COVERS AND THE CURVE COMPLEX

KASRA RAFI AND SAUL SCHLEIMER

Abstract. We propose a program of studying the coarse geometry of combinatorial moduli spaces of surfaces by classifying the quasi-isometric embeddings between them. We provide the first non-trivial examples of quasi-isometric embeddings between curve complexes. These are induced either via orbifold coverings or by puncturing a closed surface. As a corollary, we give new quasi-isometric embeddings between mapping class groups.

1. Introduction

The coarse structure of the complex of curves was first studied in [MM99]. It is central in low-dimensional topology, shedding light on the algebra of the mapping class group, the global topology of Teichmüller space and the fine structure of hyperbolic three-manifolds. It is also closely related to the geometric study of other moduli spaces of a surface: for example the pants complex, the Hatcher-Thurston complex and the disk complex.

Little is known about the subspace structure of these combinatorial moduli spaces. Thus, we propose the following:

Problem 1.1. Classify quasi-isometric embeddings between combinatorial moduli spaces of surfaces.

The most important of these moduli spaces are the complex of curves and the mapping class group itself. In this paper we produce quasi-isometric embeddings of lower into higher complexity curve complexes and similarly for the mapping class group. These arise from two topological operations on surfaces: covering and puncturing. The case of taking subsurfaces is studied elsewhere. (See below.)

Coverings. Suppose that S is a compact connected orientable orbifold, of dimension two, with non-positive orbifold Euler characteristic. Let S° denote the surface with boundary obtained by removing an open

Date: July 29, 2008.

neighborhood of the orbifold points from S. Define $\mathcal{C}(S)$ to be the curve complex of S° (see Definition 2.1). We also define the complexity

$$\xi(S) = 3\,\mathrm{genus}(S^\circ) + |\partial S^\circ| - 3.$$

Let $P\colon \Sigma \to S$ be an orbifold covering map. The covering P defines a relation $\Pi\colon \mathcal{C}(S) \to \mathcal{C}(\Sigma)$; the curve $b \in \mathcal{C}(S)$ is related to $\beta \in \mathcal{C}(\Sigma)$ if $P(\beta) = b$. That this relation is well-defined is proved in Lemma 4.1.

Theorem 10.1. *The covering relation $\Pi\colon \mathcal{C}(S) \to \mathcal{C}(\Sigma)$ is a Q–quasi-isometric embedding. The constant Q depends only on $\xi(S)$ and the degree of the covering map P.*

Theorem 10.1 is surprising in light the fact that the commonly discussed subspaces of the curve complex, such as the complex of separating curves, the disk complex of a handlebody and so on, are *not* quasi-isometrically embedded. We remark that the orbifold covering map cannot be replaced by a branched cover. The orbifold structure keeps track of which boundary components of the cover of S° must be capped off to obtain Σ°. Also, geometric structure lifts via orbifold covering; this is used, in an essential way, in the proof that the relation Π is well-defined.

Let $\mathcal{MCG}(S)$ be the orbifold mapping class group. We prove:

Theorem 11.1. *The covering P induces a quasi-isometric embedding*

$$\Pi_*\colon \mathcal{MCG}(S) \to \mathcal{MCG}(\Sigma).$$

When the cover is regular, a stronger statement holds:

Theorem 11.10. *Suppose that $\Delta \subset \mathcal{MCG}(\Sigma)$ is a finite subgroup. Then the normalizer of Δ in undistorted in $\mathcal{MCG}(\Sigma)$.*

Note that many algebraically defined subgroups of the mapping class group, such as the Torelli group, are distorted [BFP].

Puncturing. Suppose that S is a closed orientable surface of genus $g \geq 2$ and Σ is the surface of genus g with one puncture. The following theorem is inspired by Harer's paper [Har86, Lemma 3.6].

Theorem 3.1. *Lifting geodesics gives an isometric embedding of $\mathcal{C}(S)$ into $\mathcal{C}(\Sigma)$.*

There are an uncountable number of such isometric embeddings. A similar construction induces an uncountable family of quasi-isometric embeddings of the mapping class group of S into that of Σ. However, this was previously obtained by Mosher via a different technique [Mos96, quasi-isometric section lemma]. We therefore omit our construction for the mapping class group.

Subsurfaces. For completeness, we briefly mention another topological construction. Suppose that Σ is a compact orientable surface, $S \subset \Sigma$ is a cleanly embedded subsurface, and $\Sigma \smallsetminus S$ has no annular components. The inclusion $S \to \Sigma$ induces an obvious, but important, simplicial injection of curve complexes. This injection is far from being a quasi-isometric embedding; the image has diameter two. However the inclusion does induce a quasi-isometric embedding of mapping class groups. That is, these subgroups are undistorted. This follows directly from the summation formula of Masur and Minsky (see [MM00], Theorems 7.1, 6.10, and 6.12) and was independently obtained by Hamenstädt [Hama, Theorem B, Corollary 4.6].

Quasi-isometry group. A special, but quite deep, instance of Problem 1.1 is the computation of the quasi-isometry group of a combinatorial moduli space. This has recently been obtained for the mapping class group by Behrstock, Kleiner, Minsky, and Mosher [BKMM] and also by Hamenstädt [Hamb]. They show that the quasi-isometry group is virtually equal to the isometry group; that is, the mapping class group is rigid.

Using the rigidity of the mapping class group, the structure of the boundary of the curve complex and an understanding of cobounded laminations we show in [RS08] that the quasi-isometry group of the curve complex is again the mapping class goup.

Outline of the proof of Theorem 10.1. Suppose that $P\colon \Sigma \to S$ is a covering map. In this outline we assume that $\xi(S) > 1$. We deal with special cases in the body of the proof.

To prove that $\Pi\colon \mathcal{C}(S) \to \mathcal{C}(\Sigma)$ is a quasi-isometric embedding we must show, for $a, b \in \mathcal{C}(S)$ and lifts $\alpha, \beta \in \mathcal{C}(\Sigma)$, that $d_S(a, b)$ is comparable to $d_\Sigma(\alpha, \beta)$. The inequality $d_S(a, b) \geq d_\Sigma(\alpha, \beta)$ is clear; the relation Π is simplicial when $\xi(S) > 1$.

The content of the paper lies in obtaining the other direction. There are two steps: we first give a new estimate of distance in the complex of curves (Theorem 8.1). We then analyze of the behavior of our estimate under lifting.

In more detail: choose $x, y \in \mathcal{T}(S)$, the Teichmüller space of S, so that a is has bounded length in x and the same holds for b in y. Let G be the Teichmüller geodesic connecting x and y. The part of G lying in the thick part of Teichmüller space contributes linearly to the distance in $\mathcal{C}(S)$ between a and b (Lemma 6.4).

Next, we introduce $(\mathsf{T}_0, \mathsf{T}_1)$–*antichains* in the poset of subsurfaces of S (Section 7). The size of the antichain linearly estimates the number of vertices appearing in a $\mathcal{C}(S)$–geodesic while G travels through the

thin part of Teichmüller space. The sum of the estimates in the thick and thin parts is then comparable to the distance in the curve complex (Theorem 8.1).

Let Γ be the lift of G, which is again a Teichmüller geodesic with identical parametrization. This follows from the well-known fact that coverings induce isometric embeddings of the associated Teichmüller spaces (Section 6). The curves α and β have bounded lengths at the endpoints of Γ. We now estimate $d_\Sigma(\alpha, \beta)$ as above. When G is in the thick part, the same holds for Γ. Thus the thick part of Γ contributes at least as much to $d_\Sigma(\alpha, \beta)$ as the thick part of G contributes to $d_S(a, b)$.

We next prove that the lift of an antichain is again an antichain, with perhaps weaker thresholds. A key fact here is Lemma 9.3, stating that any subsurface Ω of Σ where $d_\Omega(\alpha, \beta)$ is large is symmetric. In fact, it is a lift of a subsurface Z where $d_Z(a, b)$ is large. Therefore, the estimate for $d_S(a, b)$ and $d_\Sigma(\alpha, \beta)$ are comparable (Theorem 10.1).

