

Covers and the curve complex

KASRA RAFI
SAUL SCHLEIMER

We provide the first nontrivial examples of quasi-isometric embeddings between curve complexes; these are induced by orbifold covers. This leads to new quasi-isometric embeddings between mapping class groups. As a corollary, in the mapping class group normalizers of finite subgroups are undistorted.

57M99; 30F99

1 Introduction

The coarse structure of the complex of curves was first studied by Masur and Minsky [14]. Central in low-dimensional topology, the curve complex sheds light on the algebra of the mapping class group, the global geometry of Teichmüller space and the fine structure of hyperbolic three-manifolds. It is also relevant to the geometric study of other combinatorial spaces with a mapping class group action such as the mapping class group itself, the pants complex, the Hatcher–Thurston complex and the arc complex. Little is known about the subspace structure of these; thus, we propose:

Problem 1.1 Classify quasi-isometric embeddings between combinatorial spaces, as given above.

In this paper we produce the first examples of quasi-isometric embeddings of lower into higher complexity curve complexes, yielding new embeddings between mapping class groups. Our embeddings are induced by orbifold covering maps. We also briefly discuss other topological operations, namely puncturing and taking subsurfaces.

Covering

Suppose that S is a compact connected orientable orbifold of dimension two with nonpositive orbifold Euler characteristic. Let S° denote the surface with boundary obtained by removing an open neighborhood of the orbifold points. Define $\mathcal{C}(S)$ to be the curve complex of S° (see Definition 2.1). We also define the complexity of S :

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

Let $P: \Sigma \rightarrow S$ be an orbifold covering map. The covering P defines a one-to-many relation $\Pi: \mathcal{C}(S) \rightarrow \mathcal{C}(\Sigma)$; the curve $b \in \mathcal{C}(S)$ is related to $\beta \in \mathcal{C}(\Sigma)$ if $P(\beta) = b$. We will call β a *lift* of the curve b and say that β is *symmetric*.

Theorem 8.1 *The covering relation $\Pi: \mathcal{C}(S) \rightarrow \mathcal{C}(\Sigma)$ is a Q -quasi-isometric embedding. The constant Q depends only on $\xi(S)$ and the degree of the covering map.*

Theorem 8.1 is surprising in light of 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 the proof that the relation Π is everywhere defined.

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

Theorem 9.1 *The covering P induces a quasi-isometric embedding*

$$\Pi_*: \mathcal{MCG}(S) \rightarrow \mathcal{MCG}(\Sigma).$$

When the cover is regular, a stronger statement holds.

Theorem 9.6 *Suppose that $\Delta \subset \mathcal{MCG}(\Sigma)$ is a finite subgroup. Then the normalizer of Δ is undistorted in $\mathcal{MCG}(\Sigma)$.*

Note that many algebraically defined subgroups of the mapping class group, such as the Torelli group, are distorted; see Broaddus, Farb and Putman [4].

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 follows directly from a construction of Harer [9, Lemma 3.6]: choose a hyperbolic metric on S and remove a point in the complement of all simple closed geodesics. Choosing an identification of Σ and the punctured surface now gives an embedding $\Pi: \mathcal{C}(S) \rightarrow \mathcal{C}(\Sigma)$.

Theorem 1.2 *The construction above yields uncountably many isometric embeddings of $\mathcal{C}(S)$ into $\mathcal{C}(\Sigma)$. □*

The same construction gives quasi-isometric embeddings of $\mathcal{MCG}(S)$ into $\mathcal{MCG}(\Sigma)$. This result for mapping class groups was previously obtained by Mosher via a different technique [18, Quasi-isometric section lemma].

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 $\Sigma \setminus S$ has no annular components. The inclusion of S into Σ induces an obvious, but important, simplicial injection of curve complexes. This injection is not 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 [15, Theorems 7.1, 6.10 and 6.12] and was independently obtained by Hamenstädt [7, Theorem B and Corollary 4.6].

Quasi-isometry group

A special, but quite deep, instance of Problem 1.1 is the computation of the quasi-isometry group. This has recently been obtained for the mapping class group by Behrstock, Kleiner, Minsky and Mosher [1] and also by Hamenstädt [8]. They show that the quasi-isometry group is virtually equal to the isometry group; in other words, 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 [21] that the quasi-isometry group of the curve complex is again the mapping class group.

