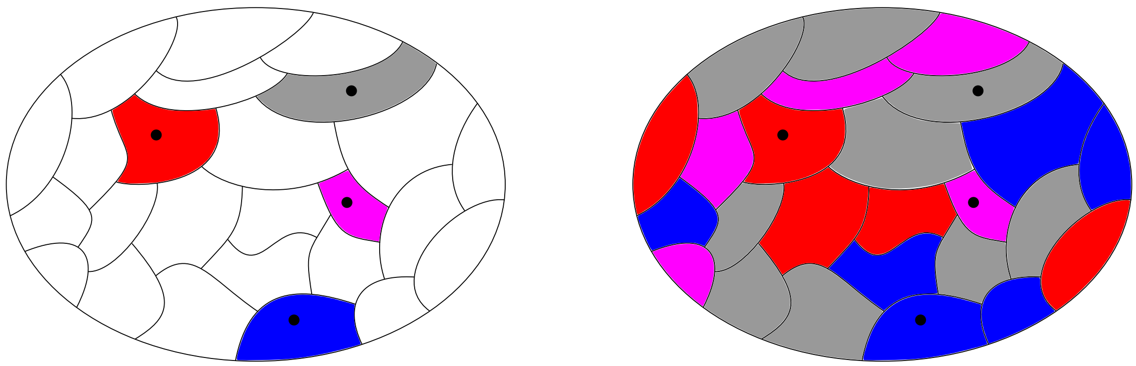

Figure 3.5: From the proof of Lemma 3.4.5: on the left, the end space $E(T)$ is divided into finitely many disjoint regions by compactness. Each shaded region is a stable neighborhood of a maximal point, shown as a black dot. Each unshaded region satisfies Fact 3.2.7 with respect to one of the shaded regions. On the right, the shaded regions have been expanded by repeated application of Fact 3.2.7 so that they now cover the whole surface. Note that each shaded region is still a stable neighborhood of the maximal point shown as a dot; in fact, each large shaded region on the right is homeomorphic to the similarly-colored small region on the left.

end space homeomorphic to $V_{f_i} \sqcup \bigsqcup_{j=1}^{m} V_{c_j}$. If f_i or any of the c_j is accumulated by genus, then by construction T_i will have infinite genus; if none of them are, we further specify that T_i have genus zero. If none of the f_i or c_i are accumulated by genus, but the surface Σ has positive genus—in other words, if S has finite positive genus—then let T_{n+1} be the surface with two boundary components, genus 1, and end space homeomorphic to $\bigsqcup_{j=1}^{m} V_{c_j}$. Let $\mathcal{S} = \{T_1, \dots, T_n, (T_{n+1})\}$, including T_{n+1} if it has been defined.

We need to handle an edge case: if $n = 0$ and S has 0 or infinite genus, the above construction will give $\mathcal{S} = \emptyset$, which is not desirable; so in this case we let $\mathcal{S} = \{S\}$, noting that S is in fact a surface with two boundary components and end space homeomorphic to $\bigsqcup_{j=1}^{m} V_{c_j}$ by Lemma 3.4.5. It can be seen without too much trouble that in this case $\mathcal{TC}(\Sigma, \mathcal{S})$ has diameter 2; this is consistent with the fact that under these conditions the end space of Σ is either self-similar (if the maximal ends are all equivalent to e_+ and e_-) or *telescoping* with respect to e_+ and e_- (if they are predecessors), and so by Proposition 3.5 of [MR19] the mapping class group $\mathrm{MCG}(\Sigma)$ is coarsely bounded.

From here on we assume $\mathcal{S}$ is the set of subsurfaces just constructed, and write $\mathcal{TC}(\Sigma)$ to mean $\mathcal{TC}(\Sigma, \mathcal{S})$.

Before proving connectedness we introduce the following construction, which allows us to produce subsurfaces of Σ with nearly arbitrary genus and end space.

Lemma 3.4.6. *Let Σ be a translatable surface with a curve α separating e_+ and e_-.*

Then for any clopen subset $V \subseteq E(\alpha_+)$*, there is a curve* β *also separating* e_+ *and* e_- *and such that* $E((\alpha, \beta)) = V$*. Furthermore, if no end of* V *is accumulated by genus, but* e_+ *is, there is for every* $n \in \mathbb{N}$ *a choice of* β *such that* (α, β) *has genus* n*.*

Proof. Recall that for any separating curve γ on Σ, the end sets of the two components of $\Sigma \setminus \gamma$ are both clopen subsets of $E(\Sigma)$, and that these clopen subsets form a basis for the topology of $E(\Sigma)$.

By picking clopen neighborhoods of this kind for each end in V and applying compactness, we can describe V as a disjoint union of clopen sets, each of which is bounded by a curve. These can then be combined so that V is bounded by a single curve, as in Figure 3.6. Call this curve η. Draw an arc λ connecting the curves α and η, and let β be the curve following along α, λ, and η as in Figure 3.7. By construction, β separates e_+ and e_-, and $E((\alpha, \beta)) = V$.

If no end of V is accumulated by genus but e_+ is, then (α, β) must have finite genus and β_+ must have infinite genus. By picking a curve ζ separating a single handle from the rest of Σ and then applying the construction of Figure 3.7 with β replacing α and ζ replacing η, we get a new curve β' such that (α, β) and (α, β') have the same end space but genus differing by 1. Doing this finitely many times lets us achieve arbitrary genus for (α, β) in this case. $\square$

Now that we can construct appropriate subsets, we will be able to build paths between curves in $\mathcal{TC}(\Sigma)$.

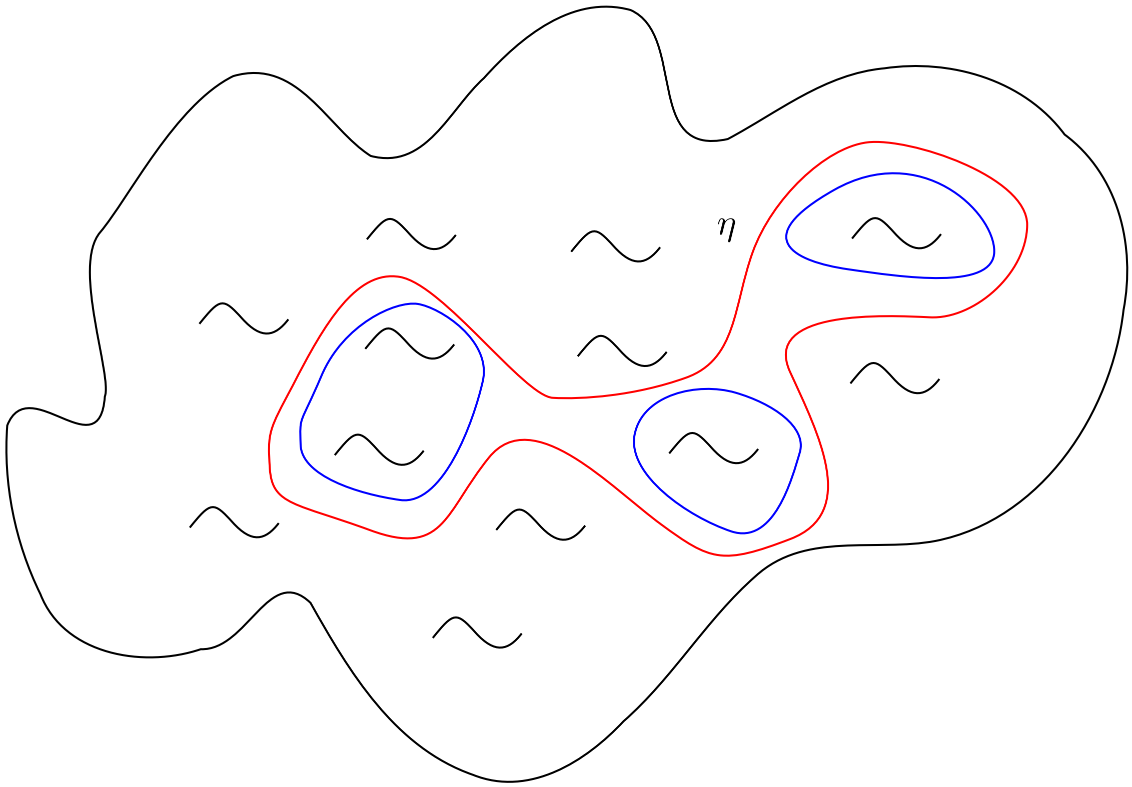

Figure 3.6: If V is the union of finitely many sets bounded by blue curves, we can find a single red curve η bounding V. Each squiggle represents a possibly complicated clopen set of ends.