Our use of Teichmüller geodesics appears unavoidable: for example, Masur-Minsky heirarchies do not *a priori* have good properties vis-à-vis covers. The main geodesic does lift to a quasi-geodesic, but this only becomes clear *a posteriori*. The antichains we use are, in fact, a subset of the domains mentioned in a hierarchy. However, the antichain chooses the *correct* subset.

Acknowledgments. We thank the referee for useful comments on an early version of this paper.

2. Background

Suppose that Σ is a compact orientable orbifold, of dimension two, with non-positive orbifold Euler characteristic. For definitions and discussion of orbifolds we refer the reader to Scott's excellent article [Sco83]. Recall that Σ° is the surface obtained by removing an open neighborhood of the orbifold points from Σ. In many respects there is no difference between Σ and Σ°; we will use whichever is convenient and remark on the few subtle points as they arise.

A simple closed curve $\alpha \subset \Sigma$, avoiding the orbifold points, is *inessential* if α bounds a disk in Σ containing one or zero orbifold points. The curve α is *peripheral* if α is isotopic to a boundary component. Note that isotopies of curves are not allowed to cross orbifold points.

Definition 2.1. When $\xi(\Sigma) > 1$ the *complex of curves* $\mathcal{C}(\Sigma^\circ)$ has as its vertices isotopy classes of essential, non-peripheral curves. A collection of $k+1$ distinct vertices spans a k–simplex if every pair of vertices has disjoint representatives.

There is a different definition when $\xi(\Sigma) \leq 1$. When Σ° is a torus, once-holed torus or a four-holed sphere the curve complex of Σ° is the well-known *Farey graph*; since all curves intersect, edges are instead placed between curves that intersect exactly once or exactly twice, respectively. The curve complex of the three-holed sphere is empty.

If Σ° is an annulus, then vertices of $\mathcal{C}(\Sigma^\circ)$ are essential arcs in Σ°, considered up to isotopy relative to their boundary. Edges are placed between vertices with representatives having disjoint interiors. The result will be quasi-isometric to $\mathbb{Z}$. The assumption on the Euler characteristic of Σ prevents Σ° from being a disk or a sphere.

To obtain a metric, give all edges of $\mathcal{C}(\Sigma)$ length one and denote distance between vertices by $d_\Sigma(\cdot, \cdot)$. It will be enough to study only the one-skeleton of $\mathcal{C}(\Sigma)$, for which we use the same notation. This is because the one-skeleton and the entire complex are *quasi-isometric*.

We now turn to convenient piece of notation: if A, B, c are non-negative real numbers with $\mathsf{c} > 0$ and if $A \leq \mathsf{c}B + \mathsf{c}$, then we write $A \prec_\mathsf{c} B$. If $A \prec_\mathsf{c} B$ and $B \prec_\mathsf{c} A$, then we write $A \asymp_\mathsf{c} B$. Suppose $\mathcal{X}$ and $\mathcal{Y}$ are metric spaces and $f\colon \mathcal{X} \to \mathcal{Y}$ is a relation. We say that f is a c–*quasi-isometric embedding* if for all $x, x' \in \mathcal{X}$ and for all $y \in f(x), y' \in f(x')$ we have $d_\mathcal{X}(x, x') \asymp_\mathsf{c} d_\mathcal{Y}(y, y')$. We say that f is a c–*quasi-isometry* if additionally a c–neighborhood of $f(\mathcal{X})$ equals $\mathcal{Y}$.

3. Puncturing

Before dealing with orbifold covers we discuss puncturing closed surfaces. Let S be a closed surface of genus $g \geq 2$ and Σ be the surface of genus g with one puncture.

Theorem 3.1. *$\mathcal{C}(S)$ embeds isometrically into $\mathcal{C}(\Sigma)$.*

As we shall see, there are uncountably many such embeddings.

Proof of Theorem 3.1. Pick a hyperbolic metric on S. By the Baire category theorem, the union of geodesic representatives of simple closed curves does not cover S. (In fact, this union has Hausdorff dimension one. See Birman and Series [BS85].) Let $*$ be a point in the complement and identify Σ with $S \smallsetminus \{*\}$. For a vertex of $\mathcal{C}(S)$, consider its geodesic representative in the given hyperbolic metric. This is an essential curve in $S \smallsetminus \{*\}$. So it is identified with an essential curve in Σ and gives a vertex of $\mathcal{C}(\Sigma)$. This defines an embedding $\Pi\colon \mathcal{C}(S) \to \mathcal{C}(\Sigma)$ that depends on the choice of metric, point and identification. Let $P\colon \mathcal{C}(\Sigma) \to \mathcal{C}(S)$ be the map obtained by filling the point $*$. Note that $P \circ \Pi$ is the identity map.

We observe, for $a, b \in \mathcal{C}(S)$, $\alpha = \Pi(a)$ and $\beta = \Pi(b)$, that

$$d_S(a, b) = d_\Sigma(\alpha, \beta).$$

This is because P and Π send disjoint curves to disjoint curves. Therefore, if $L \subset \mathcal{C}(S)$ is a geodesic connecting a and b, then $\Pi(L)$ is a path in $\mathcal{C}(\Sigma)$ of the same length connecting α to β. Conversely, if $\Lambda \subset \mathcal{C}(\Sigma)$ is a geodesic connecting α to β, then $P(\Lambda)$ is a path in $\mathcal{C}(S)$ of the same length connecting a to b. □

4. Covering

We now turn to the main topic of the paper. Suppose that S is a compact connected orientable orbifold, of dimension two, with non-positive orbifold Euler characteristic. Let $P\colon \Sigma \to S$ be an orbifold covering map. At a first reading it is simplest to assume that Σ and S are both surfaces.

Recall that the covering P defines a relation $\Pi\colon \mathcal{C}(S) \to \mathcal{C}(\Sigma)$: a curve $b \in \mathcal{C}(S)$ is related to $\beta \in \mathcal{C}(\Sigma)$ if $P(\beta) = b$.

Lemma 4.1. *The covering relation* Π *is well-defined.*

Proof. We will show that if a is an essential non-peripheral curve then every component of $P^{-1}(a)$ is essential and non-peripheral. Since S has non-positive orbifold Euler characteristic, choose a Euclidean or a hyperbolic metric on S with totally geodesic boundary. Replace a by its geodesic representative, a^*. Since a is simple, a^* misses all cone points of order greater than two. In fact, the only way a^* meets a cone point is when a bounds a disk with exactly two orbifold points of order two; here a^* is a geodesic arc connecting these two points. In any case, the lift of a is an essential non-peripheral simple closed curve that is homotopic to the lift of a^*. The conclusion follows. □

5. Subsurface projection

Suppose that Σ is a compact connected orientable orbifold. A strict suborbifold Ψ is *cleanly embedded* if every component of $\partial\Psi$ is either a boundary component of Σ or is an essential non-peripheral curve in Σ. All suborbifolds considered will be cleanly embedded.

From [MM99], recall the definition of the *subsurface projection* relation

$$\pi_\Psi\colon \mathcal{C}(\Sigma) \to \mathcal{C}(\Psi),$$

When $\mathcal{C}(\Psi)$ is empty, the projection is not defined. Thus assume that $\mathcal{C}(\Psi)$ is non-empty.

Suppose that Σ has negative orbifold Euler characteristic. Choose a complete finite volume hyperbolic metric on the interior of Σ. Let Σ' be the Gromov compactification of the cover of Σ corresponding to the inclusion $\pi_1^{\text{orb}}(\Psi) \to \pi_1^{\text{orb}}(\Sigma)$ (defined up to conjugation). Thus Σ' is homeomorphic to Ψ; this gives a canonical identification of $\mathcal{C}(\Psi)$ with $\mathcal{C}(\Sigma')$. For any $\alpha \in \mathcal{C}(\Sigma)$ let α' be the closure of the preimage of α in Σ'. If every component of α' is properly isotopic into the boundary then α is not related to any vertex of $\mathcal{C}(\Psi)$; in this case we write $\pi_\Psi(\alpha) = \emptyset$. Otherwise, let α'' be a component of α' that is not properly isotopic into the boundary. Let N be a closed regular neighborhood of $\alpha'' \cup \partial\Sigma'$. As $\mathcal{C}(\Psi) \neq \emptyset$, the suborbifold Ψ° is not a thrice-holed sphere; so there is a boundary component α''' of N which is essential and non-peripheral. We then write $\pi_\Psi(\alpha) = \alpha'''$.