Outline of the proof of Theorem 8.1

Suppose that $P: \Sigma \rightarrow S$ 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 the relation $\Pi: \mathcal{C}(S) \rightarrow \mathcal{C}(\Sigma)$ 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 $\xi(S) > 1$.

Two steps are required to reverse the inequality. We first give a new estimate of distance in the complex of curves (Theorem 6.1). We then analyze 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 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 4.4).

Next, we introduce (T_0, T_1) -*antichains*: this is an antichain, in the poset of subsurfaces of S , with *thresholds* T_0 and T_1 (Section 5). 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 6.1).

Let Γ be the lift of G , which is again a Teichmüller geodesic with identical parameterization. This follows from the well-known fact that covering maps induce isometric embeddings of Teichmüller spaces (Section 7). 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, perhaps with weaker thresholds. A key point is Lemma 7.2, stating that any subsurface Ω of Σ , having $d_\Omega(\alpha, \beta)$ large, is symmetric. In fact, such Ω is a lift of a subsurface Z where $d_Z(a, b)$ is large and we may apply induction. Therefore, the estimate for $d_S(a, b)$ and $d_\Sigma(\alpha, \beta)$ are comparable (Theorem 8.1).

Our use of Teichmüller geodesics appears unavoidable: for example, Masur–Minsky hierarchies do not, a priori, have good properties with respect to covering maps. 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, antichains choose the *correct* subset.

Acknowledgments We thank the referees for useful comments. We thank the Mathematical Sciences Research Institute for its hospitality in the fall of 2007.

2 Background

We begin by introducing some convenient notation: if A, B, c are nonnegative real numbers with $c > 0$ and if $A \leq cB + c$, then we write $A \prec_c B$. If $A \prec_c B$ and $B \prec_c A$, then we write $A \asymp_c B$.

Suppose \mathcal{X} and \mathcal{Y} are metric spaces and $f: \mathcal{X} \rightarrow \mathcal{Y}$ is a relation. Then 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 $d_{\mathcal{X}}(x, x') \asymp_c d_{\mathcal{Y}}(y, y')$. We say that f is a c -*quasi-isometry* if additionally a c -neighborhood of $f(\mathcal{X})$ equals \mathcal{Y} .

Suppose that Σ is a compact orientable orbifold, of dimension two, with nonpositive orbifold Euler characteristic. For definitions and discussion of orbifolds we refer the

reader to Scott's excellent article [23]. 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 [10] When $\xi(\Sigma) > 1$ the *complex of curves* $\mathcal{C}(\Sigma^\circ)$ has as its vertices isotopy classes of essential, nonperipheral 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 *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 resulting complex is 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 suffices to study the one-skeleton of $\mathcal{C}(\Sigma)$, for which we use the same notation, because the one-skeleton and the entire complex are quasi-isometric.

3 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 nonperipheral curve in Σ . All suborbifolds considered will be cleanly embedded.

From Masur and Minsky [14], recall the definition of the *subsurface projection* relation

$$\pi_\Psi: \mathcal{C}(\Sigma) \rightarrow \mathcal{C}(\Psi)$$

defined when Ψ° is not a three-holed sphere. Suppose first that Σ has negative orbifold Euler characteristic and Ψ° is not an annulus. 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 $\pi_1^{\text{orb}}(\Psi) \rightarrow \pi_1^{\text{orb}}(\Sigma)$. 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 $\alpha'' \cup \partial\Sigma'$. Since Ψ° is not a three-holed sphere there is a boundary component α''' of N which is essential and nonperipheral. We then write $\alpha''' \in \pi_\Psi(\alpha)$.

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

It remains to deal with the case where $\chi^{\text{orb}}(\Sigma) = 0$ and Ψ is a cleanly embedded annulus. Here Σ is either a torus or a *square pillow*: a sphere with four orbifold points of order two. In either case, fix a flat metric on Σ . Let ψ be the core curve of the annulus Ψ . Isotope α and ψ to be geodesic. When Σ is a torus we cut along ψ and take the closure of the resulting annulus; $\pi_\psi(\alpha)$ is the remains of α . When Σ is a square pillow, the geodesic ψ is replaced by the parallel geodesic arcs connecting pairs of orbifold points.