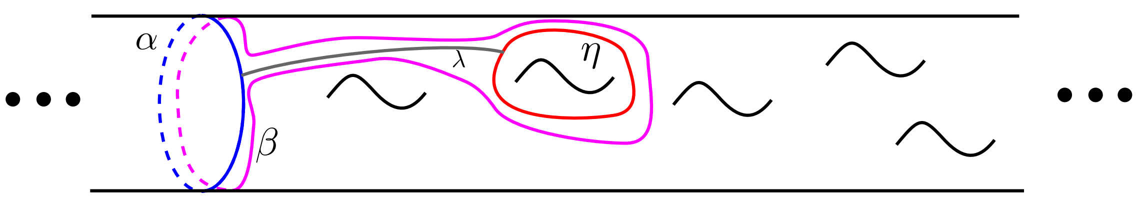

Figure 3.7: Given α in blue, η in red, and λ in gray, we draw β in magenta. Again, each squiggle represents a possibly complicated clopen set of ends.

Lemma 3.4.7. *If Σ is a translatable surface with tame end space then $\mathcal{TC}(\Sigma)$ is connected.*

Proof. Given $\alpha, \beta \in \mathcal{TC}(\Sigma)$, let γ be a curve in $\alpha_+ \cap \beta_+$ such that the subsurfaces (α, γ) and (β, γ) both have end spaces containing representatives of each $E(f_i)$ and $E(c_i)$, using the representatives $\{f_1, \ldots, f_n, c_1, \ldots, c_m\}$ defined above; such a curve γ can always be found by looking in a small enough neighborhood of e_+. We will show that α is connected to γ; by symmetry, this will imply that β is connected to γ and so α is connected to β.

Let g equal the genus of the subsurface (α, γ) if it has finite genus, and 0 if it has infinite genus. By Lemma 3.4.5 the end space of (α, γ) can be written as $\bigsqcup_{i=1}^{k} V_i$ where each V_i is a stable neighborhood of a maximal end. Let's rewrite this disjoint union as $(\bigsqcup_{i=1}^{m} V_i) \sqcup \left(\bigsqcup_{j=1}^{\ell} V_j\right)$ where the first m stable neighborhoods contain points equivalent to each of the c_i, and the subsequent ℓ stable neighborhoods contain points equivalent to the f_i. Since $E(c_i) \cup V_i$ is a Cantor set, we can identify $\ell + g$ elements of $E(c_i)$ inside V_i and write V_i as $\bigsqcup_{j=1}^{\ell+g} V_{i,j}$ where each $V_{i,j}$ is homeomorphic to V_i. Then for each $1 \leq j \leq \ell$, let $W_j = V_j \sqcup \bigsqcup_{i=1}^{m} V_{i,j}$, and for $\ell + 1 \leq j \leq \ell + g$ let $W_j = \bigsqcup_{i=1}^{m} V_{i,j}$. By construction, the end space of (α, γ) is $\bigsqcup_{j=1}^{\ell} W_j$.

Finally, define $\{\alpha_0 = \alpha, \alpha_1, \ldots, \alpha_{\ell+g} = \gamma\}$ such that the end space of (α_{j-1}, α_j) is W_j and such that the genus of (α_{j-1}, α_j) is 0 or infinite for $j \leq \ell$ and 1 for $j \geq \ell$, which is possible by Lemma 3.4.6. Then each (α_{j-1}, α_j) is homeomorphic to one of the subsurfaces $T_i \in \mathcal{S}$, and so α_{j-1} and α_j are adjacent in $\mathcal{TC}(\Sigma)$. $\square$

This leaves us quite close to fulfilling all the hypotheses of Lemma 3.2.2. We need one more ingredient:

Lemma 3.4.8. *Let Σ be a translatable surface with tame end space. Then for any vertex α of $\mathcal{TC}(\Sigma)$, the set $H = \{f \in \mathrm{MCG}(\Sigma) \mid d(f(\alpha), \alpha) \leq 1\}$ is coarsely bounded.*

Proof. We will start by defining some helpful mapping classes. Using Proposition 3.3.5, let $\Sigma = S^{\natural\mathbb{Z}} = \natural_{j\in\mathbb{Z}} S_j$ where all the S_j are homomorphic and $S_0 = (\alpha, h^N(\alpha))$ for some N. As above let $\{f_1, \ldots, f_n, c_1, \ldots, c_m\}$ be representatives of the equivalence classes of maximal ends of S_0. We may choose these so that each f_i and c_i is actually an end of S_0 itself. Let $\{V_{f_1}, \ldots, V_{f_n}, \ldots, V_{c_1}, \ldots, V_{c_m}\}$ be a set of disjoint stable neighborhoods of these ends, also contained in the end space of S_0. For each x equal to some f_i or c_i and each $j \in \mathbb{Z}$, let $V_{x,j}$ be the homeomorphic copy of V_x in the end space of the subsurface S_j.

Note that the sequence of sets $V_{f_i,j}$ converges to $e_\pm$ as j goes to $\infty_\pm$, and likewise for $V_{c_i,j}$. That means that for each maximal end x equal to some f_i or c_i there is a homeomorphism of the end space of Σ taking each $V_{x,j}$ to $V_{x,j+1}$ and fixing the rest of the end space pointwise. This homeomorphism of the end space extends to a homeomorphism h_x of Σ; if the end x is not accumulated by genus we may also construct h_x so that $(\alpha, h_x(\alpha))$ has genus 0. For each $1 \leq i \leq n$ let $h_i = h_{f_i} \circ \prod_{k=1}^m h_{c_k}$. If S_0 has finite positive genus, let h_{genus} be a map that moves a single handle from each S_j to S_{j+1} and fixes the end space of Σ, and then let $h_{n+1} = h_{\text{genus}} \circ \prod_{k=1}^m h_{c_k}$.

Observe that by construction $(\alpha, h_i(\alpha))$ and $(h_i^{-1}(\alpha), \alpha)$ are both homeomorphic to the subsurface $T_i \in \mathcal{S}$, where $\mathcal{S}$ is the canonical set of subsurfaces used to define $\mathcal{TC}(\Sigma) = \mathcal{TC}(\Sigma, \mathcal{S})$.

We are now ready to prove that H is coarsely bounded using Lemma 3.2.1. Let V be an identity neighborhood in $\mathrm{MCG}(\Sigma)$, and find $n \in \mathbb{N}$ as in Lemma 3.3.8 so that $V_{\alpha_-} \subseteq h^{-n}Vh^n$ and $V_{\alpha_+} \subseteq h^nVh^{-n}$. Let $F = \{r^{-1}, h^{\pm n}, h_1^{\pm 1}, \ldots, h_n^{\pm 1}, (h_{n+1}^{\pm 1})\}$, where r is the mapping class defined in Lemma 3.3.7 and including the maps $h_{n+1}^{\pm 1}$ if they are defined. We claim that $H \subseteq (FV)^8$.