If Ψ is an annulus the projection map is defined as above, n omitting the final steps involving the regular neighborhood N.

It remains to deal with the case where $\chi(\Sigma) = 0$ and Ψ is a cleanly embedded annulus. Here Σ is either a flat torus or a *pillowcase*. The torus is a double cover of the pillowcase and the cover induces an isomorphisms of curve complexes. In this case $\mathcal{C}(\Psi)$ is quasi-isometric to $\mathbb{Z}$. We take the projection of α to Ψ to be the integer part of the slope of α with respect to Ψ.

Generally, the curve $\alpha \in \mathcal{C}(\Sigma)$ *cuts* the suborbifold Ψ if $\pi_\Psi(\alpha) \neq \emptyset$. Otherwise, α *misses* Ψ. Suppose now that $\alpha, \beta \in \mathcal{C}(\Sigma)$ both cut Ψ. Define the *projection distance* to be

$$d_\Psi(\alpha, \beta) = d_\Psi(\pi_\Psi(\alpha), \pi_\Psi(\beta)).$$

The Bounded Geodesic Image Theorem states:

Theorem 5.1 (Masur-Minsky [MM00])**.** *Fix a surface Σ. There is a constant $\mathsf{M} = \mathsf{M}(\Sigma)$ with the following property. Suppose that $\alpha, \beta \in \mathcal{C}(\Sigma)$ are vertices, $\Lambda \subset \mathcal{C}(\Sigma)$ is a geodesic connecting α to β and $\Omega \subsetneq \Sigma$ is a subsurface. If $d_\Omega(\alpha, \beta) \geq \mathsf{M}$ then there is a vertex of Λ which misses Ω.* □

6. Teichmüller space

For this section, we take Σ to be a surface. Let $\mathcal{T}(\Sigma)$ denote the Teichmüller space of Σ: the space of complete hyperbolic metrics on the interior of Σ, up to isotopy. For background, see [Ber60, Gar87].

There is a uniform upper bound on the length of the shortest closed curve in any hyperbolic metric on Σ. For any metric σ on Σ, a curve γ has *bounded length* in σ if the length of γ in σ is less than this constant. Let $\mathsf{e}_0 > 0$ be a constant such that, for curves γ and δ, if γ

has bounded length in σ and δ has a length less than e_0 then γ and δ have intersection number zero.

Suppose that α and β are vertices of $\mathcal{C}(\Sigma)$. Fix metrics σ and τ in $\mathcal{T}(\Sigma)$ so that α and β have bounded length at σ and τ respectively. Let $\Gamma\colon [t_\sigma, t_\tau] \to \mathcal{T}(S)$ be a geodesic connecting σ to τ. For any curve γ let $l_t(\gamma)$ be the length of its geodesic representative in the hyperbolic metric $\Gamma(t)$. The following theorems are consequences of Theorem 6.2 and Lemma 7.3 in [Raf05].

Theorem 6.1 ([Raf05])**.** *For e_0 as above there exists a threshold $\mathsf{T}_{\min}$ such that, for a strict subsurface Ω of Σ, if $d_\Omega(\alpha, \beta) \geq \mathsf{T}_{\min}$ then there is a time t_Ω so that the length of each boundary component of Ω in $\Gamma(t_\Omega)$ is less than e_0.* □

Theorem 6.2 ([Raf05])**.** *For every threshold T_1, there is a constant e_1 such that if $l_t(\gamma) \leq \mathsf{e}_1$, for some curve γ and for some time t, then there exists a subsurface Ψ disjoint from γ such that $d_\Psi(\alpha, \beta) \geq \mathsf{T}_1$.* □

The *shadow* of the Teichmüller geodesic Γ inside of $\mathcal{C}(\Sigma)$ is the set of curves γ so that γ has bounded length in $\Gamma(t)$ for some $t \in [t_\sigma, t_\tau]$. The following is a consequence of the fact that the shadow is an unparameterized quasi-geodesic. (See Theorem 2.6 and then apply Theorem 2.3 in [MM99].)

Theorem 6.3 ([MM99])**.** *The shadow of a Teichmüller geodesic inside of $\mathcal{C}(\Sigma)$ does not backtrack and so satisfies the reverse triangle inequality. That is, there exists a backtracking constant $\mathsf{B} = \mathsf{B}(\Sigma)$ such that if $t_\sigma \leq t_0 \leq t_1 \leq t_2 \leq t_\tau$ and if γ_i has bounded length in $\Gamma(t_i)$, $i = 0, 1, 2$ then*

$$d_\Sigma(\gamma_0, \gamma_2) \geq d_\Sigma(\gamma_0, \gamma_1) + d_\Sigma(\gamma_1, \gamma_2) - \mathsf{B}.$$

We say that $\Gamma(t)$ is e*–thick* if the shortest closed geodesic γ in $\Gamma(t)$ has a length of at least e.

Lemma 6.4. *For every $\mathsf{e} > 0$ there is a progress constant $\mathsf{P} > 0$ so that if $t_\sigma \leq t_0 \leq t_1 \leq t_\tau$, if $\Gamma(t)$ is e–thick at every time $t \in [t_0, t_1]$, and if γ_i has bounded length in $\Gamma(t_i)$ ($i = 0, 1$) then*

$$d_\Sigma(\gamma_0, \gamma_1) \asymp_{\mathsf{P}} t_1 - t_0.$$

Proof. As above, using Theorem 6.2 and Lemma 7.3 in [Raf05] and the fact that $\Gamma(t)$ is e–thick at every time $t \in [t_0, t_1]$, we can conclude that $d_\Omega(\gamma_0, \gamma_1)$ is uniformly bounded for any strict subsurface of Ω of Σ. The lemma is then a consequence of Theorem 1.1 and Remark 5.5 in [Raf07]. (Referring to the statement and notation of [Raf07, Theorem 1.1]: Extend γ_i to a short marking μ_i. Take k large enough such

that the only non-zero term in the right hand side of [Raf07, Equation (1)] is $d_\Sigma(\mu_0, \mu_1)$.) □

In general the geodesic Γ may stray into the thin part of $\mathcal{T}(S)$. We take $\Gamma^{\geq \mathsf{e}}$ to be the set of times in the domain of Γ which are e–thick. Notice that $\Gamma^{\geq \mathsf{e}}$ is a union of closed intervals. Let $\Gamma(\mathsf{e}, \mathsf{L})$ be the union of intervals of $\Gamma^{\geq \mathsf{e}}$ which have length at least L. We use $|\Gamma(\mathsf{e}, \mathsf{L})|$ to denote the sum of the lengths of the components of $\Gamma(\mathsf{e}, \mathsf{L})$.

Lemma 6.5. *For every* e *there exists* L_0 *such that if* $\mathsf{L} \geq \mathsf{L}_0$*, then*

$$d_\Sigma(\alpha, \beta) \geq |\Gamma(\mathsf{e}, \mathsf{L})|/2\mathsf{P}.$$

Proof. Pick L_0 large enough so that, for $\mathsf{L} \geq \mathsf{L}_0$,

$$(\mathsf{L}/2\mathsf{P}) \geq \mathsf{P} + 2\mathsf{B}.$$

Realize $\Gamma(\mathsf{e}, \mathsf{L})$ as the union of intervals $[t_i, s_i]$, $i = 1, \ldots, m$. Let γ_i be a curve of bounded length in $\Gamma(t_i)$ and δ_i be a curve of bounded length in $\Gamma(s_i)$.