We say a 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

$$d_\Psi(\alpha, \beta) = \text{diam}_\Psi(\pi_\Psi(\alpha) \cup \pi_\Psi(\beta)).$$

The bounded geodesic image theorem, due to Masur and Minsky [15], states:

Theorem 3.1 *Fix a surface Σ . There is a constant $M = M(\Sigma)$ 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 $d_\Omega(\alpha, \beta) \geq M$ then there is a vertex of Λ which misses Ω . \square*

4 Teichmüller space

In 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 Bers [2] and Gardiner [6].

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 $e_0 = e_0(\Sigma) > 0$ 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 $\Gamma: [t_\sigma, t_\tau] \rightarrow \mathcal{T}(S)$ 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 next result follows from the proof of [20, Proposition 3.7]; see also [19, Theorem 5.5].

Theorem 4.1 (Large implies short.) *For every positive $e \leq e_0$ there is a threshold T_e such that, for a strict subsurface Ω of Σ , if $d_\Omega(\alpha, \beta) \geq T_e$ 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 . \square*

There is also a converse [19, Theorems 6.1, 7.1, 7.3].

Theorem 4.2 (Short implies large.) *For every threshold T_1 there is a constant e_1 such that if $l_t(\gamma) \leq e_1$ for some curve γ and for some time t , then there exists a subsurface Ψ disjoint from γ having $d_\Psi(\alpha, \beta) \geq T_1$. \square*

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 [14].)

Theorem 4.3 *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 $B = B(\Sigma)$ 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*

$$d_\Sigma(\gamma_0, \gamma_2) \geq d_\Sigma(\gamma_0, \gamma_1) + d_\Sigma(\gamma_1, \gamma_2) - B. \quad \square$$

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 4.4 *For every $e > 0$ there is a progress constant $P > 0$ 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*

$$d_\Sigma(\gamma_0, \gamma_1) \asymp_P t_1 - t_0.$$

Proof Since $\Gamma(t)$ is ϵ -thick at every time $t \in [t_0, t_1]$, from Theorem 4.1 we deduce that $d_\Omega(\gamma_0, \gamma_1) \leq T_\epsilon$ for every strict subsurface Ω of Σ . The lemma is then a consequence of Theorem 1.1 and Remark 5.5 in [20]. (Referring to the statement and notation of [20, Theorem 1.1]: Extend γ_i to a short marking μ_i . Take k large enough such that the only nonzero term in the right hand side of [20, Equation (1)] is $d_\Sigma(\mu_0, \mu_1)$.) \square

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

Lemma 4.5 *For every ϵ there exists L_ϵ such that if $L \geq L_\epsilon$, then*

$$d_\Sigma(\alpha, \beta) \geq |\Gamma(\epsilon, L)|/2P.$$

Proof Pick L_ϵ large enough so that, for $L \geq L_\epsilon$,

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

Realize $\Gamma(\epsilon, L)$ as the union of intervals $[t_i, s_i]$, $i = 1, \dots, m$. 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 4.3 we have

$$d_\Sigma(\alpha, \beta) \geq \left(\sum_i d_\Sigma(\gamma_i, \delta_i) \right) - 2mB.$$

From Lemma 4.4 we deduce

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

Rearranging, we find

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

Thus, as desired:

$$d_\Sigma(\alpha, \beta) \geq \frac{1}{2P}|\Gamma(\epsilon, L)|. \quad \square$$

5 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 a time, Theorem 4.2 gives a suborbifold Ω with $d_\Omega(\alpha, \beta)$ large. However, the number of these suborbifolds is not a good estimate for the distance in the complex of curves; many suborbifolds 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 $T_1 \geq T_0 > 0$. We say that a set \mathcal{J} of suborbifolds $\Omega \subsetneq \Sigma$, is a (T_0, 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 $d_\Omega(\alpha, \beta) \geq T_0$.
- If $\Psi \subsetneq \Sigma$ and $d_\Psi(\alpha, \beta) \geq T_1$ then Ψ is a suborbifold of some element of \mathcal{J} .