Let $f \in H$. If $d(\alpha, f(\alpha)) = 0$, then $f \in (FV)^5$ as shown in Corollary 3.3.9. If not, then $d(\alpha, f(\alpha)) = 1$ and so α and $f(\alpha)$ have disjoint representatives. Assume $f \subseteq \alpha_+$—if not, we will merely have to reverse some signs. By the definition of adjacency in $\mathcal{TC}(\Sigma) = \mathcal{TC}(\Sigma, \mathcal{S})$, we know that $(\alpha, f(\alpha))$ is homeomorphic to some $T_i \in \mathcal{S}$. Since $(\alpha, f(\alpha))$ is homeomorphic to $(\alpha, h_i(\alpha))$ by construction, and $f(\alpha)_+$ is homeomorphic to $h_i(\alpha)_+$ by Lemma 3.3.6, there is a map g taking $h_i(\alpha)$ to $f(\alpha)$ and restricting to the identity on α_-—that is, $g(h_i(\alpha)) = f(\alpha)$ and $g \in V_{\alpha_-} \subseteq (FV)^2$.

Let $f_0 = h_i^{-1}g^{-1}f$. By construction, $f_0(\alpha) = \alpha$, so by Corollary 3.3.9 $f_0 \in (FV)^5$. Then we have $f = gh_if_0$, where $g \in (FV)^2$, $h_i \in F$, and $f_0 \in (FV)^5$, so $f \in (FV)^8$. $\square$

Putting this all together gives us

Theorem 3.4.9. *If Σ is a translatable surface with tame end spcae, then $\mathcal{TC}(\Sigma)$ equipped with the edge metric is quasi-isometric to* $\mathrm{MCG}(\Sigma)$.

Proof. The group $\mathrm{MCG}(\Sigma)$ is locally coarsely bounded by Corollary 3.3.10. The graph $\mathcal{TC}(\Sigma)$ is connected by Lemma 3.4.7. The action of $\mathrm{MCG}(\Sigma)$ on it is transitive by Lemma 3.3.6. Finally, the set of mapping classes that moves a vertex a distance at most 1 is coarsely bounded by Lemma 3.4.8. Thus by Lemma 3.2.2 the action induces a quasi-isometry. □

3.5 Equivalent definitions of translatability

We have just seen that a translatable surface Σ with tame end space is quasi-isometric to the translatable curve graph $\mathcal{TC}(\Sigma)$, which is a graph whose vertices are curves. This section establishes that the existence of such a graph is nearly unique to translatable surfaces.

We say "nearly unique" because there is one other example: if $\mathrm{MCG}(\Sigma)$ is coarsely bounded, then it is quasi-isometric to any finite-diameter graph. In particular, the curve graph $\mathcal{C}(\Sigma)$ of an infinite-type surface always has diameter 2, giving a trivial quasi-isometry. For this reason, coarsely bounded mapping class groups are excluded in the hypothesis of Theorem 3.5.3.

Another condition which we show to be equivalent to translatability is the following, due to Horbez, Qing, and Rafi [HQR20].

Definition 3.5.1 (Definition 1.8 of [MR19])**.** A connected, finite-type subsurface S of a surface Σ is called *nondisplaceable* if $f(S) \cap S \neq \emptyset$ for each $f \in \mathrm{MCG}(\Sigma)$. A non-connected surface is nondisplaceable if, for every $f \in \mathrm{MCG}(\Sigma)$ and every connected component S_i of S, there is a connected component S_j of S such that

$f(S_i) \cap S_j \neq \emptyset$.

Definition 3.5.2 (Definition 4.4 of [HQR20])**.** An *avenue surface* is a connected, orientable surface Σ which does not contain any nondisplaceable finite-type subsurfaces, whose end space is tame, and whose mapping class group $\mathrm{MCG}(\Sigma)$ admits a coarsely bounded generating set but is not itself coarsely bounded.

Theorem 3.5.3. *Let Σ be an infinite-type surface with tame end space such that $\mathrm{MCG}(\Sigma)$ is locally coarsely bounded and admits a coarsely bounded generating set—and thus has a well-defined quasi-isometry type—but is not itself coarsely bounded. Then the following are equivalent:*

1. *There exists a graph Γ whose vertices are curves and such that the action of $\mathrm{MCG}(\Sigma)$ on Γ induces a quasi-isometry.*

2. *Σ is translatable.*

3. *Σ is an avenue surface.*

Proof. $2 \implies 1$: This is Theorem 3.4.9.

$2 \implies 3$: We are already assuming that Σ has tame end space and that $\mathrm{MCG}(\Sigma)$ admits a coarsely bounded generating set but is not itself coarsely bounded. By the definition of a translation map, a translatable surface cannot have any finite-type nondisplaceable surfaces, and so Σ is an avenue surface.

$3 \implies 2$: Lemma 4.5 of [HQR20] says that an avenue surface has zero or infinite genus and exactly two maximal ends, while Lemma 4.6 of [HQR20] says that every nonmaximal end of an avenue surface precedes both maximal ends under the standard ordering. It follows by Lemma 3.5.9 that Σ is translatable.

$1 \implies 2$: We divide our work into three cases, depending on the genus and maximal ends of Σ:

1. If Σ has zero or infinite genus and one or a Cantor set of equivalent maximal ends, then by Corollary 3.5.6 the group $\mathrm{MCG}(\Sigma)$ is coarsely bounded, which is excluded by the hypothesis of the theorem.

2. If Σ has zero or infinite genus and two equivalent maximal ends, then by Proposition 3.5.7 it is translatable.

3. If Σ has finite positive genus or any other structure of maximal ends, then by Proposition 3.5.17 there is no graph whose vertices are curves and on which the action of $\mathrm{MCG}(\Sigma)$ induces a quasi-isometry, contradicting our assumption.

Thus the only remaining possibility is that Σ is translatable. □

The following three subsections correspond to the three cases in the last step of the proof of Theorem 3.5.3.

3.5.1 Coarsely bounded mapping class groups

The first case is essentially a rehash of the following facts. Recall that $\mathcal{M}(\Sigma)$ is the set of maximal ends of Σ.

Fact 3.5.4 (Proposition 3.1 of [MR19]). *If Σ has zero or infinite genus and self-similar end space, then* $\mathrm{MCG}(\Sigma)$ *is coarsely bounded.*

Fact 3.5.5 (Proposition 4.8 of [MR19]). *If Σ has no nondisplaceable finite-type subsurfaces and $\mathcal{M}(\Sigma)$ consists of either a singleton or a Cantor set of equivalent ends, then its end space is self-similar.*

To link these two facts together we need to add the assumption of tameness:

Corollary 3.5.6. *If Σ has zero or infinite genus and tame end space, and $\mathcal{M}(\Sigma)$ consists of either a singleton or a Cantor set of equivalent ends, then* $\mathrm{MCG}(\Sigma)$ *is coarsely bounded.*

Proof. Given the previous facts, we need only show that Σ has no nondisplaceable finite-type subsurfaces. Let S be a finite-type subsurface of Σ. By expanding S we may assume that S is connected with all its boundary curves essential and separating. Let $E_1, \dots, E_n$ be the end spaces of the complementary components of S, and E_0 the end space of S itself, which may be empty or contain a finite set of punctures. Since $E_0 \sqcup \cdots \sqcup E_n = E$, there is a maximal end x in some E_i; without loss of generality we may assume $x \in E_n$. Let α be the boundary component of S corresponding to E_n.