By Theorem 6.3 we have

$$d_\Sigma(\alpha, \beta) \geq \left(\sum_i d_\Sigma(\gamma_i, \delta_i) \right) - 2m\mathsf{B}.$$

From Lemma 6.4 we deduce

$$d_\Sigma(\alpha, \beta) \geq \left(\sum_i \frac{1}{\mathsf{P}}(s_i - t_i) - \mathsf{P} \right) - 2m\mathsf{B}.$$

Rearranging, we find

$$d_\Sigma(\alpha, \beta) \geq \frac{1}{\mathsf{P}} |\Gamma(\mathsf{e}, \mathsf{L})| - m(\mathsf{P} + 2\mathsf{B}).$$

Thus, as desired:

$$d_\Sigma(\alpha, \beta) \geq \frac{1}{2\mathsf{P}} |\Gamma(\mathsf{e}, \mathsf{L})|.$$ □

7. Antichains

Consider two curves $\alpha, \beta \in \mathcal{C}(\Sigma)$. As discussed in the introduction, we would like to estimate the length of the geodesic $[\alpha, \beta]$ in $\mathcal{C}(\Sigma)$ corresponding to the times when the Teichmüller geodesic $\Gamma = [\sigma, \tau]$ is in the thin part of $\mathcal{T}(\Sigma)$. At such time, Theorem 6.2 gives a suborbifolds Ω where $d_\Omega(\alpha, \beta)$ is large. However, the number of these suborbifolds is not a good estimate for the distance in the complex of curves; many subsurfaces with high projection distance may be disjoint from a single curve in the geodesic $[\alpha, \beta]$. Nonetheless, by carefully choosing a subcollection of such suborbifolds, we can find a suitable estimate.

Fix α and β in $\mathcal{C}(\Sigma)$ and *thresholds* $\mathsf{T}_1 \geq \mathsf{T}_0 > 0$. We say that a set $\mathcal{J}$ of suborbifolds $\Omega \subsetneq \Sigma$, is a $(\mathsf{T}_0, \mathsf{T}_1)$–*antichain* for (Σ, α, β) if $\mathcal{J}$ satisfies the following properties.

- $\mathcal{J}$ is an antichain in the poset of suborbifolds ordered by inclusion: if $\Omega, \Omega' \in \mathcal{J}$ then Ω is not a strict suborbifold of Ω'.
- If $\Omega \in \mathcal{J}$ then $d_\Omega(\alpha, \beta) \geq \mathsf{T}_0$.
- For any $\Psi \subsetneq \Sigma$, if $d_\Psi(\alpha, \beta) \geq \mathsf{T}_1$ then Ψ is a suborbifold of some element of $\mathcal{J}$.

Notice that there may be many different antichains for the given data $(\Sigma, \alpha, \beta, \mathsf{T}_0, \mathsf{T}_1)$. One particularly nice example is when $\mathsf{T}_0 = \mathsf{T}_1 = \mathsf{T}$ and $\mathcal{J}$ is defined to be the maxima of the set

$$\{\Omega \subsetneq \Sigma \mid d_\Omega(\alpha, \beta) \geq \mathsf{T}\}$$

as ordered by inclusion. We call this the T–*antichain of maxima* for (Σ, α, β). By $|\mathcal{J}|$ we mean the number of elements of $\mathcal{J}$. We may now prove:

Lemma 7.1. *For every orbifold Σ and for every pair of sufficiently large thresholds $\mathsf{T}_0, \mathsf{T}_1$, there is an* accumulation constant $\mathsf{A}_\Sigma = \mathsf{A}(\Sigma, \mathsf{T}_0, \mathsf{T}_1)$ *so that if $\mathcal{J}$ is an $(\mathsf{T}_0, \mathsf{T}_1)$–antichain for (Σ, α, β) then*

$$d_\Sigma(\alpha, \beta) \geq |\mathcal{J}|/\mathsf{A}_\Sigma.$$

Proof. We proceed via induction on the complexity of Σ. In the base case, when $\mathcal{C}(\Sigma^\circ)$ is the Farey graph, $\mathcal{J}$ is the set of annuli whose core curves γ have the property that $d_\gamma(\alpha, \beta) \geq \mathsf{T}_0$. In this case, assuming $\mathsf{T}_0 > 3$, every such curve γ is a vertex of every geodesic connecting α to β (see [Min99, §4]). Therefore the lemma holds for Farey graphs with $\mathsf{A}_\Sigma = 1$.

Now assume the lemma is true for all surfaces with lower complexity than Σ. Let C be a constant so that: if $\Omega \subset \Psi \subset \Sigma$ and α', β' are the projections of α, β to Ψ then

$$|d_\Omega(\alpha, \beta) - d_\Omega(\alpha', \beta')| \leq \mathsf{C}.$$

We take thresholds T_0 and T_1 for Σ large enough so that for the $(\mathsf{T}_0, \mathsf{T}_1)$–antichain $\mathcal{J}$ we have:

- the lemma applies to any strict suborbifold Ψ with thresholds $\mathsf{T}_0 - \mathsf{C}, \mathsf{T}_1 + \mathsf{C}$ and
- $\mathsf{T}_0 \geq \mathsf{M}(\Sigma)$; thus by Theorem 5.1 for any orbifold in $\Omega \in \mathcal{J}$ and for any geodesic $\Lambda = [\alpha, \beta]$ in $\mathcal{C}(\Sigma)$ there is a curve γ in Λ so that γ misses Ω.

For $\Psi \subsetneq \Sigma$, define

$$\mathsf{A}_\Psi = \mathsf{A}(\Psi, \mathsf{T}_0 - \mathsf{C}, \mathsf{T}_1 + \mathsf{C}) \quad \text{and} \quad \mathcal{J}_\Psi = \{\Omega \in \mathcal{J} \mid \Omega \subsetneq \Psi\}.$$

Claim. Suppose that γ be a vertex in $\Lambda = [\alpha, \beta]$ and let Ψ be a component of $\Sigma \setminus \gamma$. Then

$$|\mathcal{J}_\Psi| \leq \mathsf{A}_\Psi \cdot (\mathsf{T}_1 + \mathsf{C}).$$

Proof of claim. If Ψ is a suborbifold of an element of $\mathcal{J}$ then $\mathcal{J}_\Psi$ is the empty set and the claim holds vacuously. Thus we may assume that

$$d_\Psi(\alpha, \beta) < \mathsf{T}_1.$$

Let α' and β' be the projections of α and β to Ψ. From the definition of C, $\mathcal{J}_\Psi$ is a $(\mathsf{T}_0 - \mathsf{C}, \mathsf{T}_1 + \mathsf{C})$–antichain for Ψ, α' and β'. Thus,

$$\mathsf{T}_1 > d_\Psi(\alpha, \beta) \geq d_\Psi(\alpha', \beta') - \mathsf{C} \geq |\mathcal{J}_\Psi|/\mathsf{A}_\Psi - \mathsf{C},$$

with the last inequality being the induction hypothesis. Hence,

$$\mathsf{T}_1 + \mathsf{C} \geq |\mathcal{J}_\Psi|/\mathsf{A}_\Psi. \qquad \square$$

Now consider a vertex $\gamma \in \Lambda$. Note that $\Sigma \setminus \gamma$ has at most two components, say Ψ and Ψ'. Any element of $\mathcal{J}$ not cut by γ is either a strict suborbifold of Ψ or Ψ', an annular neighborhood of γ, or Ψ or Ψ' itself. Therefore, by the above claim, the maximum number of orbifolds in $\mathcal{J}$ that are disjoint from γ is

$$(\mathsf{A}_\Psi + \mathsf{A}_{\Psi'})(\mathsf{T}_1 + \mathsf{C}) + 3.$$

Since every orbifold in $\mathcal{J}$ is disjoint from some vertex of Λ, the lemma holds for $\mathsf{A}(\Sigma, \mathsf{T}_0, \mathsf{T}_1)$ equal to

$$2 \cdot \max\{\mathsf{A}_\Psi \mid \Psi \subsetneq \Sigma\} \cdot (\mathsf{T}_1 + \mathsf{C}) + 3. \qquad \square$$

8. An estimate of distance

Again, take Σ to be a surface. In this section we provide the main estimate for $d_\Sigma(\alpha, \beta)$. Let e_0 be as before. We choose thresholds $\mathsf{T}_0 \geq \mathsf{T}_{\min}$ (see Theorem 6.1) and T_1 so that Lemma 7.1 holds. Let e_1 be the constant provided in Lemma 6.4 and let $\mathsf{e} > 0$ be any constant smaller than $\min\{\mathsf{e}_0, \mathsf{e}_1\}$. Finally, we pick L_0 such that Lemma 6.5 holds and that $\mathsf{L}_0/2\mathsf{P} > 4$. Let L be any length larger than L_0.