Notice that there may be many different antichains for the given data $(\Sigma, \alpha, \beta, T_0, T_1)$. One particularly nice example is when $T_0 = T_1 = T$ and \mathcal{J} is defined to be the maxima of the set

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

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 5.1 *For every orbifold Σ and for every pair of sufficiently large thresholds T_0, T_1 , there is an accumulation constant $A_\Sigma = A(\Sigma, T_0, T_1)$ so that if \mathcal{J} is an (T_0, T_1) -antichain for (Σ, α, β) , then*

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

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 projection distance $d_\gamma(\alpha, \beta) \geq T_0$. In this case, assuming $T_0 > 3$, every such curve γ is a vertex of every geodesic connecting α to β [17, Section 4]. Therefore the lemma holds for Farey graphs with $A_\Sigma = 1$.

Now assume the lemma is true for all surfaces with complexity less than $\xi(\Sigma)$. Note that there exists a constant C so that if $\Omega \subset \Psi \subset \Sigma$ and α', β' are the projections of α, β to Ψ then

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

We take thresholds T_0 and T_1 for Σ large enough so that for the (T_0, T_1) -antichain \mathcal{J} we have:

- The lemma applies to any strict suborbifold Ψ with thresholds $T_0 - C, T_1 + C$.
- $T_0 \geq M(\Sigma)$; thus by Theorem 3.1 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

$$A_\Psi = A(\Psi, T_0 - C, T_1 + C) \quad \text{and} \quad \mathcal{J}_\Psi = \{\Omega \in \mathcal{J} \mid \Omega \subsetneq \Psi\}.$$

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

$$|\mathcal{J}_\Psi| \leq A_\Psi \cdot (T_1 + C).$$

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

$$d_\Psi(\alpha, \beta) < T_1.$$

Let α' and β' be the projections of α and β to Ψ . From the definition of C , \mathcal{J}_Ψ is a $(T_0 - C, T_1 + C)$ -antichain for Ψ, α' and β' . Thus,

$$T_1 > d_\Psi(\alpha, \beta) \geq d_\Psi(\alpha', \beta') - C \geq |\mathcal{J}_\Psi|/A_\Psi - C,$$

with the last inequality being the induction hypothesis. Hence,

$$T_1 + C \geq |\mathcal{J}_\Psi|/A_\Psi. \quad \square$$

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

$$(A_\Psi + A_{\Psi'})(T_1 + C) + 3.$$

Since every orbifold in \mathcal{J} is disjoint from some vertex of Λ , the lemma holds for $A(\Sigma, T_0, T_1)$ equal to

$$2 \cdot \max\{A_\Psi \mid \Psi \subsetneq \Sigma\} \cdot (T_1 + C) + 3. \quad \square$$

6 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 defined in Section 4. We choose thresholds $T_0 \geq T_{e_0}$ (see Theorem 4.1) and T_1 so that Lemma 5.1 holds. Let e_1 be the constant provided in Theorem 4.2 and let $e > 0$ be any constant smaller than $\min\{e_0, e_1\}$. Finally, we pick L_e such that Lemma 4.5 holds and such that $L_e/2P > 4$. Let L be any length larger than L_e .

Theorem 6.1 *Let T_0, T_1, e and L be constants chosen as above. There is a constant $K = K(\Sigma, T_0, T_1, e, L)$ such that for any curves α and β , any (T_0, T_1) -antichain \mathcal{J} and any Teichmüller geodesic Γ , chosen as above, we have*

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

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

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

follows from Lemmas 5.1 and 4.5. It remains to show that

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

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 4.1). Let \mathcal{E} be the union:

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

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

Claim *The number of intervals in $\Gamma(e, L)$ 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 $I, J \subset \Gamma(e, L)$ when some curve γ becomes e -short (and hence e_1 -short). Therefore, by Theorem 4.2, γ is disjoint from a subsurface Ψ where $d_\Psi(\alpha, \beta) \geq T_1$. Since \mathcal{J} is an (T_0, 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 4.5 to Γ' , we find

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

and therefore Ω is not equal to Ω' . \square

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

Claim

$$d_\Sigma(\gamma_i, \gamma_{i+1}) \leq \begin{cases} P(t_{i+1} - t_i) + P, & \text{if } [t_i, t_{i+1}] \subset \Gamma(e, L), \\ 2B + PL + P + 4, & \text{otherwise.} \end{cases}$$

Proof The first case follows from Lemma 4.4. So suppose that the interior of $[t_i, t_{i+1}]$ is disjoint from the interior of $\Gamma(e, L)$.