Since Σ has tame end space, E_n contains a stable neighborhood of x; and since for every end y, we know that x is an accumulation point of $E(y)$, we can use Fact 3.2.7 and compactness to find a clopen subset $F \subseteq E_n$ homeomorphic to $E \setminus E_n$.

Let β be a separating curve whose complementary components have end spaces F and $E \setminus F$, and such that the component of $\Sigma \setminus \beta$ with end space F has the same genus as the component of $\Sigma \setminus \alpha$ containing S.

The complementary components of the curves α and β have the same genus and end space by construction, so we can find some $f \in \mathrm{MCG}(\Sigma)$ exchanging α and β. Then $f(S)$ and S are in distinct components of $\Sigma \setminus \alpha$, and so the subsurface S is not nondisplaceable; thus Σ has no nondisplaceable finite-type subsurfaces. It follows that the end space of Σ is self-similar, and so $\mathrm{MCG}(\Sigma)$ is coarsely bounded. $\square$

3.5.2 Translatable surfaces

Proposition 3.5.7. *Suppose Σ has tame end space, zero or infinite genus, and exactly two equivalent maximal ends, e_+ and e_-. If $\mathrm{MCG}(\Sigma)$ is locally coarsely bounded and has a coarsely bounded generating set, then Σ is translatable with respect to the ends e_+ and e_-.*

Our first step towards proving Proposition 3.5.7 will be to find the immediate predecessors of the maximal ends of Σ.

Lemma 3.5.8. *Let Σ be a surface with tame end space and two maximal ends, e_+ and e_-. If $\mathrm{MCG}(\Sigma)$ is loaclly coarsely bounded and admits a coarsely bounded generating set, then there is a finite set of ends $x_1, \ldots, x_n$ such that each x_i is an immediate predecessor of e_+, and every end $y \prec e_+$ satisfies $y \preccurlyeq x_i$ for some i.*

Proof. Find K as in Fact 3.2.9, with complementary region A containing the end

e_+, and let U_0 be the end space of A. Fix $y \prec e_+$; by possibly replacing y with an equivalent end, we may assume $y \in U_0$. Let U_1 be a clopen subset of $U_0 \setminus \{y\}$ containing e_+, and construct a neighborhood basis $U_0 \supseteq U_1 \supseteq U_2 \supseteq \cdots$ of clopen sets such that $\bigcap_{n \in \mathbb{N}} U_n = \{e_+\}$.

The set $\{x \in U_0 \setminus U_1 \mid y \preccurlyeq x\}$ is compact and nonempty, and so it has a (not necessarily unique) maximal element which we will call z_0. Similarly, for $n > 0$ let z_n be a maximal element of the set $\{x \in U_n \setminus U_{n-1} \mid z_{n-1} \preccurlyeq x\}$.

We claim there is some z such that $z_n \sim z$ for all sufficiently high n. If not, then up to taking a subsequence we may assume the z_n are pairwise inequivalent. By construction each $z_n \in U_n$, so $E(z_n) \cap U_n \neq \emptyset$. For any $m < n$, $z_m \prec z_n$ but z_m was maximal in $U_m \setminus U_{m+1}$, so $E(z_n) \cap (U_m \setminus U_{m+1}) = \emptyset$, and thus in general $E(z_n) \cap (U_0 \setminus U_n) = \emptyset$, or in other words $E(z_n) \subseteq (U_n \cup U_0^c)$. Finally, let B be a subsurface of A with end space U_{n+1}. By Fact 3.2.9 there is a mapping class $f \in \mathrm{MCG}(\Sigma)$ so that $A \subseteq f(B)$. Then $f(z_n) \in U_0^c \cap E(z_n)$ so this set is not empty.

Let $G = \{f \in \mathrm{MCG}(\Sigma) \mid f(e_+) = e_+\}$. Since the only end of Σ that might be equivalent to e_+ is e_-, G has index at most two in $\mathrm{MCG}(\Sigma)$, and the set $\{e_+\}$ is G-invariant. Thus we have fulfilled the definition of limit type, and so by Fact 3.2.11 $\mathrm{MCG}(\Sigma)$ cannot admit a coarsely bounded generating set. Since we assumed otherwise, this is a contradiction. This proves our claim that $z_n \sim z$ for all sufficiently high n. This z must be an immediate predecessor of e_+ (otherwise it would not be maximal in some $U_n \setminus U_{n+1}$) and by construction $y \preccurlyeq z$.

We now claim that there is a clopen subset $F \subseteq U_0$ such that $e_+ \notin F$ but for every immediate predecessor z of e_+, $F \cap E(z) \neq \emptyset$. Suppose not. Then we can pick a sequence of immediate predecessors $\{z_n\}_{n\in\mathbb{N}}$ of e_+ and clopen sets $U_0 \supseteq U_1 \supseteq U_2 \supseteq \cdots$ with $\bigcap_{n\in\mathbb{N}} U_n = \{e_+\}$ such that each $z_n \in U_n$ but $z_n \notin (U_0 \setminus U_n)$. As above we can use Fact 3.2.9 to find an element of $E(z_n)$ in U_0^c, so this would again show that E has limit type, a contradiction by Fact 3.2.11. This proves our claim.

For each end $y \in F$, let $x_y \in E$ be an immediate predecessor of e_+ with $y \preccurlyeq x_y$. That means that for a stable neighborhood V_{x_y} of x_y—which must exist because E is tame—there is some clopen neighborhood U_y of y such that U_y is homeomorphic to a clopen subset of V_{x_y}. The sets $\{U_y\}_{y\in F}$ cover the compact set F, so we can pick a finite collection $U_1, \ldots, U_n$ with corresponding $x_1, \ldots, x_n$ predecessors to e_+ so that the U_i cover F and each U_i is homeomorphic to a clopen subset of V_{x_i}. In particular, by stability $z \preccurlyeq x_i$ for every $z \in U_i$ and so every end in F is bounded above by one of the x_i. If z is an immediate predecessor of e_+, then by construction there is some $z' \sim z$ in F, and so $z \sim z' \preccurlyeq x_i$ for some i. Since z is an immediate predecessor of e_+, it follows that $z \sim x_i$, so there are only finitely many equivalence classes of immediate predecessors. $\square$

The following lemma is nearly identical to Proposition 3.5.7; we list it separately so that it can be used in other parts of this section.

Lemma 3.5.9. *Let Σ be a surface of zero or infinite genus with tame end space and* *two maximal ends, e_+ and e_-, with the property that for every end $y \in E \setminus \{e_+, e_-\}$, $y \preccurlyeq e_+$ and $y \preccurlyeq e_-$. If* $\mathrm{MCG}(\Sigma)$ *is locally coarsely bounded and admits a coarsely bounded generating set, then* Σ *is translatable.*

Proof. First, note that e_+ is accumulated by genus if and only if e_- is, by the following argument: suppose e_+ is accumulated by genus but e_- is not. If some $y \prec e_+$ were accumulated by genus, then e_- would have to be as well since it is an accumulation point of $E(y)$. So e_+ is the only end of Σ accumulated by genus. Since $\mathrm{MCG}(\Sigma)$ is locally coarsely bounded, we can find a surface K as in Fact 3.2.9. Let A be the component of $\Sigma \setminus K$ containing e_+, and note that A is the only such component with nonzero genus. Pick a subsurface $V \subseteq A$ containing e_+ and such that $A \setminus V$ has positive genus. By Fact 3.2.9, there is some $f \in \mathrm{MCG}(\Sigma)$ such that $A \subseteq f(V)$. But that would mean $\Sigma \setminus V$ has positive genus while $\Sigma \setminus f(V)$ has zero genus, a contradiction. Thus e_+ is accumulated by genus if and only if e_- is.

Let $x_1, \ldots, x_k$ be the immediate predecessors to e_+ found via Lemma 3.5.8. Fix a curve α_0 separating e_+ and e_-. Pick a sequence of pairwise disjoint curves $\alpha_1, \alpha_2, \ldots$ such that $\lim_{n\to\infty} \alpha_n = e_+$—this is possible by definition for any end—and likewise $\alpha_{-1}, \alpha_{-2}, \ldots$ such that $\lim_{n\to-\infty} \alpha_n = e_-$. By moving to a subsequence, we may assume that for each $n \in \mathbb{Z}$ the subsurface (α_n, α_{n+1}) has positive (possibly infinite) genus if Σ has infinite genus, and furthermore that the end space of this subsurface includes an element of each $E(x_i)$.

By Lemma 3.4.5, write the end space of (α_n, α_{n+1}) as $\bigsqcup_{j=1}^{p} V_{n,j}$, where each $V_{n,j}$

is a stable neighborhood of an immediate predecessor of e_+. Note that each $V_{n,j}$ is homeomorphic to some U_i, a stable neighborhood of the immediate predecessor x_i. Thus we can describe the end space of Σ as follows:

$$E = \{e_-, e_+\} \sqcup \bigsqcup_{n\in\mathbb{Z}} \left(\bigsqcup_{i=1}^{k} U_{i,n} \right)$$

where each $U_{i,n}$ is a copy of U_i, and the only additional topology is given by the limits

$$\lim_{n\to\pm\infty} V_{i,n} = e_\pm$$

for each $1 \leq i \leq k$.

Let S be a surface with end space $\bigsqcup_{i=1}^{k} U_i$ and genus defined as follows: if some x_i is accumulated by genus, S will have infinite genus by definition. If Σ has zero genus, let S also have zero genus. If Σ has infinite genus but no x_i is accumulated by genus—which implies that only e_+ and e_- are accumulated by genus—then let S have genus 1. Then the surface $S^{\natural\mathbb{Z}}$, which is translatable by construction, has genus and end space matching that of Σ, and so they are homeomorphic. Thus Σ is translatable. □

The proof of Proposition 3.5.7 follows directly:

Proof of Proposition 3.5.7. Since e_+ and e_- are the only maximal ends of Σ, and $e_+ \sim e_-$, every end $y \in E$ has $y \preccurlyeq e_+$ and $y \preccurlyeq e_-$. Then we can apply Lemma 3.5.9. □

3.5.3 All other surfaces

Remark 3.5.10. Many of the proofs in this subsection are inspired by and to some extent duplicate the work in Mann and Rafi's proof of Fact 3.2.9. They are included for completeness.

We have covered the cases where Σ has zero or infinite genus and either one maximal end, two equivalent maximal ends, or a Cantor set of equivalent maximal ends. We now show that if the maximal ends of Σ have any other structure, or if Σ has finite positive genus, there is no graph whose vertices are curves onto which the action of $\mathrm{MCG}(\Sigma)$ induces a quasi-isometry. Our main tool will be the following observation of Mann and Rafi:

Fact 3.5.11 (Lemma 5.2 of [MR19])**.** *Let $K \subseteq \Sigma$ be a finite-type subsurface. If there exists a finite-type, nondisplaceable (possibly disconnected) subsurface $S \subseteq \Sigma \setminus K$, then V_K is not coarsely bounded.*

Corollary 3.5.12. *Let Σ be a surface. If for every curve α on Σ, $\Sigma \setminus \alpha$ contains a finite-type nondisplaceable surface, then there is no graph Γ whose vertices are curves on Σ such that the orbit map $\mathrm{MCG}(\Sigma) \to \Gamma$ is a quasi-isometry.*

Proof. For the orbit map to be a quasi-isometry, the preimage of every bounded set in Γ must be coarsely bounded in $\mathrm{MCG}(\Sigma)$. In particular, the stabilizer of a curve is the preimage of a single vertex, so it must be coarsely bounded. □

Mann and Rafi give three basic examples (Examples 2.4 and 2.5 of [MR19]) of nondisplaceable surfaces, all of which we will use:

1. If Σ has finite positive genus, then any subsurface of Σ with the same genus as Σ is nondisplaceable.

2. If X is a $\mathrm{MCG}(\Sigma)$-invariant, finite set of ends of Σ of cardinality at least 3, then any surface that separates the elements of X into different complementary components is nondisplaceable.

3. If X and Y are disjoint, closed $\mathrm{MCG}(\Sigma)$-invariant sets of ends of Σ with X homeomorphic to a Cantor set, then a subsurface homeomorphic to a pair of pants containing elements of X in two complementary components, and all of Y in the third, is nondisplaceable.

The easiest place to apply Corollary 3.5.12 is in the case of finite-genus surfaces:

Lemma 3.5.13. *If Σ has finite positive genus, then for any graph Γ whose vertices are curves on Σ, the orbit map $\mathrm{MCG}(\Sigma) \to \Gamma$ is not a quasi-isometry.*

Proof. Let S be a connected, finite-type subsurface of Σ with the same genus as Σ. If α is disjoint from S, or if α is nonseparating in S, $S \setminus \alpha$ is still connected and nondisplaceable. If α separates S into two components, one of which has the same genus as Σ, then that component is connected and nondisplaceable.

Finally, if α separates S into two components, both of which have positive genus, consider the surface $S \setminus \alpha$. For any $f \in \mathrm{MCG}(\Sigma)$, both components of $f(S \setminus \alpha)$ contain nonseparating curves, and every nonseparating curve on Σ intersects $S \setminus \alpha$.

Therefore both components of $f(S \setminus \alpha)$ intersect $S \setminus \alpha$, making it a nondisplaceable surface. The result follows by Corollary 3.5.12. $\square$

Next consider the case where Σ has at least three—but finitely many—maximal ends.

Lemma 3.5.14. *If Σ has at least 3 but finitely many maximal ends, then for any graph Γ whose vertices are curves on Σ, the orbit map $\mathrm{MCG}(\Sigma) \to \Gamma$ is not a quasi-isometry.*

Proof. Let S be a finite-type surface separating the maximal ends of Σ into distinct complementary components. If α is nonseparating in or disjoint from S, then $S \setminus \alpha$ is still connected and nondisplaceable. If α separates S into two components, one of which still separates the maximal ends of Σ into distinct complementary components, then that component is connected and nondisplaceable.

Otherwise, there are at least two maximal ends of Σ in both components of $\Sigma \setminus \alpha$. Fix $f \in \mathrm{MCG}(\Sigma)$. Since there are at least two maximal ends of Σ in both components of $\Sigma \setminus f(\alpha)$, either $f(\alpha) = \alpha$ or $f(\alpha)$ intersects $S \setminus \alpha$. Since $f(\alpha)$ is a boundary component of both components of $S \setminus \alpha$, it follows that $S \setminus \alpha$ is nondisplaceable. The result follows by Corollary 3.5.12. $\square$

Now we move to the case of infinitely many maximal ends:

Lemma 3.5.15. *If Σ has infinitely many maximal ends, not all equivalent, then for any graph Γ whose vertices are curves on Σ, the orbit map $\mathrm{MCG}(\Sigma) \to \Gamma$ is not a* *quasi-isometry.*

Proof. By Fact 3.2.4 the equivalence class of every maximal end is either finite or a Cantor set. If every such equivalence class is finite then there are infinitely many of them; in particular let x, y, z be three nonequivalent maximal ends with $E(x)$, $E(y)$, and $E(z)$ all finite. Then let $X = E(x)$, $Y = E(y)$, and $Z = E(z)$. If on the other hand there is some maximal end x such that $E(x)$ is a Cantor set, pick a maximal end z not equivalent to x, and let $X \sqcup Y$ be nonempty sets partitioning $E(x)$, and let $Z = E(z)$.

In either case above, let S be a finite-type surface with X, Y, and Z in distinct complementary components. For any curve α, one component of $S \setminus \alpha$ still has X, Y, and Z in distinct complementary components, so we may assume S is contained in $\Sigma \setminus \alpha$. By construction, S is connected and nondisplaceable, and the result follows by Corollary 3.5.12. $\square$

There is only one more case, which requires a bit more subtlety as well as the condition of tameness:

Lemma 3.5.16. *If Σ has tame end space and two non-equivalent maximal ends e_+ and e_-, then for any graph Γ whose vertices are curves on Σ, the orbit map $\mathrm{MCG}(\Sigma) \to \Gamma$ is not a quasi-isometry.*

Proof. If $\mathrm{MCG}(\Sigma)$ does not have a well-defined quasi-isometry type, then such a quasi-isometry cannot be defined and we are done; so we may assume it does. Then we can apply Lemma 3.5.8 to find immediate predecessors to e_+ and e_-. If e_+ and e_- had the same predecessors, Σ would be translatable by Lemma 3.5.9, which would imply $e_+ \sim e_-$, a contradiction. Thus without loss of generality we may assume there is some immediate predecessor x of e_+ such that $x \not\preccurlyeq e_-$.

Let V be a stable neighborhood of x. Since x is maximal in V, $E(x) \cap V$ is either a singleton or a Cantor set. We claim it is in fact a Cantor set. Suppose by contradiciton that $E(x) \cap V$ is discrete; since x is an immediate predecessor of e_+ and $x \not\preccurlyeq e_-$, this means that $E(x)$ is countable, with a unique accumulation point at e_+. Find a subsurface K as in Fact 3.2.9, with complementary components A_+ and A_- containing e_+ and e_- respectively. Note that all but finitely many elements of $E(x)$ are in the end set of A_+. Let B be a subsurface of A_+ containing e_+, and such that the end space of $A_+ \setminus V$ contains a single element of $E(x)$. Then there is some $f \in \mathrm{MCG}(\Sigma)$ such that $A_+ \subseteq f(B)$. But $\Sigma \setminus A_+$ and $\Sigma \setminus B$ have a different number of elements of $E(x)$. This contradiction proves our claim.

Since $E(x) \cap V$ is a Cantor set, x is an immediate predecessor of e_+, and $x \not\preccurlyeq e_-$, $E(x)$ must be a countable sequence of disjoint Cantor sets converging to e_+. Let $X \sqcup Y$ be a partition of $E(x) \cup \{e_+\}$ into nonempty clopen sets, and let $Z = \{e_-\}$. Then a finite-type surface S that has X, Y, and Z in distinct complementary components will be nondisplaceable. As in the proof of Lemma 3.5.15, removing a single curve α from S does not change this property, and so the result follows by Corollary 3.5.12. $\square$

These lemmas together give the main result of this subsection:

Proposition 3.5.17. *If Σ has tame end space and either finite positive genus, two or infinitely many maximal ends that are not all equivalent, or at least three but finitely many maximal ends, then for any graph Γ whose vertices are curves on Σ, the orbit map $\mathrm{MCG}(\Sigma) \to \Gamma$ is not a quasi-isometry.*

Proof. If Σ has finite positive genus, this is Lemma 3.5.13. If it has two maximal ends, this is Lemma 3.5.16. If it has at least three but finitely many maximal ends, this is Lemma 3.5.14. If it has infinitely many maximal ends, this is Lemma 3.5.15. □

3.6 The plane minus a Cantor set

We now turn from the general case of translatable surfaces, of which there are uncountably many examples only a few of which have received specific notice, to a much more specific but more well-studied case. In this section we focus exclusively on the surface $\Sigma = \mathbb{R}^2 \setminus C$, where C is a Cantor set embedded in the plane. In this instance we will not have to go looking for a suitable graph, as one has been provided for us in the form of the *loop graph* defined by Bavard [Bav16]. We will show in this section that the mapping class group of this surface is quasi-isometric to its loop graph.

Note that the surface Σ has a unique isolated end, usually called ∞ because it is the "point at infinity" of $\mathbb{R}^2$.

Definition 3.6.1. A *loop* in Σ is an embedded line in Σ with both ends approaching ∞, considered up to isotopy and orientation reversal. The *loop graph* $\mathcal{L}(\Sigma)$ of Σ is the graph whose vertices are loops in Σ, with two loops connected by an edge if they have disjoint representatives.

It was shown by Bavard [Bav16] that the loop graph[5] is connected and Gromov-hyperbolic. A subsequent paper of Bavard and Walker [BW18] characterized the Gromov boundary of $\mathcal{L}(\Sigma)$. The high degree of symmetry possessed by Σ also makes the following transitivity lemma possible.

Lemma 3.6.2. *If α, β, and γ are loops in Σ, with β and γ both in the same component of $\Sigma \setminus \alpha$, then there is a mapping class $f \in \mathrm{MCG}(\Sigma)$ such that $f(\alpha) = \alpha$, $f(\beta) = \gamma$, and f restricts to the identity on the component of $\Sigma \setminus \alpha$ not containing β and γ.*

Proof. Observe that every loop on Σ is separating. If we cut Σ along α and β, we get three subsurfaces: one whose only boundary component is α, one whose only boundary component is β, and one with both boundary components. The end space of each of these subsurfaces is a nonempty clopen subset of a Cantor set, which must be itself a Cantor set. Since the surface Σ has no genus, this is a complete

[5]Many papers deal interchangably with the loop graph as defined here and the *ray graph*, whose vertices are embedded lines with one end at ∞ and the other in the Cantor set. These are shown by Bavard [Bav16] to be quasi-isometric.

description of the subsurfaces. The same argument applies when cutting the surface along α and γ, so we can fix the surface bounded by α, map the surface bounded by β to that bounded by γ, and map the surface bounded by both α and β to that bounded by α and γ. $\square$

The following lemmas are analogs of Lemmas 3.3.7 and 3.3.8 in the setting of Σ:

Lemma 3.6.3. *Let α be a loop on Σ and α_- and α_+ the two components of $\Sigma \setminus \alpha$. Then there is a mapping class $r \in \mathrm{MCG}(\Sigma)$ such that after an isotopy $r(\alpha_+) = \alpha_-$, $r(\alpha_-) = \alpha_+$, and $r_{|\alpha}$ is orientation-reversing.*

Proof. In a tubular neighborhood of α, which is a punctured anulus, r is just a rotation by π about the line running down the middle of that punctured annulus, as in Figure 3.4. Since α separates the end space of Σ into two nonempty clopen sets, the end spaces of α_- and α_+ are homeomorphic and so this r can be extended to all of Σ. $\square$

Lemma 3.6.4. *Let α be a loop on Σ and V an identity neighborhood in $\mathrm{MCG}(\Sigma)$. Let α_- and α_+ be the two components of $\Sigma \setminus \alpha$. Then there are mapping classes $h_+, h_- \in \mathrm{MCG}(\Sigma)$ such that $V_{\alpha_-} \subseteq h_+^{-1} V h_+$ and $V_{\alpha_+} \subseteq h_-^{-1} V h_-$. In addition, $h_+(\alpha) \subseteq \alpha_+$ and $h_-(\alpha) \subseteq \alpha_-$.*

Proof. Since the sets $\{V_S \mid S \subseteq \Sigma \text{ has finite type}\}$ form a neighborhood basis of the identity in $\mathrm{MCG}(\Sigma)$, there is some finite-type $S \subseteq \Sigma$ such that $V_S \subseteq V$. By growing

S—and thus shrinking V_S—we can ensure that the loop α and its basepoint are included in S.

S is a finite-type surface of genus zero, with n boundary components for some n. In particular, it must have at least one boundary component in α_- and at least one boundary component in α_+. Pick two new arcs $\beta \subseteq \alpha_-$ and $\gamma \subseteq \alpha_+$ and such that both $\beta_- \cap S$ and $\gamma_+ \cap S$ are disks, as in Figure 3.8. Using Lemma 3.6.2, let $h_+, h_- \in \mathrm{MCG}(\Sigma)$ such that h_+ fixes β and maps α to γ, while h_- fixes γ and maps α to β.

It is not quite true that $S \subseteq h_+(\alpha)_- = \gamma_-$ as in the proof of Lemma 3.3.8. However, the intersection $S \cap \gamma_+$ is a disk, and so any homeomorphism that restricts to the identity on γ_- can be homotoped to one restricting to the identity on S, and thus $V_{(h_+(\alpha))_-} = V_{\gamma_-} \subseteq V_S \subseteq V$. It follows that $V_{\alpha_-} \subseteq h_+^{-1} V h_+$ and likewise $V_{\alpha_+} \subseteq h_-^{-1} V h_-$. □

These are enough ingredients to prove our main theorem for this section:

Theorem 3.6.5. *Let $\Sigma = \mathbb{R}^2 \setminus C$ be the plane minus a Cantor set. Then* $\mathrm{MCG}(\Sigma)$ *is quasi-isometric to $\mathcal{L}(\Sigma)$.*

Proof. The loop graph is known to be connected by work of Bavard [Bav16], and the action of $\mathrm{MCG}(\Sigma)$ on it is transitive by Lemma 3.6.2. To apply Lemma 3.2.2 it remains to show that for α a loop on Σ, the set $A = \{f \in \mathrm{MCG}(\Sigma) \mid d(\alpha, f(\alpha)) \leq 1\}$ is coarsely bounded. Fix such an α, and refer to the components of $\Sigma \setminus \alpha$ as α_+ and α_-.

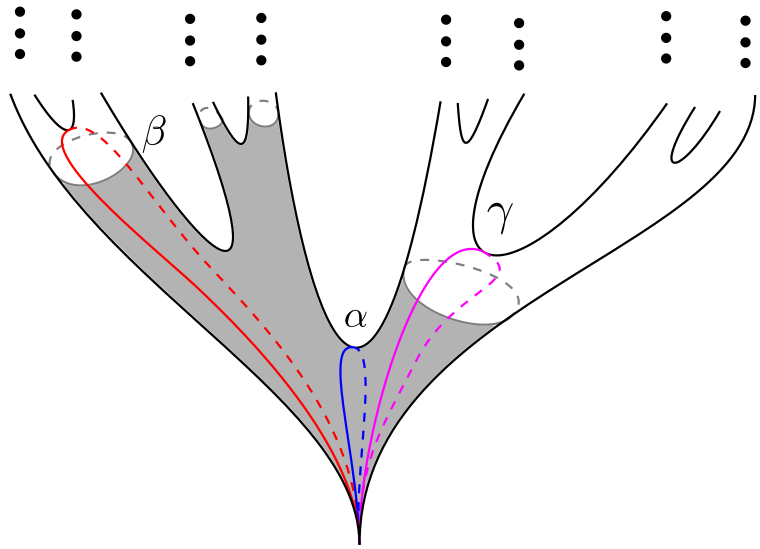

Figure 3.8: A finite-type subsurface (shaded), with the loop α in blue, the loop β in red, and the loop γ in magenta.

We will of course be using Lemma 3.2.1. Fix an identity neighborhood V in $\mathrm{MCG}(\Sigma)$, and let r and h be as in Lemmas 3.6.3 and 3.6.4. Let $F = \{r^{-1}, h_+, h_-, h_+^{-1}, h_-^{-1}\}$. We will show that $A \subseteq (FV)^8$.

Fix $f \in A$. First consider the case where $d(\alpha, f(\alpha)) = 0$. After possibly replacing f with rf, we may assume f restricts to the identity on α, and so it decomposes as $f = f_- f_+$, where $f_- \in V_{\alpha_-}$ and $f_+ \in V_{\alpha_+}$. Then $f = f_- f_+ \in V_{\alpha_-} V_{\alpha_+} \subseteq h_+^{-1} V h_+ h_-^{-1} V h_- \subseteq (FV)^4$. Since we may have replaced f with rf, this gives $f \in (FV)^5$ in general when $d(\alpha, f(\alpha)) = 0$.

Now suppose $d(\alpha, f(\alpha)) = 1$. That means α and $f(\alpha)$ are disjoint. Without loss of generality we assume that $f(\alpha) \subseteq \alpha_+$; if not then we need merely replace h_+ with h_- below. By Lemma 3.6.2 there is some $g \in V_{\alpha_-}$ such that $g(\alpha) = \alpha$ and

$g(h_+(\alpha)) = f(\alpha)$. Let $f_0 = h_+^{-1} g^{-1} f$. By construction $f_0(\alpha) = \alpha$ so by the previous paragraph $f_0 \in (FV)^5$. Then $f = gh_+ f_0 \in V_{\alpha_-} h_+(FV)^5 \subseteq h_+^{-1} V h_+ h_+ (FV)^5 \subseteq (FV)^8$.

Thus $A \subseteq (FV)^8$, so A is coarsely bounded, and then by Lemma 3.2.2 the action of $\mathrm{MCG}(\Sigma)$ on $\mathcal{L}(\Sigma)$ induces a quasi-isometry. □

3.6.1 Some consequences of this quasi-isometry

Theorem 3.6.5 has some interesting immediate consequences. The first is hyperbolicity; as mentioned in the introduction, $\mathcal{L}(\Sigma)$ is known to be δ-hyperbolic.

Corollary 3.6.6. *Let $\Sigma = \mathbb{R}^2 \setminus C$. Then* $\mathrm{MCG}(\Sigma)$ *is non-elementary δ-hyperbolic.*

Proof. The mapping class group is quasi-isometric to the loop graph, which was shown by Bavard [Bav16] to be non-elementary δ-hyperbolic. □

Since the translatable surfaces are known to have non-hyperbolic mapping class groups, this proves that the mapping class groups are not quasi-isometric.

Corollary 3.6.7. *The mapping class group of $\mathbb{R}^2 \setminus C$ is not quasi-isometric to that of any translatable surface.*

Proof. The translatable curve graph is never non-elementary hyperbolic by the results of Horbez, Qing, and Rafi [HQR20], and thus neither is the mapping class group of any translatable surface. Thus by Corollary 3.6.6 the mapping class group of a translatable surface is not quasi-isometric to that of $\mathbb{R}^2 \setminus C$. □

For the final interesting consequence, we introduce some concepts from the world of locally compact groups. A generating set S for a group G can be thought of as a homeomorphism $\varphi\colon F_S \to G$ from the free group on the set S to G. A collection of words $R \subseteq F_S$ that normally generates the kernel of this map is called a *set of relators* and we often write G as a *group presentation* $G = \langle S \mid R \rangle$. When the sets S and R are both finite, we say the group G is *finitely presented*. Cornulier and de la Harpe [CdlH16] introduce the following generalization of this notion.

Definition 3.6.8. A group presentation $G = \langle S \mid R \rangle$ is a *bounded presentation* if the words in R have bounded length. In this case we say G is *boundedly presented* over the set S.

Note that a finite presentation is simply a bounded presentation over a finite generating set. Cornulier and de la Harpe call a group *compactly presented* if it has a bounded presentation over a compact generating set, and by analogy we might call a group *coarse-boundedly presented* if it has a bounded presentation over a coarsely bounded generating set. Crucially, Cornulier and de la Harpe show the following close relationship between bounded presentations word metrics.

Fact 3.6.9 (Proposition 7.B.1 of [CdlH16])**.** *Let G be a group endowed with a generating set S. Then G is boundedly presented over S if and only if the Rips complex* $\mathrm{Rips}_c(G, d_S)$ *is simply connected for some c.*

It follows directly that the mapping class group of the plane minus a Cantor set has a coarsely bounded presentation.

Corollary 3.6.10. *Let* $\Sigma = \mathbb{R}^2 \setminus C$ *be the plane minus a Cantor set. Then* $\mathrm{MCG}(\Sigma)$ *has a coarsely bounded presentation.*

Proof. By Corollary 3.6.6, the mapping class group $\mathrm{MCG}(\Sigma)$ is δ-hyperbolic with respect to (any) coarsely bounded generating set S. Then for high enough c, the Rips complex $\mathrm{Rips}_c(G, d_S)$ is contractible, and so by Fact 3.6.9 $\mathrm{MCG}(\Sigma)$ has a bounded presentation over S. □

Bibliography

[AFP17] Javier Aramayona, Ariadna Fossas, and Hugo Parlier. Arc and curve graphs for infinite-type surfaces. *Proc. Amer. Math. Soc.*, 145(11):4995–5006, 2017.

[APV20] Tarik Aougab, Priyam Patel, and Nicholas G. Vlamis. Isometry groups of infinite-genus hyperbolic surfaces. arXiv, 2020, 2007.01982.

[Bav16] Juliette Bavard. Hyperbolicité du graphe des rayons et quasi-morphismes sur un gros groupe modulaire. *Geom. Topol.*, 20(1):491–535, 2016.

[BDR18] Juliette Bavard, Spencer Dowdall, and Kasra Rafi. Isomorphisms Between Big Mapping Class Groups. *International Mathematics Research Notices*, 05 2018, http://oup.prod.sis.lan/imrn/advance-article-pdf/doi/10.1093/imrn/rny093/24860689/rny093.pdf.

[Bra21] H. R. Brahana. Systems of circuits on two-dimensional manifolds. *Ann. of Math. (2)*, 23(2):144–168, 1921.

[BW18] Juliette Bavard and Alden Walker. The Gromov boundary of the ray graph. *Trans. Amer. Math. Soc.*, 370(11):7647–7678, 2018.

[CdlH16] Yves Cornulier and Pierre de la Harpe. *Metric geometry of locally compact groups*, volume 25 of *EMS Tracts in Mathematics*. European Mathematical Society (EMS), Zürich, 2016. Winner of the 2016 EMS Monograph Award.

[Deh38] M. Dehn. Die Gruppe der Abbildungsklassen (Das arithmetische Feld auf Flächen). *Acta Math.*, 69(1):135–206, 1938. Translated in Dehn, *Papers on Group Theory and Topology*, Springer-Verlag 1987.

[DFV18] Matthew Gentry Durham, Federica Fanoni, and Nicholas G. Vlamis. Graphs of curves on infinite-type surfaces with mapping class group actions. *Ann. Inst. Fourier (Grenoble)*, 68(6):2581–2612, 2018.

[Eps66] D. B. A. Epstein. Curves on 2-manifolds and isotopies. *Acta Math.*, 115:83–107, 1966.

[FGM20] Federica Fanoni, Tyrone Ghaswala, and Alan McLeay. Homeomorphic subsurfaces and the omnipresent arcs. arXiv, 2020, 2003.04750.

[FW99] George K. Francis and Jeffrey R. Weeks. Conway's ZIP proof. *Amer. Math. Monthly*, 106(5):393–399, 1999.

[GX13] Jean Gallier and Dianna Xu. *A guide to the classification theorem for compact surfaces*, volume 9 of *Geometry and Computing*. Springer, Heidelberg, 2013.

[Hat91] Allen Hatcher. On triangulations of surfaces. *Topology Appl.*, 40(2):189–194, 1991.

[HHMV18] Jesús Hernández Hernández, Israel Morales, and Ferrán Valdez. Isomorphisms between curve graphs of infinite-type surfaces are geometric. *Rocky Mountain J. Math.*, 48(6):1887–1904, 2018.

[HHMV19] Jesús Hernández Hernández, Israel Morales, and Ferrán Valdez. The Alexander method for infinite-type surfaces, 11 2019.

[HQR20] Camille Horbez, Yulan Qing, and Kasra Rafi. Big mapping class groups with hyperbolic actions: classification and applications. arXiv, 2020, 2005.00428.

[IM10] Elmas Irmak and John D. McCarthy. Injective simplicial maps of the arc complex. *Turkish J. Math.*, 34(3):339–354, 2010.

[Iva97] Nikolai V. Ivanov. Automorphism of complexes of curves and of Teichmüller spaces. *Internat. Math. Res. Notices*, (14):651–666, 1997.

[JDD17] Francis Jennings, Richard S Dunn, and Mary Maples Dunn. 12. brother miquon: Good lord! In *The World of William Penn*. University of Pennsylvania Press,, Philadelphia :, 2017.

[Ker23] Béla Kerékjártó. *Vorlesungen über Topologie*, volume I. Springer, 1923.

[Kor99] Mustafa Korkmaz. Automorphisms of complexes of curves on punctured spheres and on punctured tori. *Topology Appl.*, 95(2):85–111, 1999.

[Luo00] Feng Luo. Automorphisms of the complex of curves. *Topology*, 39(2):283–298, 2000.

[Mil68] J. Milnor. A note on curvature and fundamental group. *J. Differential Geometry*, 2:1–7, 1968.

[MM00] H. A. Masur and Y. N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.

[MR19] Kathryn Mann and Kasra Rafi. Large scale geometry of big mapping class groups. arXiv, 2019, 1912.10914.

[Mö61] A.F. Möbius. Zur Theorie der Polyëder und der Elementarverwandtschaft. In *Oeuvres Complètes*, volume 2, pages 519–559. 1861.

[Ras20] Alexander J. Rasmussen. Uniform hyperbolicity of the graphs of nonseparating curves via bicorn curves. *Proc. Amer. Math. Soc.*, 148(6):2345–2357, 2020.

[Ric63] Ian Richards. On the classification of noncompact surfaces. *Trans. Amer. Math. Soc.*, 106:259–269, 1963.

[Ros18] Christian Rosendal. Coarse geometry of topological groups. `http://homepages.math.uic.edu/~rosendal/PapersWebsite/Coarse-Geometry-Book23.pdf`, 2018.

[SC20] Anschel Schaffer-Cohen. Graphs of curves and arcs quasi-isometric to big mapping class groups. arXiv, 2020, 2006.14760.

[SC22] Anschel Schaffer-Cohen. Automorphisms of the loop and arc graph of an infinite-type surface. *Proc. Amer. Math. Soc.*, 2022. to appear.

[Vla19] Nicholas G. Vlamis. Notes on the topology of mapping class groups. `http://qcpages.qc.cuny.edu/~nvlamis/Papers/AIM_Notes.pdf`, 2019.

[Š55] A. S. Švarc. A volume invariant of coverings. *Dokl. Akad. Nauk SSSR (N.S.)*, 105:32–34, 1955.