Theorem 8.1. *Let T_0, T_1, e and L be constants chosen as above. There is a constant $\mathsf{K} = \mathsf{K}(\Sigma, \mathsf{T}_0, \mathsf{T}_1, \mathsf{e}, \mathsf{L})$ such that for any curves α and β, any $(\mathsf{T}_0, \mathsf{T}_1)$–antichain $\mathcal{J}$ and any Teichmüller geodesic Γ, chosen as above, we have:*

$$d_\Sigma(\alpha, \beta) \asymp_{\mathsf{K}} |\mathcal{J}| + |\Gamma(\mathsf{e}, \mathsf{L})|.$$

Proof. For $\mathsf{K} \geq 2 \cdot \max(\mathsf{A}, 2\mathsf{P})$, the inequality

$$d_\Sigma(\alpha, \beta) \succ_\mathsf{K} |\mathcal{J}| + |\Gamma(\mathsf{e}, \mathsf{L})|$$

follows from Lemmas 7.1 and 6.5. It remains to show that

$$d_\Sigma(\alpha, \beta) \prec_\mathsf{K} |\mathcal{J}| + |\Gamma(\mathsf{e}, \mathsf{L})|.$$

For each $\Omega \in \mathcal{J}$ fix a time $t_\Omega \in [t_\sigma, t_\tau]$ so that all boundary components of Ω are e_0–short in $\Gamma(t_\Omega)$ (see Theorem 6.1). Let $\mathcal{E}$ be the union:

$$\Big\{ t_\Omega \;\Big|\; \Omega \in \mathcal{J},\ t_\Omega \notin \Gamma(\mathsf{e}, \mathsf{L}) \Big\} \cup \Big\{ \partial I \;\Big|\; I \text{ a component of } \Gamma(\mathsf{e}, \mathsf{L}) \Big\}.$$

We write $\mathcal{E} = \{t_0, \ldots, t_n\}$, indexed so that $t_i < t_{i+1}$.

Claim. The number of intervals in $\Gamma(\mathsf{e}, \mathsf{L})$ is at most $|\mathcal{J}| + 1$. Hence, $|\mathcal{E}| \leq 3|\mathcal{J}| + 1$.

Proof. There is at least one moment s between any two consecutive intervals $I, J \subset \Gamma(\mathsf{e}, \mathsf{L})$ when some curve γ becomes e–short (and hence e_1–short). Therefore, by Theorem 6.2, γ is disjoint from a subsurface Ψ where $d_\Psi(\alpha, \beta) \geq \mathsf{T}_1$. Since $\mathcal{J}$ is an $(\mathsf{T}_0, \mathsf{T}_1)$–antichain, Ψ is a subsurface of some element $\Omega \in \mathcal{J}$. It follows that $d_\Sigma(\gamma, \partial\Omega) \leq 2$. This defines a one-to-many relation from pairs of consecutive intervals to $\mathcal{J}$. To see the injectivity consider another such pair of consecutive intervals I' and J', a moment s' between them and a corresponding curve γ' and subsurface Ω'. Let $\Gamma' = \Gamma|_{[s,s']}$. Applying Lemma 6.5 to Γ', we find

$$d_\Sigma(\gamma, \gamma') \geq \mathsf{L}/2\mathsf{P} > 4$$

and therefore Ω is not equal to Ω'. □

Let γ_i be a curve of bounded length in $\Gamma(t_i)$.

Claim.

$$d_\Sigma(\gamma_i, \gamma_{i+1}) \leq \begin{cases} \mathsf{P}(t_{i+1} - t_i) + \mathsf{P}, & \text{if } [t_i, t_{i+1}] \subset \Gamma(\mathsf{e}, \mathsf{L}) \\ 2\mathsf{B} + \mathsf{PL} + \mathsf{P} + 2, & \text{otherwise.} \end{cases}$$

Proof. The first case follows from Lemma 6.4. So suppose that the interior of $[t_i, t_{i+1}]$ is disjoint from the interior of $\Gamma(\mathsf{e}, \mathsf{L})$.

We define sets $I_+, I_- \subset [t_i, t_{i+1}]$ as follows: A point $t \in [t_i, t_{i+1}]$ lies in I_- if

- there is a curve γ which is e–short in $\Gamma(t)$ and
- for some $\Omega \in \mathcal{J}$, so that $d_\Sigma(\partial\Omega, \gamma) \leq 2$, we have $t_\Omega \leq t_i$.

If instead $t_\Omega \geq t_{i+1}$ then we place t in I_+. Finally, we place t_i in I_- and t_{i+1} in I_+.

Notice that if $\Omega \in \mathcal{J}$ then t_Ω does not lie in the open interval (t_i, t_{i+1}). It follows that every e–thin point of $[t_i, t_{i+1}]$ lies in I_-, I_+, or both. If $t \in I_-$ and γ is the corresponding e–short curve then $d_\Sigma(\gamma_i, \gamma) \leq \mathsf{B}+2$. This is because either $t = t_i$ and so γ and γ_i are in fact disjoint, or there is a surface $\Omega \in \mathcal{J}$ as above with

$$2 \geq d_\Sigma(\partial\Omega, \gamma) \geq d_\Sigma(\gamma_i, \gamma) - \mathsf{B},$$

Similarly if $t \in I_+$ then $d_\Sigma(\gamma_{i+1}, \gamma) \leq \mathsf{B} + 2$.

If I_+ and I_- have non-empty intersection then $d_\Sigma(\gamma_i, \gamma_{i+1}) \leq 2B + 4$ by the triangle inequality.

Otherwise, there is an interval $[s, s']$ that is e–thick, has length less than L such that $s \in I_-$ and $s' \in I_+$. Let γ and γ' be the corresponding short curves in $\Gamma(s)$ and $\Gamma(s')$. Thus

$$d_\Sigma(\gamma_i, \gamma) \leq \mathsf{B} + 2 \quad \text{and} \quad d_\Sigma(\gamma', \gamma_{i+1}) \leq \mathsf{B} + 2.$$

We also know from Lemma 6.4 that

$$d_\Sigma(\gamma, \gamma') \leq \mathsf{PL} + \mathsf{P}.$$

This finishes the proof of our claim. □

It follows that

$$\begin{aligned} d_\Sigma(\alpha, \beta) &\leq d_\Sigma(\gamma_0, \gamma_1) + \ldots + d_\Sigma(\gamma_{n-1}, \gamma_n) \\ &\leq |\mathcal{E}|(2\mathsf{B} + \mathsf{PL} + \mathsf{P} + 2) + \mathsf{P}|\Gamma(\mathsf{e}, \mathsf{L})| + |\mathcal{E}|\mathsf{P} \\ &\prec_{\mathsf{K}} |\mathcal{J}| + |\Gamma(\mathsf{e}, \mathsf{L})|, \end{aligned}$$

for an appropriate choice of K. This proves the theorem. □

9. Symmetric curves and surfaces

Let $P\colon \Sigma \to S$ be an orbifold covering map.

Definition 9.1. A curve $\alpha \in \mathcal{C}(\Sigma)$ is *symmetric* if there is a curve $a \in \mathcal{C}(S)$ so that $P(\alpha) = a$. We make the same definition for a suborbifold $\Omega \subset \Sigma$ lying over a suborbifold $Z \subset S$.

As is well-known, coverings of surfaces induce isometric embeddings of the associated Teichmüller spaces. For completeness and to establish notation we include a proof below.