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 $d_\Sigma(\gamma_i, \gamma) \leq B + 2$. 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

$$2 \geq d_\Sigma(\partial\Omega, \gamma) \geq d_\Sigma(\gamma_i, \gamma) - B.$$

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

If I_+ and I_- have nonempty 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 with 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

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

We also know from Lemma 4.4 that

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

This finishes the proof of our claim. □

It follows that

$$\begin{aligned} d_\Sigma(\alpha, \beta) &\leq d_\Sigma(\gamma_0, \gamma_1) + \dots + d_\Sigma(\gamma_{n-1}, \gamma_n) \\ &\leq (2B + PL + P + 4)|\mathcal{E}| + P|\Gamma(e, L)| + P|\mathcal{E}| \\ &<_K |\mathcal{J}| + |\Gamma(e, L)| \end{aligned}$$

for an appropriate choice of K . This proves Theorem 6.1. □

7 Symmetric curves, surfaces and metrics

Suppose that S is a compact connected orientable orbifold, of dimension two, with nonpositive orbifold Euler characteristic. Let $P: \Sigma \rightarrow S$ be an orbifold covering map. Recall that the covering P defines a relation $\Pi: \mathcal{C}(S) \rightarrow \mathcal{C}(\Sigma)$: a curve $b \in \mathcal{C}(S)$ is related to $\beta \in \mathcal{C}(\Sigma)$ if $P(\beta) = b$. If $\beta \in \Pi(b)$ then we call β *symmetric* and say that β is a *lift* of b . Similarly, a suborbifold $\Omega \subset \Sigma$ is *symmetric* if it is a component of $P^{-1}(Z)$ where Z is a suborbifold of S .

Lemma 7.1 *The covering relation Π is everywhere defined.*

Proof We will show that if a is an essential nonperipheral curve then every component of $P^{-1}(a)$ is essential and nonperipheral. Since S has nonpositive 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 nonperipheral simple closed curve that is homotopic to the lift of a^* . The conclusion follows. \square

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.

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 $G: [t_x, t_y] \rightarrow \mathcal{T}(S^\circ)$ 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 $\Gamma: [t_x, t_y] \rightarrow \mathcal{T}(\Sigma^\circ)$. 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 $k|q|/q$ 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 $k|\theta|/\theta$, where the quadratic differential θ is the pullback of q to $\Gamma(t)$. That is, the lift of the Teichmüller map

from $G(t)$ 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 $\Gamma(s)$.

Lemma 7.2 *For ϵ small enough, T_ϵ as in Theorem 4.1 and any suborbifold $\Omega \subset \Sigma$ if $d_\Omega(\alpha, \beta) \geq 2T_\epsilon + 1$, then Ω is symmetric.*

Proof Consider the first and last times t^\pm that every component of $\partial\Omega$ is ϵ -short in $\Gamma(t^\pm)$. Thus the image $P(\partial\Omega)$ is ϵ -short in $G(t^\pm)$. Therefore, all components of the image are simple. (This is a version of the Collar Lemma. For example, see [5, Theorem 4.2.2].) It follows that the boundary of Ω is symmetric.

To show that Ω is symmetric it now suffices to show that $\Omega \cap P^{-1}(P(\partial\Omega)) = \partial\Omega$. So suppose that there is a curve γ in the intersection which is not a boundary component. Then γ has bounded length (in fact, length at most $d \cdot \epsilon$) at the times t^\pm . It follows from Theorem 4.1 that $d_\Omega(\alpha, \gamma), d_\Omega(\gamma, \beta) \leq T_\epsilon$, contradicting the assumption $d_\Omega(\alpha, \beta) \geq 2T_\epsilon + 1$. \square

8 The quasi-isometric embedding

Theorem 8.1 *The covering relation $\Pi: \mathcal{C}(S) \rightarrow \mathcal{C}(\Sigma)$, corresponding to the covering map $P: \Sigma \rightarrow S$, is a Q -quasi-isometric embedding. The constant Q depends only on $\xi(S)$ and the degree of the covering map.*

Remark 8.2 The constant Q may go to infinity with the degree of the covering map; the distance $d_\Sigma(\alpha, \beta)$ between the lifts may be significantly larger or smaller than $d_S(a, b)$.

To see that the distance may decrease suppose that S is a surface with negative Euler characteristic. For any pair of curves a, b there is a cover admitting lifts α, β which are disjoint. In fact a cover of degree at most 2^{d-1} , where $d = d_S(a, b)$, suffices [11, Lemma 2.3].

In the other direction, suppose that S is a torus and a, b are the curves of slope ± 1 . Let Σ be the pq -cover obtained by unwrapping the zero and infinity slopes p and q times, respectively. The lifts α, β have slopes $\pm p/q$ and so $i(\alpha, \beta) = 2pq$. If p and q are consecutive Fibonacci numbers then the distance $d_\Sigma(\alpha, \beta)$ is essentially equal to the logarithm of the intersection number.

Remark 8.3 Note that Q does not depend directly on the topology of Σ . When $\chi^{\text{orb}}(S) = 0$ (the annulus, torus or square pillow) the degree of the covering map is not determined by the topology of Σ . When the orbifold Euler characteristic is negative the topology of Σ can be bounded in terms of the topology of S and the degree of the covering map.

Remark 8.4 When Σ is the orientation double cover of a nonorientable surface S , Theorem 8.1 is due to Masur and Schleimer [16].

Proof of Theorem 8.1 We first show that

$$d_{\Sigma}(\alpha, \beta) \prec_Q d_S(a, b).$$

When $\xi(S) > 1$ or 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 these cases. If Σ° is a torus, once-holed torus or four-holed sphere then two curves are at distance one when they intersect once or twice, depending on S . The lifts of these curves then intersect at most $2d$ times, where d is the degree of the covering. Thus, the distance between the lifts is at most $2 \log_2(2d) + 2$. (See [22, Lemma 1.21].) Therefore

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

Now we must prove the opposite inequality:

$$d_{\Sigma}(\alpha, \beta) \succ_Q d_S(a, b).$$

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,

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

Therefore, the theorem is true with $Q = d$.

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 with a degree d cover. Choose the threshold T , constant e and length L such that Theorem 6.1 holds for both the data (S, T, T, e, L) as well as $(\Sigma, (T/Q') - 1, Q'T + Q', e, L)$. We also assume that $Q'T + Q' \geq 2T_e + 1$ (Lemma 7.2). 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}_Σ be the set of components of preimages of elements of \mathcal{J}_S .

Claim *The set \mathcal{J}_Σ is a $((T/Q') - 1, Q'T + Q')$ -antichain for (Σ, α, β) .*

Proof 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 $d_\Omega(\alpha, \beta) \geq (T/Q') - 1$ is the induction hypothesis. Now suppose $\Psi \subset \Sigma$ with $d_\Psi(\alpha, \beta) \geq Q'T + Q'$. By Lemma 7.2, Ψ is symmetric. That is, Ψ is a component of the preimage of some $Y \subset S$ and by induction

$$Q'd_Y(a, b) + Q' \geq d_\Psi(\alpha, \beta) \geq Q'T + Q'.$$

Thus $d_Y(a, b) \geq T$ and so $Y \subset Z$ 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. \square

Hence, there are constants K and K' such that

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

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

Note that $|\mathcal{J}_S| \leq |\mathcal{J}_\Sigma|$ as distinct suborbifolds have distinct preimages. Note also that $|G(e, L)| \leq |\Gamma(e, L)|$ because $\Gamma(t)$ is at least as thick as $G(t)$. Therefore

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

for $Q = K(K' + 1)$. This finishes the proof of Theorem 8.1. \square

9 An application to mapping class groups

Suppose that $P: \Sigma \rightarrow S$ 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 8.1 we prove the following theorem:

Theorem 9.1 *The covering P induces a quasi-isometric embedding*

$$\Pi_*: \mathcal{MCG}(S) \rightarrow \mathcal{MCG}(\Sigma).$$

We will use the language of markings from [15]. 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 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 9.2 *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 disks and annuli in $S^\circ \setminus m$. These are glued along edges whose number is also bounded. \square

Lemma 9.3 *For every marking m and any $N > 0$ there are only finitely many markings n with $i(m, n) \leq N$.*

Proof The restriction of n to a component of $S^\circ \setminus m$ 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 $i(m, n)$. \square

Lemma 9.4 *For every $N_1 > 0$ there is an $N_2 > 0$ such that if m and n are two markings, each with self-intersection number less than N_1 , then there is a mapping class $x \in \mathcal{MCG}(S)$ such that $i(x(m), n) \leq N_2$.*