For the rest of this section, fix symmetric curves α and β. Pick $x, y \in \mathcal{T}(S^\circ)$ so that $a = P(\alpha)$ has bounded length in x and $b = P(\beta)$ is bounded in y. Let $G\colon [t_x, t_y] \to \mathcal{T}(S^\circ)$ be the Teichmüller geodesic connecting x to y. For every $t \in [t_x, t_y]$ let q_t be the terminal quadratic differential of the Teichmüller map from $G(t_x)$ to $G(t)$. We lift q_t to

the surface $P^{-1}(S^\circ)$, fill the punctures not corresponding to orbifold points and so obtain a parameterized family θ_t of quadratic differentials on Σ°. Notice that θ_t is indeed a quadratic differential: suppose that $p \in S$ is a orbifold point and q_t has a once-pronged singularity at p. For every regular point π in the preimage of p the differential θ_t has at least a twice-pronged singularity at π.

Uniformize the associated flat structures to obtain hyperbolic metrics on Σ°. This gives a path $\Gamma\colon [t_x, t_y] \to \mathcal{T}(\Sigma^\circ)$. The path Γ is a geodesic in $\mathcal{T}(\Sigma^\circ)$. This is because, for $t, s \in [t_x, t_y]$, the Teichmüller map from $G(t)$ to $G(s)$ has Beltrami coefficient $\mathsf{k}\,|q|/q$ where q is an integrable holomorphic quadratic differential in $G(t)$. This map lifts to a map from $\Gamma(t)$ to $\Gamma(s)$ with Beltrami coefficient $\mathsf{k}\,|\theta|/\theta$, where the quadratic differential θ is the pullback of q to $\Gamma(t)$. That is, the lift of the Teichmüller map from from $G(t)$ to $G(s)$ is the Teichmüller map from $\Gamma(t)$ to $\Gamma(s)$ with the same quasi-conformal constant. Therefore, the distance in $\mathcal{T}(S^\circ)$ between $G(t)$ and $G(s)$ equals the distance in $\mathcal{T}(\Sigma^\circ)$ between $\Gamma(t)$ and $\Gamma(s)$.

Proposition 9.2 (Proposition 3.7 [Raf07])**.** *For any* e*, there is a constant* N *such that the following holds. Assume that, for all* $t \in [r, s]$*, there is a component of* $\partial\Omega$ *whose length in* $\Gamma(t)$ *is larger than* e*. Suppose* γ *has bounded length in* $\Gamma(r)$ *and* δ *has bounded length in* $\Gamma(s)$*. Then*

$$d_\Omega(\gamma, \delta) \leq \mathsf{N}.$$

Lemma 9.3. *For* e *small enough,* N *as above and any suborbifold* $\Omega \subset \Sigma$*, if* $d_\Omega(\alpha, \beta) \geq 2\mathsf{N} + 1$*, then* Ω *is symmetric.*

Proof. Consider the first time t^- and last time t^+ that the boundary of Ω is e–short. Since every component of $\partial\Omega$ is short in $\Gamma(t^\pm)$, so is the image $P(\partial\Omega)$ in $G(t^\pm)$. Therefore, all components of the image are simple. (This is a version of the Collar Lemma. For example, see [Bus92, Theorem 4.2.2].) It follows that the boundary of Ω is symmetric. This is because choosing e small enough will ensure that curves in $P^{-1}(P(\Omega))$ have bounded length at both t^- and t^+. (The length of each is at most the degree of the covering map times e.) If any such curve γ intersects Ω we have $d_\Omega(\gamma, \alpha) \leq \mathsf{N}$ and $d_\Omega(\gamma, \beta) \leq \mathsf{N}$, contradicting the assumption $d_\Omega(\alpha, \beta) \geq 2\mathsf{N}+1$. Thus, the suborbifold Ω is symmetric. □

10. The quasi-isometric embedding

Theorem 10.1. *The covering relation* $\Pi\colon \mathcal{C}(S) \to \mathcal{C}(\Sigma)$*, corresponding to the covering map* $P\colon \Sigma \to S$*, is a* Q*–quasi-isometric embedding. The*

constant Q *depends only on* $\xi(S)$ *and the degree of the covering map* P*.*

Remark 10.2. Note that Q does not depend directly on the topology of Σ. When S is an annulus, the degree of covering is not determined by the topology of Σ. Conversely, when S is not an annulus, the topology of Σ can be bounded in terms of the topology of S and the degree of the covering.

Remark 10.3. The constant Q may go to infinity with the degree of the covering. For example, any pair of distinct curves a, b in a surface S may be made disjoint in some cover. In fact a cover of degree at most 2^{d-1}, where $d = d_S(a, b)$, will suffice [Hem01, Lemma 2.3].

Remark 10.4. When Σ is the orientation double cover of a nonorientable surface S, Theorem 10.1 is due to Masur-Schleimer [MS07].

Proof of Theorem 10.1. We first show that

$$d_\Sigma(\alpha, \beta) \prec d_S(a, b).$$

When $\xi(S) > 1$ and when S is an annulus, two vertices of $\mathcal{C}(S)$ have distance one when they have intersection number zero. But disjoint curves in S have disjoint preimages in Σ. Therefore, a path connecting a to b lifts to a path of equal length connecting α to β . This implies the desired inequality in this case. In all other cases, two curves are at distance one when they intersect once or twice, depending on S. The lifts of these curves then intersect at most $2\mathsf{d}$ times, where d is the degree of the covering. Thus, the distance between the lifts is at most $2\log_2(2\mathsf{d}) + 2$. (See [Sch, Lemma 1.21].) Therefore

$$d_\Sigma(\alpha, \beta) \leq (2\log_2(2\mathsf{d}) + 2) \cdot d_S(a, b).$$

Now we must prove the opposite inequality:

$$d_\Sigma(\alpha, \beta) \succ_{\mathsf{Q}} d_S(a, b).$$

Suppose that d is the degree of the covering. We prove the theorem by induction on the complexity of S. In the case where S is an annulus without orbifold points, the cover Σ is also an annulus and the distances in $\mathcal{C}(\Sigma)$ and $\mathcal{C}(S)$ are equal to the intersection number plus one. But, in this case,

$$i(\alpha, \beta) \geq i(a, b)/\mathsf{d}.$$

Therefore, the theorem is true with $\mathsf{Q} = \mathsf{d}$.

Now assume the theorem is true for all strict suborbifolds of S. Let Q' be the largest constant of quasi-isometry necessary for such suborbifolds. Choose the threshold T, constant e and length L such

that Theorem 8.1 holds for both the data $(S, \mathsf{T}, \mathsf{T}, \mathsf{e}, \mathsf{L})$ as well as $(\Sigma, (\mathsf{T}/\mathsf{Q}') - \mathsf{Q}', \mathsf{T}, \mathsf{e}, \mathsf{L})$. We also assume that $\mathsf{T} \geq 2\mathsf{N} + 1$ (Proposition 9.2). All of the constants depend only on the topology of S and the degree d, because these last two bound the topology of Σ.

Let $\mathcal{J}_S$ be the T–antichain of maxima for S, a and b and let $\mathcal{J}_\Sigma$ be the set of components of preimages of elements of $\mathcal{J}_S$.

Claim. The set $\mathcal{J}_\Sigma$ is a $((\mathsf{T}/\mathsf{Q}') - \mathsf{Q}', \mathsf{T})$–antichain for (Σ, α, β).

We check the conditions of being an antichain. Since elements of $\mathcal{J}_S$ are not subsets of each other, the same holds for their preimages. The condition $d_\Omega(\alpha, \beta) \geq (\mathsf{T}/\mathsf{Q}') - \mathsf{Q}'$ is the induction hypothesis. Now suppose $\Psi \subset \Sigma$ with $d_\Psi(\alpha, \beta) \geq \mathsf{T}$. By Lemma 9.3, Ψ is symmetric. That is, it is a component of the preimage of an orbifold $Y \subset S$ and

$$d_Y(a, b) \geq d_\Psi(\alpha, \beta) \geq \mathsf{T}.$$

This implies that $Y \subset Z$ for some $Z \in \mathcal{J}_S$. Therefore, taking Ω to be the preimage of Z, we have $\Psi \subset \Omega \in \mathcal{J}_\Sigma$. This proves the claim.

Hence, there are constants K and K' such that

$$d_S(a, b) \asymp_{\mathsf{K}} |\mathcal{J}_S| + |G(\mathsf{e}, \mathsf{L})|,$$

and

$$d_\Sigma(\alpha, \beta) \asymp_{\mathsf{K}'} |\mathcal{J}_\Sigma| + |\Gamma(\mathsf{e}, \mathsf{L})|.$$

Note that $|\mathcal{J}_S| \leq \mathsf{d}|\mathcal{J}_\Sigma|$ as a suborbifold of S has at most d preimages. Note also that $|G(\mathsf{e}, \mathsf{L})| \leq |\Gamma(\mathsf{e}, \mathsf{L})|$ because $\Gamma(t)$ is at least as thick as $G(t)$. Therefore

$$d_S(a, b) \prec_{\mathsf{Q}} d_\Sigma(\alpha, \beta),$$

for $\mathsf{Q} = \mathsf{d}\,\mathsf{K}\,\mathsf{K}'$. This finishes the proof. □

11. An application to mapping class groups

Suppose that $P \colon \Sigma \to S$ is an orbifold covering map. Let $\mathcal{MCG}(\Sigma)$ be the orbifold mapping class group of Σ: isotopy classes of homeomorphisms of Σ restricting to the identity on ∂S and respecting the set of orbifold points and their orders. Here all isotopies must fix all boundary components and all orbifold points. As an application of Theorem 10.1 we prove the following theorem:

Theorem 11.1. *The covering P induces a quasi-isometric embedding*

$$\Pi_* \colon \mathcal{MCG}(S) \to \mathcal{MCG}(\Sigma).$$

We will use the language of markings from [MM00]. Recall that a *marking* m of S is a collection of curves which fill S. That is, cutting S° along m results in a collection of disks and boundary parallel annuli. If m, n are both markings then we define $i(m,n)$ to be their *intersection number*: the sum of intersection numbers of pairs of curves coming from m and n. Notice for any marking m that there are only finitely many mapping classes $x \in \mathcal{MCG}(S)$ with $x(m) = m$.

Here we establish a few properties of markings.

Lemma 11.2. *For every* N *there are finitely many markings of self-intersection number less than* N *up to the action of the mapping class group.*

Proof. The bound on intersection number provides an upper bound on the number of disk and annuli in $S^\circ \smallsetminus m$. These are glued along edges whose number is also bounded. □

Lemma 11.3. *For every marking* m *and any* $\mathsf{N} > 0$ *there are only finitely many markings* n *with* $i(m,n) \leq \mathsf{N}$.

Proof. The restriction of n to a component of $S^\circ \smallsetminus m$ is a union of arcs. The number of these arcs is bounded by $i(m,n)$. Therefore, the combinatorial type of the collection of arcs is bounded depending on m and $i(m,n)$. □

Lemma 11.4. *For every* $\mathsf{N}_1 > 0$ *there is* $\mathsf{N}_2 > 0$ *such that if* m *and* n *are two markings with self-intersection numbers less than* N_1 *then there is a mapping class* $x \in \mathcal{MCG}(S)$ *such that* $i\big(x(m), n\big) \leq \mathsf{N}_2$.

Proof. Let $[m_1], \ldots, [m_k]$ be the homeomorphism classes of markings that have self-intersection number less than N_1; by Lemma 11.2 there are finitely many such classes. Define $i\big([m_i],[m_j]\big)$ to be the minimum intersection number between a marking in $[m_i]$ and a marking in $[m_j]$. The the lemma now holds for

$$\mathsf{N}_2 = \max_{i,j} i\big([m_i],[m_j]\big). \qquad \square$$

Lemma 11.5. *Let* Θ *be a generating set of* $\mathcal{MCG}(\Sigma)$ *and let* μ *be a marking of* Σ. *For every* $\mathsf{N} > 0$ *there is* $\mathsf{W} > 0$ *with the following property: for any* $\zeta \in \mathcal{MCG}(\Sigma)$,

$$i\big(\mu, \zeta(\mu)\big) \leq \mathsf{N} \quad \Longrightarrow \quad ||\zeta||_\Theta \leq \mathsf{W}.$$

Here, W *depends on* $\xi(\Sigma)$, Θ, μ *and* N *but is independent of* ζ.

Proof. Lemma 11.3 implies that there are only finitely many marking μ' so that μ' is a homeomorphic image of μ and $i(\mu,\mu') \leq \mathsf{N}$. For each

such μ' there are only finitely many mapping classes taking μ to μ' (the marking μ may have symmetries). Let W be the maximum word length of all these mapping classes. □

Proof of Theorem 11.1. Fix, for the remainder of the proof, a marking m of S. Let $\mu = \Pi(m)$ be the lift of m to Σ. Note that μ fills Σ and so is a marking.

We construct $\Pi_\star$ as follows: Let x be an element of $\mathcal{MCG}(S)$ and let μ' be the lift of $x(m)$. The markings m and $x(m)$ have equal self-intersection. Therefore, the same holds for μ and μ'. By Lemma 11.4, there is an N_2 depending only on the self-intersection number of μ such that one can always find $\xi \in \mathcal{MCG}(\Sigma)$ where $i(\mu', \xi(\mu)) \leq \mathsf{N}_2$. Also, it follows from Lemma 11.3 that there are only finitely many such mapping classes. We define $\Pi_\star(x)$ to be any such mapping class ξ.

Let T be a finite generating set for $\mathcal{MCG}(S)$ and Θ be a finite generating set for $\mathcal{MCG}(\Sigma)$. Let $||x||_T$ and $||\xi||_\Theta$ denote the word lengths of x and ξ with respect to T and Θ respectively. To prove the proposition it suffices to show that, for $\xi = \Pi_\star(x)$,

$$||x||_T \asymp_{\mathsf{W}} ||\xi||_\Theta, \tag{11.6}$$

where W is a constant that does not depend on x.

By [MM00], Theorems 7.1, 6.10, and 6.12, we have

$$||x||_T \asymp_{\mathsf{W}_1} \sum \Big[d_Z\big(m, x(m)\big) \Big]_{\mathsf{k}_1}. \tag{11.7}$$

Here the sum ranges over all sub-orbifolds $Z \subset S$. The constant W_1 depends on k_1 which in turn depends on our choice of the marking m and the generating set T. However, all of the choices are independent of the group element x. Finally, $[r]_\mathsf{k} = r$ if $r \geq \mathsf{k}$ and $[r]_\mathsf{k} = 0$ if $r < \mathsf{k}$.

As above, after fixing a large enough constant k_2 (see below) and an appropriate W_2, we have

$$||\xi||_\Theta \asymp_{\mathsf{W}_2} \sum \Big[d_\Omega\big(\mu, \xi(\mu)\big) \Big]_{\mathsf{k}_2}. \tag{11.8}$$

But $\xi(\mu)$ and $\Pi(x(m))$ have bounded intersection. Therefore, their projection distance in every subsurface Ω is *a priori* bounded. Hence we can write

$$||\xi||_\Theta \asymp_{\mathsf{W}_3} \sum \Big[d_\Omega\big(\mu, \Pi(x(m))\big) \Big]_{\mathsf{k}_2}, \tag{11.9}$$

for a slightly larger constant W_3.

We prove equation (11.6) by comparing the terms of the the right hand side of (11.7) with those on the right hand side of (11.9). Note that $\mu = \Pi(m)$ is a union of symmetric orbits and the same holds for $\Pi(x(m))$. Therefore, we can choose k_2 large enough such that if

$d_\Omega(\mu, g(\mu))$ is larger than k_2 then Ω is itself symmetric (see Lemma 9.3). Taking $Z = P(\Omega)$, it follows from Theorem 10.1 that

$$d_Z(m, x(m)) \asymp d_\Omega(\mu, \Pi(x(\mu))).$$

On the other hand, Theorem 10.1 also tells us that large projection distance in any $Z \subset S$ implies large projection distance in all the components of the pre-image of Z. Therefore, there is a finite-to-one correspondence between the surfaces that appear in (11.9) and in (11.7) and the corresponding projection distances are comparable. We conclude that $||x||_T \asymp_{\mathsf{W}} ||\xi||_\Theta$ for some W. This finishes the proof. □

Assume now $\Delta < \mathcal{MCG}(\Sigma)$ is a finite subgroup. Applying Nielsen Realization [Ker83] the group Δ can be realized as a group of homeomorphisms of Σ. Let S be the quotient and let $P\colon \Sigma \to S$ be the regular covering with deck group Δ. Let $N(\Delta)$ be the normalizer of Δ inside of $\mathcal{MCG}(\Sigma)$ and let $M < \mathcal{MCG}(S)$ be the finite index subgroup of mapping classes that lift. MacLachlan and Harvey [MH75, Theorem 10] give a short exact sequence:

$$1 \to \Delta \to N(\Delta) \xrightarrow{p} M \to 1.$$

Theorem 11.10. *Suppose that $\Delta \subset \mathcal{MCG}(\Sigma)$ is a finite group. Then the normalizer of Δ in undistorted in $\mathcal{MCG}(\Sigma)$.*

Proof of Theorem 11.10. Choose finite generating sets Θ for $\mathcal{MCG}(\Sigma)$ and Θ' for $N(\Delta)$. Equip the groups with the word metric. For $\zeta \in N(\Delta)$, we must show that the word length of ζ with respect to Θ is comparable to its word length with respect to Θ'.

Let M be as in the MacLachlan-Harvey short exact sequence. Choose finite generating sets T for $\mathcal{MCG}(S)$ and T' for M. Again, equip these groups with the word metric.

$$\begin{array}{ccccccccc} & & & & \mathcal{MCG}(\Sigma) & \xleftarrow{\Pi_\star} & \mathcal{MCG}(S) & & \\ & & & & \uparrow{\scriptstyle \iota} & & \uparrow{\scriptstyle \iota} & & \\ 1 & \longrightarrow & \Delta & \longrightarrow & N(\Delta) & \xrightarrow{p} & M & \longrightarrow & 1 \end{array}$$

The map $p\colon N(\Delta) \to M$ is a quasi-isometry because Δ is finite. Therefore,

$$||\zeta||_{\Theta'} \asymp ||p(\zeta)||_{T'}.$$

Also, since M is a finite index subgroup of $\mathcal{MCG}(S)$, the word metric in M and the metric it inherits from $\mathcal{MCG}(S)$ are comparable. Hence,

$$||p(\zeta)||_{T'} \asymp ||p(\zeta)||_T.$$

Let $\xi = \Pi_\star p(\zeta)$, where $\Pi_\star$ is an in the proof of Theorem 11.1. Theorem 11.1 states that $\Pi_\star$ is a quasi-isometric embedding. That is,

$$||p(\zeta)||_T \asymp ||\Pi_\star\, p(\zeta)||_\Theta.$$

Also, by the definition of $\Pi_\star$, the intersection number of $\xi(\mu)$ and $\zeta(\mu)$ is bounded. It follows that the intersection number of μ and $\xi^{-1}\,\zeta(\mu)$ is also bounded. Lemma 11.5 implies that ξ and ζ are close in the mapping class group. That is

$$||\Pi_\star\, p(\zeta)||_\Theta \asymp ||\zeta||_\Theta.$$

The theorem follows from the last four equations. □

As a a special case, let Σ be the closed orientable surface of genus g and let $\phi\colon \Sigma \to \Sigma$ be a hyperelliptic involution. Let $S = \Sigma/\phi$ and let $P\colon \Sigma \to S$ be the induced orbifold cover. Birman and Hilden [BH73] provide a short exact sequence:

$$1 \to \langle\phi\rangle \to N(\phi) \to \mathcal{MCG}(S) \to 1$$

which has a group-theoretic section. Notice that $\mathcal{MCG}(S)$ is the spherical braid group on $2g + 2$ strands. Theorem 11.10 now answers a question of Luis Paris:

Corollary 11.11. *The section of the Birman-Hilden map induces a quasi-isometric embedding of the spherical braid group on $2g+2$ strands into the mapping class group of the closed surface of genus g.* □

References

[Ber60] Lipman Bers. Quasiconformal mappings and Teichmüller's theorem. In *Analytic functions*, pages 89–119. Princeton Univ. Press, Princeton, N.J., 1960.

[BFP] Nathan Broaddus, Benson Farb, and Andrew Putman. Irreducible sp– representations and subgroup distortion in the mapping class group. arXiv:0707.2262.

[BH73] Joan S. Birman and Hugh M. Hilden. On isotopies of homeomorphisms of Riemann surfaces. *Ann. of Math. (2)*, 97:424–439, 1973.

[BKMM] Jason Behrstock, Bruce Kleiner, Yair Minsky, and Lee Mosher. Geometry and rigidity of mapping class groups. arXiv:0801.2006.

[BS85] Joan S. Birman and Caroline Series. Geodesics with bounded intersection number on surfaces are sparsely distributed. *Topology*, 24(2):217–225, 1985.

[Bus92] Peter Buser. *Geometry and spectra of compact Riemann surfaces*, volume 106 of *Progress in Mathematics*. Birkhäuser Boston Inc., Boston, MA, 1992.

[Gar87] Frederick P. Gardiner. *Teichmüller theory and quadratic differentials*. Pure and Applied Mathematics (New York). John Wiley & Sons Inc., New York, 1987. , A Wiley-Interscience Publication.

[Hama] Ursula Hamenstaedt. Geometry of the mapping class groups II: (Quasi)-geodesics. arXiv:math.GR/0511349.

[Hamb] Ursula Hamenstaedt. Geometry of the mapping class groups III: Quasi-isometric rigidity. arXiv:math/0512429.

[Har86] John L. Harer. The virtual cohomological dimension of the mapping class group of an orientable surface. *Invent. Math.*, 84(1):157–176, 1986.

[Hem01] John Hempel. 3-manifolds as viewed from the curve complex. *Topology*, 40(3):631–657, 2001. arXiv:math/9712220.

[Ker83] Steven P. Kerckhoff. The Nielsen realization problem. *Ann. of Math. (2)*, 117(2):235–265, 1983.

[MH75] C. Maclachlan and W. J. Harvey. On mapping-class groups and Teichmüller spaces. *Proc. London Math. Soc. (3)*, 30(part 4):496–512, 1975.

[Min99] Yair N. Minsky. The classification of punctured-torus groups. *Ann. of Math. (2)*, 149(2):559–626, 1999. arXiv:math/9807001.

[MM99] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.*, 138(1):103–149, 1999. arXiv:math/9804098.

[MM00] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000. arXiv:math/9807150.

[Mos96] Lee Mosher. Hyperbolic extensions of groups. *J. Pure Appl. Algebra*, 110(3):305–314, 1996.

[MS07] Howard A. Masur and Saul Schleimer. The geometry of the disk complex. 2007.

[Raf05] Kasra Rafi. A characterization of short curves of a Teichmüller geodesic. *Geom. Topol.*, 9:179–202 (electronic), 2005. arXiv:math/0404227.

[Raf07] Kasra Rafi. A combinatorial model for the Teichmüller metric. *Geom. Funct. Anal.*, 17(3):936–959, 2007. arXiv:math/0509584.

[RS08] Kasra Rafi and Saul Schleimer. Curve complexes with connected boundary are rigid, 2008. arXiv:math/0710.3794.

[Sch] Saul Schleimer. Notes on the curve complex. http://math.rutgers.edu/∼saulsch/Maths/notes.pdf.

[Sco83] Peter Scott. The geometries of 3-manifolds. *Bull. London Math. Soc.*, 15(5):401–487, 1983.

E-mail address: rafi@math.uchicago.edu
URL: http://www.math.uchicago.edu/∼rafi

E-mail address: saulsch@math.rutgers.edu
URL: http://www.math.rutgers.edu/∼saulsch