Proof Let $[m_1], \dots, [m_k]$ be the homeomorphism classes of markings that have self-intersection number less than N_1 ; by Lemma 9.2 there are finitely many such classes. Define $i([m_i], [m_j])$ to be the minimum intersection number between a marking in $[m_i]$ and a marking in $[m_j]$. The lemma now holds for

$$N_2 = \max_{i,j} i([m_i], [m_j]). \quad \square$$

Lemma 9.5 *Let Θ be a generating set of $\mathcal{MCG}(\Sigma)$ and let μ be a marking of Σ . For every $N > 0$ there is $W > 0$ with the following property: for any $\zeta \in \mathcal{MCG}(\Sigma)$,*

$$i(\mu, \zeta(\mu)) \leq N \implies \|\zeta\|_\Theta \leq W.$$

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

Proof Lemma 9.3 implies that there are only finitely many markings μ' so that μ' is a homeomorphic image of μ and $i(\mu, \mu') \leq N$. 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. \square

Proof of Theorem 9.1 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 Π_\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 9.4, 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 $i(\mu', \xi(\mu)) \leq N_2$. Also, it follows from Lemma 9.3 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 $\xi = \Pi_\star(x)$,

$$(1) \quad \|x\|_T \asymp_W \|\xi\|_\Theta,$$

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

By [15, Theorems 7.1, 6.10 and 6.12] we have

$$(2) \quad \|x\|_T \asymp_{W_1} \sum [d_Z(m, x(m))]_{k_1}.$$

Here the sum ranges over all suborbifolds $Z \subset S$. 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, $[r]_k = r$ if $r \geq k$ and $[r]_k = 0$ if $r < k$.

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

$$\|\xi\|_\Theta \asymp_{W_2} \sum [d_\Omega(\mu, \xi(\mu))]_{k_2}.$$

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

$$(3) \quad \|\xi\|_\Theta \asymp_{W_3} \sum [d_\Omega(\mu, \Pi(x(m)))]_{k_2},$$

for a slightly larger constant W_3 .

We prove Equation (1) by comparing the terms of the right hand side of (2) with those on the right hand side of (3). 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 7.2). Taking $Z = P(\Omega)$, it follows from Theorem 8.1 that

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

On the other hand, Theorem 8.1 also tells us that large projection distance in any $Z \subset S$ implies large projection distance in all the components of the preimage of Z . Therefore, there is a finite-to-one correspondence between the surfaces that appear in (3) and in (2) and the corresponding projection distances are comparable. We conclude that $\|x\|_T \asymp_W \|\xi\|_\Theta$ for some W . This finishes the proof. \square

Assume now $\Delta < \mathcal{MCG}(\Sigma)$ is a finite subgroup. Applying Nielsen Realization [12] the group Δ can be realized as a group of homeomorphisms of Σ . Let S be the quotient and let $P: \Sigma \rightarrow S$ 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 [13, Theorem 10] give a short exact sequence:

$$1 \rightarrow \Delta \rightarrow N(\Delta) \xrightarrow{P} M \rightarrow 1.$$

Theorem 9.6 *Suppose that $\Delta \subset \mathcal{MCG}(\Sigma)$ is a finite group. Then the normalizer of Δ is undistorted in $\mathcal{MCG}(\Sigma)$.*

Proof 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.

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

The map $p: N(\Delta) \rightarrow M$ 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 Π_\star is as in the proof of Theorem 9.1. Thus Π_\star is a quasi-isometric embedding. That is,

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

Also, by the definition of Π_\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 9.5 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. \square

As a special case, let Σ be the closed orientable surface of genus g and let $\phi: \Sigma \rightarrow \Sigma$ be a hyperelliptic involution. Let $S = \Sigma/\phi$ and let $P: \Sigma \rightarrow S$ be the induced orbifold cover. Birman and Hilden [3] provide a short exact sequence:

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

which has a group-theoretic section. Notice that $\mathcal{MCG}(S)$ is the spherical braid group on $2g + 2$ strands. Theorem 9.6 now answers a question of Luis Paris:

Corollary 9.7 *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 .* \square

References

- [1] **J Behrstock, B Kleiner, Y N Minsky, L Mosher**, *Geometry and rigidity of mapping class groups* arXiv:0801.2006
- [2] **L Bers**, *Quasiconformal mappings and Teichmüller’s theorem*, from: “Analytic functions”, Princeton Univ. Press (1960) 89–119 MR0114898
- [3] **J S Birman, H M Hilden**, *On the mapping class groups of closed surfaces as covering spaces*, from: “Advances in the theory of Riemann surfaces (Proc. Conf., Stony Brook, N.Y., 1969)”, Ann. of Math. Studies 66, Princeton Univ. Press (1971) 81–115 MR0292082
- [4] **N Broaddus, B Farb, A Putman**, *Irreducible Sp -representations and subgroup distortion in the mapping class group* arXiv:0707.2262
- [5] **P Buser**, *Geometry and spectra of compact Riemann surfaces*, Progress in Math. 106, Birkhäuser, Boston (1992) MR1183224
- [6] **F P Gardiner**, *Teichmüller theory and quadratic differentials*, Pure and Applied Math., Wiley-Interscience, John Wiley & Sons, New York (1987) MR903027

- [7] **U Hamenstädt**, *Geometry of the mapping class groups II: (Quasi)-geodesics* arXiv: math.GR/0511349
- [8] **U Hamenstädt**, *Geometry of the mapping class groups III: Quasi-isometric rigidity* arXiv:math/0512429
- [9] **J L Harer**, *The virtual cohomological dimension of the mapping class group of an orientable surface*, Invent. Math. 84 (1986) 157–176 MR830043
- [10] **W J Harvey**, *Boundary structure of the modular group*, from: “Riemann surfaces and related topics (Proc. Conf., Stony Brook, N.Y., 1978)”, (I Kra, B Maskit, editors), Ann. of Math. Stud. 97, Princeton Univ. Press (1981) 245–251 MR624817
- [11] **J Hempel**, *3–manifolds as viewed from the curve complex*, Topology 40 (2001) 631–657 MR1838999
- [12] **S P Kerckhoff**, *The Nielsen realization problem*, Ann. of Math. (2) 117 (1983) 235–265 MR690845
- [13] **C Maclachlan, W J Harvey**, *On mapping-class groups and Teichmüller spaces*, Proc. London Math. Soc. (3) 30 (1975) 496–512 MR0374414
- [14] **H A Masur, Y N Minsky**, *Geometry of the complex of curves. I. Hyperbolicity*, Invent. Math. 138 (1999) 103–149 MR1714338
- [15] **H A Masur, Y N Minsky**, *Geometry of the complex of curves. II. Hierarchical structure*, Geom. Funct. Anal. 10 (2000) 902–974 MR1791145
- [16] **H A Masur, S Schleimer**, *The geometry of the disk complex*, Preprint (2007)
- [17] **Y N Minsky**, *The classification of punctured-torus groups*, Ann. of Math. (2) 149 (1999) 559–626 MR1689341
- [18] **L Mosher**, *Hyperbolic extensions of groups*, J. Pure Appl. Algebra 110 (1996) 305–314 MR1393118
- [19] **K Rafi**, *A characterization of short curves of a Teichmüller geodesic*, Geom. Topol. 9 (2005) 179–202 MR2115672
- [20] **K Rafi**, *A combinatorial model for the Teichmüller metric*, Geom. Funct. Anal. 17 (2007) 936–959 MR2346280
- [21] **K Rafi, S Schleimer**, *Curve complexes with connected boundary are rigid* arXiv: 0710.3794
- [22] **S Schleimer**, *Notes on the curve complex* Available at <http://math.rutgers.edu/~saulsch/Maths/notes.pdf>
- [23] **P Scott**, *The geometries of 3–manifolds*, Bull. London Math. Soc. 15 (1983) 401–487 MR705527

*Department of Mathematics, University of Chicago
5734 S University Avenue, Chicago, Illinois 60637, USA*

*Mathematics Institute, University of Warwick
Coventry, CV4 7A, UK*

`rafi@math.uchicago.edu, s.schleimer@warwick.ac.uk`

`http://www.math.uchicago.edu/~rafi, http://www.warwick.ac.uk/~masgar/`

Proposed: Benson Farb

Received: 27 July 2008

Seconded: Walter Neumann, Danny Calegari

Revised: 17 April 2009: