Proceedings of the Royal Society of Edinburgh: Section A Mathematics, 1–71, *2026*
DOI:10.1017/prm.2025.23

Rigidity of mapping class groups MOD powers of twists

Giorgio Mangioni
Maxwell Institute and Department of Mathematics, Heriot-Watt University, Edinburgh, United Kingdom (gm2070@hw.ac.uk)

Alessandro Sisto
Maxwell Institute and Department of Mathematics, Heriot-Watt University, Edinburgh, United Kingdom (a.sisto@hw.ac.uk) (corresponding author)

(Received 8 June 2023; revised 26 May 2024; accepted 10 March 2025)

We study quotients of mapping class groups of punctured spheres by suitable large powers of Dehn twists, showing an analogue of Ivanov's theorem for the automorphisms of the corresponding quotients of curve graphs. Then we use this result to prove quasi-isometric rigidity of these quotients, answering a question of Behrstock, Hagen, Martin, and Sisto in the case of punctured spheres. Finally, we show that the automorphism groups of our quotients of mapping class groups are "small", as are their abstract commensurators. This is again an analogue of a theorem of Ivanov about the automorphism group of the mapping class group.

In the process, we develop techniques to extract combinatorial data from a quasi-isometry of a hierarchically hyperbolic space, and use them to give a different proof of a result of Bowditch about quasi-isometric rigidity of pants graphs of punctured spheres.

Keywords: mapping class group; quasi-isometric rigidity; Ivanov's Theorem; Dehn twist quotients

2020 *Mathematics Subject Classification:* Primary: 20F65
Secondary: 57K20

1. Introduction

Dehn twist quotients. Given a finite-type surface S, let $MCG(S)$ be its mapping class group. For every $K \in \mathbb{N}_{>0}$, let DT_K be the subgroup generated by the collection $\{T_\alpha^K\}$, where α varies among all isotopy classes of essential simple closed curves on S, and T_α denotes the Dehn Twist around α. In this paper, we study the properties of quotients of the form MCG/DT_K, where K is chosen to be "sufficiently large".

© The Author(s), 2025. Published by Cambridge University Press on behalf of The Royal Society of Edinburgh. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

There are plenty of reasons of interest in these quotients. Firstly, they can be regarded as "Dehn filling quotients" of mapping class groups, as pointed out and explored in [19], where they are proven to be acylindrically hyperbolic. In fact, these and similar quotients are used in [8] (where they are proven to be hierarchically hyperbolic) to relate questions on finite quotients of mapping class groups to residual finiteness of certain hyperbolic groups; this should be compared with the use of Dehn fillings in the resolution of the virtual Haken conjecture [1].

Moreover, when the surface is closed and hyperbolic, these quotients are conjectured to be the outer automorphism group of the Burnside quotient Γ/Γ^n, where Γ is the fundamental group of the surface [16]. This fact would be the analogue of the celebrated Dehn-Nielsen-Baer theorem [20, 27].

These quotients of mapping class groups also appear naturally in a somewhat unrelated area of mathematics, namely the study of topological quantum field theories (see e.g. [23]). As it turns out, if G is a compact Lie group, every homomorphism $MCG(S_g) \to G$ (with $g \geqslant 3$) factors through the quotient by large enough powers of Dehn twists [4, Corollary 2.6], and this applies in particular to the so-called quantum representations of mapping class groups.

Yet another reason of interest in these quotients is simply that given a group G and a collection of elements $g_1, \dots, g_n$, it is natural to study the normal closure of the set, and the corresponding quotient. In the case of our quotients of mapping class groups, we are considering a collection of conjugacy representatives of Dehn twists, and we are "stabilising" by taking powers.

Rigidity results. We have three main results, analogous to results for mapping class groups, that illustrate three different forms of rigidity of the groups we are considering. In all cases, we consider punctured spheres; we discuss below what would need to be done to cover surfaces with genus. We note that, up to commensurability, we could equivalently regard the quotient groups we consider as quotients of braid groups.

The first result we state is that our quotients are quasi-isometrically rigid, thus answering [8, Question 3] for the case of punctured spheres. We say that two groups G and H are *weakly commensurable* if there exist two finite normal subgroups $L \trianglelefteq H$ and $M \trianglelefteq G$ such that the quotients H/L and G/M have two finite index subgroups that are isomorphic.

THEOREM 1 (Quasi-isometric rigidity). *Let $S = S_{0,b}$ be a punctured sphere, with $b \geqslant 7$ punctures. There exists $K_0 \in \mathbb{N}_{>0}$ such that, if K is a non-trivial multiple of K_0, then $H = MCG(S)/DT_K$ is quasi-isometrically rigid, meaning that if a finitely generated group* G *is quasi-isometric to* H *then* G *and* H *are weakly commensurable.*

Quasi-isometric rigidity of mapping class groups was first proven in [9], see also [24], but our argument is closer to the proof given in [7, Section 5]. The key result from that paper is [7, Theorem 5.7], which roughly speaking allows one to extract an automorphism of a certain graph from a quasi-isometry of a hierarchically hyperbolic space satisfying suitable assumptions. We cannot use this result as stated, since our quotients do not satisfy its assumptions; instead, we will show that a similar result holds under different hypotheses that our quotient groups do satisfy. This is done in Theorem 8.7, which we believe to be of independent interest since we enlarge the family of hierarchically hyperbolic spaces to which one can apply the arguments from [7, Section 5] to study quasi-isometric rigidity.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

In our case of interest, the graph from Theorem 8.7 can be related to a certain subgraph of the quotient of the curve graph. From there we are able to prove quasi-isometric rigidity of our quotient groups, adapting arguments due to Bowditch, especially from [13], where the author proved quasi-isometric rigidity of pants graphs.

Combining quasi-isometric rigidity and results on automorphisms of acylindrically hyperbolic groups from [2], we are able to deduce the following, which can be regarded as a form of algebraic rigidity. Here, $MCG^{\pm}$ denotes the extended mapping class group, where orientation-reversing mapping classes are allowed.

THEOREM 2 (Algebraic rigidity). *Let $S = S_{0,b}$ be a punctured sphere, with $b \geqslant 7$ punctures. There exists $K_0 \in \mathbb{N}_{>0}$ such that, if K is a non-trivial multiple of K_0, then*

(1) $Aut(MCG(S)/DT_K) = MCG^{\pm}(S)/DT_K$;
(2) $Out(MCG(S)/DT_K) \cong \mathbb{Z}/2\mathbb{Z}$;
(3) *The abstract commensurator of $MCG^{\pm}(S)/DT_K$ is isomorphic to its inner automorphism group. In other words, any isomorphism between finite index subgroups of $MCG^{\pm}(S)/DT_K$ is the restriction of an inner automorphism.*

In the case of mapping class groups, the analogue of the result above was proven by Ivanov in his famous paper on automorphisms of curve graphs [28]. Indeed, our quasi-isometric rigidity and in turn algebraic rigidity results rely on an analogue of Ivanov's theorem for quotients of the *curve graph* $\mathcal{C}(S)$ of S, that is, the graph whose vertices are isotopy classes of essential simple closed curves and edges correspond to disjointness (up to isotopy). This is our final main result, and it can be thought of as a form of combinatorial rigidity:

THEOREM 3 (Combinatorial rigidity). *Let $S = S_{0,b}$ be a punctured sphere, with $b \geqslant 7$ punctures, and let $\mathcal{C}(S)$ be its curve graph. There exists $K_0 \in \mathbb{N}_{>0}$ such that, if K is a non-trivial multiple of K_0, then the natural map $MCG^{\pm}(S)/DT_K \to Aut(\mathcal{C}(S)/DT_K)$ is an isomorphism.*

In fact, the result applies to other quotients of mapping class groups of punctured spheres, see Theorem 5.11 and the discussion in the outline section. As explained below, the main tools used in the proof of Theorem 3 are the lifting techniques first used in [19] and the finite rigid set of Aramayona-Leiniger [3]. Indeed, we actually obtain a finite rigidity result as well, see Corollary 5.14. An analogue of the theorem also holds for $b = 4$, but in that case the map has finite kernel, see Theorem 5.2.

REMARK 1. We believe that an integer K_0 for which Theorems 1, 2, and 3 hold could in principle be explicitly computed. The starting point for this are explicit constants for various properties of subsurface projections, for example the bounded geodesic image theorem [29]. Then one would need to track all the constants at play in previous papers on the subject, in particular [17, Subsection 1.2.1, Section 3] (which in turn depends on the construction from [10]), [19, Proof of Theorem 2.1], and [8, Notation 8.11]. It would be interesting to have more direct arguments yielding an explicit (and hopefully low) constant, at least for low numbers of punctures in our setting.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

1.1. Open questions

Throughout the paper we only cover the case of punctured spheres/braid groups, so the first questions, which we believe to have positive answer, are the following:

QUESTION 1. *Does Theorem 3 generalise to arbitrary genus?*

QUESTION 2. *Do Theorems 1 and 2 generalise to arbitrary genus?*

It will be clear from the outline section that the key result to generalise is combinatorial rigidity, Theorem 3, which is why we stated the first question separately.

It will be very important for us that the quotients we deal with are hierarchically hyperbolic. Proving hierarchical hyperbolicity of other quotients would be key to proving further quasi-isometric and algebraic rigidity results, and there is a natural class of quotients that contains the quotients we consider in this paper:

QUESTION 3. *Let $g_1, \dots, g_n$ be elements of a mapping class group* MCG*(*S*). Does there exist $K_0 \in \mathbb{N}_{>0}$ such that for all non-trivial multiples K_i of K_0 we have that $MCG(S)/\langle\langle\{g_i^{K_i}\}\rangle\rangle$ is hierarchically hyperbolic?*

In forthcoming work, we will give a positive answer for the case where S is a five-holed sphere. Furthermore, if one allows S to be any finite-type surface, there are certain classes of quotients for which the answer is known to be affirmative, such as those covered in [8], as discussed above, and [5], where quotients by powers of pseudo-Anosovs are considered. Additional quotients covered by the question include quotients of mapping class groups by suitable powers of Dehn twists around non-separating curves only, for instance, as well as the quotients considered in [15].

1.2. Outline of proofs

We will first prove combinatorial rigidity, Theorem 3. The main idea for doing so is the following. Fix a punctured sphere S, say with at least 7 punctures, let $\mathcal{C}$ be its curve graph, and let K be a suitable large integer. While the map $\mathcal{C} \to \mathcal{C}/DT_K$ is not a covering map (as it is quite far from being locally injective), there are still various subgraphs of $\mathcal{C}/DT_K$ that can be lifted. This idea was first used in [19] to show that $\mathcal{C}/DT_K$ is hyperbolic by lifting geodesic triangles, and was developed further in [8] to show that $MCG(S)/DT_K$ is in fact hierarchically hyperbolic. We push these techniques even further to show that the graphs constructed in [3] can be lifted. These are finite rigid sets for $\mathcal{C}$, that is, all isometric embeddings of them into $\mathcal{C}$ differ by an element of the extended mapping class group. Consider now an automorphism $\bar{\phi}$ of $\mathcal{C}/DT_K$ and a copy X of one such graph in $\mathcal{C}/DT_K$. One can consider the diagram below, where the hats denote lifts and π is the projection map.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

$$\begin{array}{ccc} \widehat{X} & \xrightarrow{g} & \widehat{\phi(X)} \\ \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle \pi} \\ X & \xrightarrow{\overline{\phi}} & \phi(X) \end{array}$$

From the diagram, we obtain a candidate element of the extended mapping class group that induces ϕ, namely the element g mapping $\widehat{X}$ to $\widehat{\phi(X)}$. Showing that this candidate is the desired element requires more care, and further lifts, but this is the basic idea. We note that we do not know of a graph that cannot be lifted from $\mathcal{C}/DT_K$ to $\mathcal{C}$, but even a path of length 2 does not have a *unique* DT_K-orbit of lifts, see Remark 4.2, while uniqueness of orbits of lifts is crucial in many arguments.

In order to prove quasi-isometric rigidity we will also need combinatorial rigidity for some subgraphs of the quotient graphs. This is shown in § 6 and 7, where we combine arguments due to Bowditch with further lifting techniques.

Remark 2. The first step towards generalising our results to the case of surfaces with genus, using the same outline, would be to show that some finite rigid sets admit lifts in those cases as well. The finite rigid sets from [3] are more complicated for surfaces with genus, but they can hopefully be used for this purpose.

1.2.1. Quasi-isometric rigidity

To prove quasi-isometric rigidity, we will use a version of [7, Theorem 5.7] which states, roughly, that any quasi-isometry f of a hierarchically hyperbolic group satisfying three assumptions induces an automorphism of a certain graph called the hinge graph. This is a graph that encodes the "standard flats" of the group, as well as their intersection patterns. The three assumptions do not apply to our case, but we show that a similar statement still holds under different assumptions that do apply; this is Theorem 8.7. Afterwards, we use combinatorial arguments to show that an automorphism of the hinge graph induces automorphisms of the graphs from § 6 and 7. Then, by combinatorial rigidity of such graphs, these automorphisms are induced by some element g of the mapping class group, and with some more effort one can show that f and g are at bounded distance.

1.2.2. Algebraic rigidity

Since automorphisms of a group, and more generally isomorphisms between finite index subgroups, induce quasi-isometries, one can expect to obtain strong algebraic rigidity results from quasi-isometric rigidity. We do just that in Theorem 10.1, using results from [2]. Finite normal subgroups could cause the outer automorphism group to be finite rather than trivial, so the key technical results we need is that $MCG^{\pm}(S)/DT_K$ does not contain non-trivial finite normal subgroups, Lemma 10.8.

1.3. Outline of sections

In § 2, we introduce the basic lifting tools that we will use throughout the paper, and we also discuss finite rigid sets. § 3 studies the projection maps from curve graphs to their relevant quotients, and a key result here is that our chosen finite rigid sets map injectively, see Theorem 3.7. In § 4, we show that (graphs isomorphic to our)

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

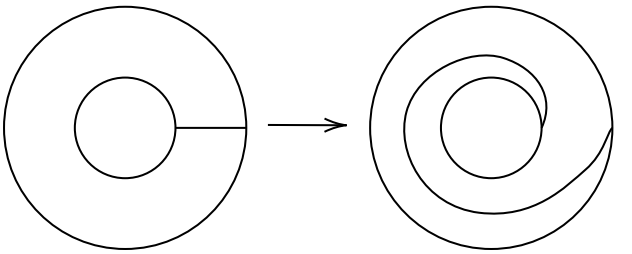


Figure 1. The left twist T of an annulus.

finite rigid sets can be lifted from quotients of curve graphs to curve graphs. We then use this in § 5 to show combinatorial rigidity, Theorem 3, see Theorem 5.11. Further combinatorial rigidity results, that will be needed in § 9, are proved in § 6 and 7.

The goal of § 8 is to show that quasi-isometries of suitable hierarchically hyperbolic spaces induce automorphisms of an associated graph, Theorem 8.7. Then in § 9 we combine Theorem 8.7 and the combinatorial rigidity results from § 6 and 7 to prove Theorem 9.20, which says that quasi-isometries are all at controlled distance from left-multiplications. Then Theorem 1 follows by combining Theorem 9.20 and the general Lemma 9.22. Finally, in § 10, we show algebraic rigidity, Theorem 2, see Theorems 10.1 and 10.9.

2. Setting and main tools

2.1. Mapping classes and curves

Let $S_b = S_{0,b}$ be the sphere with b punctures, on which we fix some orientation. By a *curve* on S_b we will always mean the isotopy class of a simple closed curve which is essential, meaning that it does not bound a disk with less than two punctures. Let $\mathcal{C}_b = \mathcal{C}(S_b)$ be the curve graph of S_b, that is, the simplicial graph whose vertices are curves and two curves are adjacent if and only they admit disjoint representatives. We will often denote S_b by S and $\mathcal{C}_b$ by $\mathcal{C}$ whenever the dependence on b is not relevant. Given two curves α, β, let $i(\alpha, \beta)$ denote their *geometric intersection number*, that is, the minimum number of transverse intersection points between a representative for α and a representative for β.

Let $MCG^{\pm}(S_b)$ be the extended mapping class group of S_b, where we allow orientation-reversing homeomorphisms. We are interested in (powers of) the following type of mapping classes:

Definition 2.1. (Dehn Twist) *Consider the Euclidean plane $\mathbb{R}^2$ with polar coordinates (θ, r), and consider the annulus $A = \{(\theta, r) | 1 \leqslant r \leqslant 2\}$. Choose the orientation on A such that a counterclockwise rotation is positive. Let $T : A \to A$ be the twist map of A given by the formula $T(\theta, r) = (\theta + 2\pi r, r)$. We call T the* left twist *of our annulus. The construction is depicted in Figure 1.*

Now let S *be an oriented surface and let α be a curve in S. Let* N *be a regular neighbourhood of α, which we call an* annulus *with* core curve α*, and choose an orientation-preserving homeomorphism $\phi : A \to N$. The* (left) Dehn twist *about α,*

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

denoted T_α, is the homeomorphism

$$x \to \begin{cases} \phi \circ T \circ \phi^{-1}(x), & \text{if } x \in N, \\ x, & \text{if } x \in S \setminus N. \end{cases}$$

The isotopy class of the Dehn twist T_α does not depend on the choice of N and ϕ, nor on the choice of α within its free isotopy class. Thus, we will often abuse notation slightly and view T_α as a mapping class.

For every $K \in \mathbb{N}_{>0}$, let DT_K be the subgroup generated by the collection $\{T_\alpha^K\}$, where α varies among all curves on S. Furthermore, let $\pi : \mathcal{C} \to \mathcal{C}/DT_K$ be the quotient projection.

DEFINITION 2.2. [26] *Let γ be a curve, and let $p_\gamma : S_\gamma \to S$ be the* cyclic covering *associated to γ. There is a natural compactification of S_γ to a closed annulus $\widehat{S_\gamma}$.*

The annular curve graph *associated to the curve γ is the simplicial graph $\mathcal{C}(\gamma)$ whose vertices are the paths connecting the two boundary components of $\widehat{S_\gamma}$, modulo homotopies that fix the endpoints, and where two such paths are adjacent if they have representatives with disjoint interiors (but may share one or both endpoints).*

If a curve α intersects γ, we denote by $\pi_\gamma(\alpha)$ the *annular projection* of α onto γ, that is, a lift of any representative of α to $\widehat{S_\gamma}$, which is an arc connecting the two boundary components. This is well-defined, as any two lifts of (any two representatives of) α are homotopic relative to the endpoints. If α and β both intersect γ, we shall write $d_\gamma(\alpha, \beta)$ as shorthand for $d_{\mathcal{C}(\gamma)}(\pi_\gamma(\alpha), \pi_\gamma(\beta))$.

2.2. Complexity reduction and lifting properties

In all our arguments, we will only need one property of DT_K (for suitable K), shown in [19] using [17], which we now state.

PROPOSITION 2.3. *There exists $K_0 \in \mathbb{N}$ such that for every $\Theta > 0$ there exists K' such that for all multiples K of K_0 larger than K' the following holds. There exists a well-ordered set $\mathcal{O}$ and a map $\alpha : DT_K \to \mathcal{O}$, which we call* complexity*, such that, for all $x \in \mathcal{C}$ and $g \in DT_K - \{1\}$ there exist $s \in \mathcal{C}$ and some $\gamma_s \in DT_K$, which is a power of the Dehn twist around* s*, such that $\alpha(\gamma_s g) < \alpha(g)$ and one of the following holds:*

- *$i(x, s) = 0$, or*
- *$d_s(x, g(x)) > \Theta$.*

Proof. This is [19, Corollary 3.6] (which applies to DT_K in view of [19, Proposition 5.1]). □

The second case will always be used in conjunction with the bounded geodesic image Theorem:

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

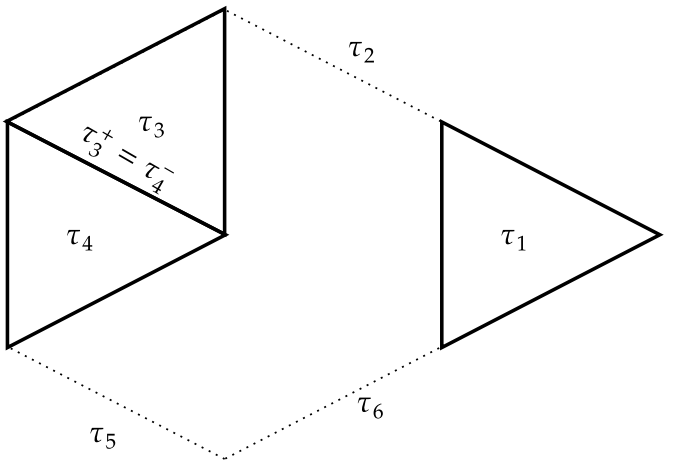


Figure 2. An example of a generalised hexagon. The dotted lines represent geodesics, while the triangles represent simplices.

THEOREM 2.4 [26] *There exists a constant* B *such that the following holds for all finite-type surfaces. For all vertices* $s, x, y \in \mathcal{C}(S)$*, if* $d_s(x, y)$ *is defined and larger than* B*, then any geodesic from* x *to* y *intersects the star of* s.

DEFINITION 2.5. *For short, in the statements below we will say that a normal subgroup* $\mathcal{N} \trianglelefteq MCG(S_b)$ *is* deep enough *if it satisfies the conclusion of Proposition 2.3 for some* Θ *depending on the data that has been fixed up to that point.*

REMARK 2.6. Notice that a deep enough subgroup $\mathcal{N}$ is generated by powers of Dehn twists. This is proved by induction on the complexity $\alpha(g)$ of an element $g \in \mathcal{N}$: if $g = 1$ we have nothing to prove, otherwise there exists $\gamma_s \in \langle T_s \rangle \cap \mathcal{N}$ for some $s \in \mathcal{C}$ such that $g' = \gamma_s g$ has lower complexity, and hence it is a product of powers of Dehn twists by induction hypothesis. Thus $g = \gamma_s^{-1} g'$ is a product of Dehn twists.

The key use of the conclusion of Proposition 2.3 will be to lift a variety of subgraphs from a quotient of the curve graph to the curve graph. This is also done in [19] as well as generalised in [8], and in fact we now state a result from the latter paper. First, a definition from [8]. A *generalised* m-gon in a graph is a sequence $\tau_0, \dots, \tau_{m-1}$ such that:

- Each τ_j is either a simplex, together with non-empty sub-simplices $\tau_j^{\pm}$, or a geodesic in $\mathrm{Lk}(\Delta_j)$ for some (possibly empty) simplex Δ_j with endpoints $\tau_j^{\pm}$.
- $\tau_j^+ = \tau_{j+1}^-$ (indices are taken modulo m).

Note that the second bullet implies that $\tau_j \cap \tau_{j+1}$ is non-empty. See Figure 2 for an example.

The following is a small variation on a result from [8].

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

THEOREM 2.7 *For all $m_0 \geqslant 1$ and all deep enough normal subgroups $\mathcal{N}$ the following hold:*

- *For $b \geqslant 4$, any ordered simplex $\overline{\Delta} \subseteq \mathcal{C}/\mathcal{N}$ admits a unique $\mathcal{N}$-orbit of lifts in $\mathcal{C}$.*
- *For $b \geqslant 4$, given a simplex $\overline{\Delta}$ in $\mathcal{C}/\mathcal{N}$, a lift Δ of $\overline{\Delta}$ in $\mathcal{C}$, and a geodesic γ in the link of $\overline{\Delta}$, we have that γ can be lifted to a geodesic in the link of Δ.*
- *For $b \geqslant 5$ and $m \leqslant m_0$, any generalised m-gon in $\mathcal{C}/\mathcal{N}$ can be lifted in $\mathcal{C}$.*

Proof. This is essentially [8, Proposition 8.31], whose proof only uses [19, Corollary 3.6] and therefore works for any deep enough subgroup. Said proposition is stated for $m_0 = 3\mathrm{Gen}(S) + 2p(S) - 3$, but the only reason for that particular threshold in [8] is that this is the maximal number of sides of generalised m-gons considered in that paper, but the proof works for any fixed m_0. Finally, [8, Proposition 8.31] as stated deals only with the case $b \geqslant 5$, but with similar tools it is easy to prove the uniqueness of the orbit of lifts also for edges and triangles in the Farey complex. This shall be done in Lemmas 2.8 and 2.9, whose proofs are prototypical of many arguments throughout the paper. □

LEMMA 2.8. *Whenever $\mathcal{N}$ is deep enough, every edge $\{\overline{x}, \overline{y}\} \subset \mathcal{C}_4/\mathcal{N}$ admits a unique $\mathcal{N}$-orbit of lifts.*

Proof. By definition of the quotient graph, there exist two adjacent vertices $x, y \in \mathcal{C}_4$ whose projections are $\overline{x}$ and $\overline{y}$, respectively. Therefore, there exists at least one lift. Now let $\{x', y'\}$ be another lift of $\{\overline{x}, \overline{y}\}$. Up to the action of $\mathcal{N}$ we can assume that $y = y'$. Therefore, we have a path $\{x, y, x'\}$ inside $\mathcal{C}_4$, and there is an element $g \in \mathcal{N}$ such that $g(x) = x'$.

If g is the identity we are done, otherwise let (s, γ_s) as in Proposition 2.3. If $d_{\mathcal{C}}(x, s) \leqslant 1$, so that γ_s fixes x, we can apply γ_s to both edges. Now we have a path $\{\gamma_s(x), \gamma_s(y), \gamma_s(x')\}$, and $x = \gamma_s(x)$ is sent to $\gamma_s(x')$ by the element $\gamma_s g$. Since $\alpha(\gamma_s g) < \alpha(g)$ we can proceed by induction on the complexity $\alpha(g)$.

Otherwise $d_s(x, x') > \Theta$. We claim that $d_{\mathcal{C}}(y, s) \leqslant 1$. If this is not the case then $\pi_s(y)$ is well-defined, and by triangle inequality for annular projections either $d_s(x, y) > \Theta/2$ or $d_s(x', y) > \Theta/2$. Without loss of generality, we can assume to be in the first case. If we choose $\Theta \geqslant 2B$, where B is the constant from the bounded geodesic image Theorem 2.4, we get that every geodesic from x to y (that is, the edge between these vertices) should pass through the star of s, which is a contradiction. Therefore, γ_s must fix y, and if we replace x' with $\gamma_s(x')$ we can again reduce the complexity of g, while preserving the fact that the two edges share an endpoint.

At the end of the inductive argument we have that $x = x'$, thus proving uniqueness of the orbit of lifts of an edge. □

LEMMA 2.9. *Whenever $\mathcal{N}$ is deep enough, every triangle $\overline{\Delta} \subset \mathcal{C}_4/\mathcal{N}$ admits a unique $\mathcal{N}$-orbit of lifts.*

Proof. First we show that $\overline{\Delta}$ has a lift. Let $\{x, y, z, x'\}$ a lift of the triangle, which we see as a closed path $\{\overline{x}, \overline{y}, \overline{z}, \overline{x}\}$. Such a lift exists by how edges in the quotient

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

graph are defined, but it is possibly open. Let $g \in \mathcal{N}$ be an element mapping x to x', and suppose that g is not the identity. Then let (s, γ_s) as in Proposition 2.3. If $d(x, s) \leqslant 1$, we can apply γ_s to the whole path, and proceed by induction on the complexity $\alpha(g)$.

Otherwise we claim that either $d_{\mathcal{C}}(s, y) \leqslant 1$ or $d_{\mathcal{C}}(s, z) \leqslant 1$. If none of this happens then $\pi_s(y)$ and $\pi_s(z)$ are both defined, and by triangle inequality we can assume, without loss of generality, that $d_s(x, y) > \Theta/3$. But then again if we choose $\Theta \geqslant 3B$, where B is the constant from the Bounded Geodesic Image Theorem 2.4, we have a contradiction. Hence γ_s fixes a *cut point* between y and z, and we can apply γ_s "beyond" this point. Then we get a new lift of the path, whose endpoints are x and $\gamma_s(x') = \gamma_s g(x)$, and again we can proceed by induction.

At the end of the inductive procedure we will have glued x to x', thus we will have a closed lift of the triangle.

Now let $\Delta = \{x, y, z\}$ and $\Delta' = \{x', y', z'\}$ be two lifts inside $\mathcal{C}_4$ of the same triangle $\overline{\Delta} = \{\overline{x}, \overline{y}, \overline{z}\}$. By Lemma 2.8 we know that the edge $\overline{y}, \overline{z}$ admits a unique $\mathcal{N}$-orbit of lifts, therefore we can assume that $y = y'$ and $z = z'$. Moreover let $g \in \mathcal{N}$ be an element mapping x to x'.

If g is the identity we are done, otherwise let (s, γ_s) as in Proposition 2.3. If $d(x, s) \leqslant 1$ then γ_s fixes x, and we can apply γ_s to both triangles and proceed by induction on the complexity $\alpha(g)$. Otherwise $d_s(x, x') > \Theta$, and with the same argument as in Lemma 2.8 we have that γ_s fixes both y and z. But then if we replace Δ' with $\gamma_s(\Delta')$ we can again reduce the complexity of g, while preserving the fact that the two triangles share an edge.

At the end of the inductive argument we have that $\Delta = \Delta'$, thus proving uniqueness of the orbit of lifts of a triangle. □

Later we will need the following refinement to Theorem 2.7, which shows that, when lifting a quadrilateral (that is, a geodesic quadrangle), two opposite sides can be chosen somewhat "rigidly":

LEMMA 2.10. *Let $\mathcal{N}$ be deep enough with respect to some $\Theta > 0$, let $\overline{Q} \subset \mathcal{C}/\mathcal{N}$ be a quadrilateral with vertices $\overline{v}_1, \overline{w}_1, \overline{v}_2, \overline{w}_2$ and let $Q \subset \mathcal{C}$ be one of its lifts. If the geodesics $[\overline{v}_1, \overline{w}_1]$, $[\overline{v}_2, \overline{w}_2]$ have lifts $[v_i, w_i]$ so that $d_s(v_i, w_i) \leqslant \Theta$ whenever the quantity is defined, then the lifts $[v_i', w_i']$ of $[\overline{v}_i, \overline{w}_i]$ contained in Q is an $\mathcal{N}$-translate of $[v_i, w_i]$.*

Proof. This is the "moreover" part of [19, Proposition 4.3], whose proof only uses [19, Corollary 3.6]. □

2.3. Special intersections

For the following definitions, let Γ be either $\mathcal{C}$ or $\mathcal{C}/\mathcal{N}$.

DEFINITION 2.11. *A* chain *is a collection of vertices $v_1, \dots, v_k \in \Gamma$ such that $d_\Gamma(v_i, v_j) > 1$ iff $|i - j| = 1$. A chain is closed if the same holds with indices mod* k.

In the curve graph, a chain is just a sequence of curves such that two of them intersect iff they are consecutive.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

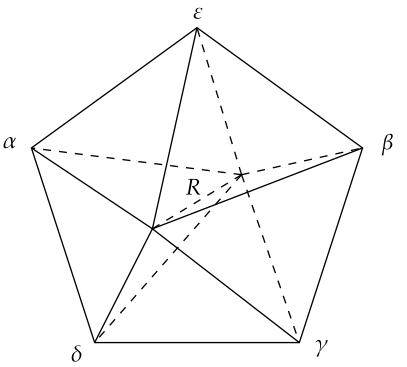


Figure 3. The generalised pentagon $\mathfrak{P}$ described in 2.12, which detects that any two non-adjacent vertices in the link of R have special intersection.

The following should be compared with the notion of $\mathcal{X}$-detectable intersection from [3]:

DEFINITION 2.12. (Special intersection) *Let $b \geqslant 5$. Given a facet $P \subseteq \Gamma$ (that is, a simplex of codimension one in a maximal simplex) we say that the vertex $\alpha \in \Gamma$ has* special intersection *with the vertex $\beta \in \Gamma$ with respect to* P *if:*

- *α and β both complete P to maximal simplices;*
- *there exist γ, δ such that γ, α, β, δ is a chain;*
- *there exists $\varepsilon \in P$ such that $Lk(\gamma) \cap P = Lk(\delta) \cap P = P - \{\varepsilon\}$.*

We say that γ and δ are the auxiliary vertices that detect the intersection of α and β. The situation is illustrated in Figure 4.

REMARK 2.13. If α and β are in $\mathcal{C}$ and have special intersection then they have intersection number 2 and lie in the four-punctured sphere that P cuts out (see the remark following [3, Lemma 2.5]).

REMARK 2.14. Definition 2.12 actually describes an isometrically embedded generalised pentagon $\mathfrak{P}$, as in Figure 3. More precisely, if we define R as the codimension 2 simplex such that $P = \varepsilon \star R$, then γ, α, β, δ, ε form a closed chain in the link of R; this becomes an isometrically embedded pentagon when the vertices are ordered as α, ε, β, γ, δ. Note that these five vertices must play symmetric roles, meaning that when two of them are not adjacent in Γ they have special intersection, and the special intersection is detected by the others. For example, γ and ε have special intersection with respect to the facet $R \star \beta$, and α and δ detect it. Notice that having special intersection is a completely combinatorial property, and hence it is preserved by graph automorphisms.

DEFINITION 2.15. *We will call an isometrically embedded generalised pentagon $\mathfrak{P} \subset \Gamma$ which is isomorphic to the one in Figure 3 a* special *pentagon.*

Another consequence of stating special intersection in combinatorial terms is the following lifting property:

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

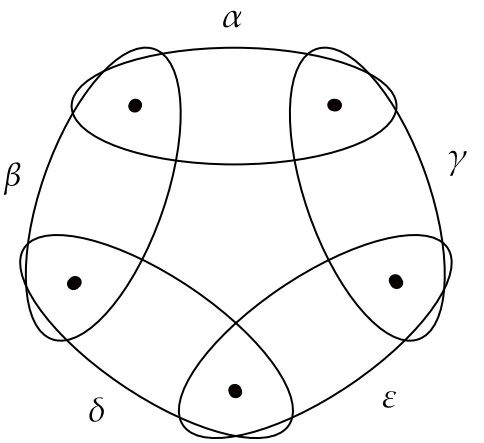


Figure 4. The five curves involved in Definition 2.12 for Γ=C, forming a chain on the five-holed sphere that R cuts out. Any intersection is special, therefore the intersection number is always 2.

LEMMA 2.16. *For every $b \geqslant 5$ and for every $\mathcal{N}$ deep enough, if $\overline{\alpha}, \overline{\beta} \in \mathcal{C}/\mathcal{N}$ have special intersection then they admit lifts α, β with special intersection.*

Proof. By Theorem 2.7, whenever $\mathcal{N}$ is deep enough every special pentagon $\overline{\mathfrak{P}}$ as in Definition 2.15 admits a lift $\mathfrak{P}$, which remains isometrically embedded since the projection map is 1-Lipschitz. More precisely, if $\overline{x}, \overline{y} \in \overline{\Gamma}$ and $x, y \in \Gamma$ are their lift then

$$d_{\mathcal{C}^1}(x, y) \geqslant d_{\mathcal{C}^1/DT_K}(\overline{x}, \overline{y}) = d_{\overline{\mathfrak{P}}}(\overline{x}, \overline{y}) = d_{\mathfrak{P}}(x, y).$$

□

For the rest of the paper, given a curve $\beta \in \mathcal{C}$ which bounds a twice-punctured disk, let H_β be the *half Dehn twist* around β, with respect to the given orientation of S_b. If $b \geqslant 5$ there is only one twice-punctured disk bounded by β, while if $b = 4$ for each disk D bounded by β we will refer to the half Dehn twist of D as H_D. The following lemma describes the set of curves with special intersection with β with respect to some facet P:

LEMMA 2.17. *Let $b \geqslant 4$ and let $\alpha, \alpha' \in \mathcal{C}$ be two curves which both have special intersection with the same β, with respect to the same P. Let S_4 be the four-holed sphere that P cuts out, and let $D \subset S_4$ be any of the two disks bounded by β inside* S_4. *Then there exists an integer $k \in \mathbb{Z}$ such that $\alpha' = H_D^k(\alpha)$.*

Proof. P cuts out a sphere S_4 with four punctures, and by Remark 2.13 having special intersection implies that the intersection number is 2. Thus we want to show that, if α, α' are both adjacent to β in the Farey complex and D is one of the two disks bounded by β, then $\alpha' = H_D^k(\alpha)$ for some $k \in \mathbb{Z}$. By properties of the Farey complex there exists a sequence of triangles $T_1, \dots, T_k$ such that each triangle contains β, $\alpha \in T_1$, $\alpha' \in T_k$ and every two consecutive triangles T_i and T_{i-1} share and edge containing β. Thus it suffices to prove that if α, α', β are the

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

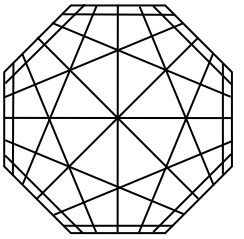


Figure 5. Doubling these arcs gives a copy of X_8.

vertices of a triangle then $\alpha' = H_D^{\pm 1}(\alpha)$. This is true, since the two curves $H_D^{\pm 1}(\alpha)$ are adjacent to both α and β, and α and β belong to exactly two triangles. □

REMARK 2.18. When $b \geqslant 5$ we have to be a little cautious. If β already bounds a twice-punctured disk D inside S_b then this disk embeds inside S_4, since it cannot contain any curve of P. Hence the map H_D, defined on S_4, extends to the usual half Dehn Twist along β, which is defined on the whole S. Therefore, Lemma 2.17 shows that, whenever α, α' are two curves with special intersection with β with respect to the same P, there exists $k \in \mathbb{Z}$ such that $\alpha' = H_\beta^k(\alpha)$.

However, if β does not bound a twice-punctured disk inside S_b there might be no way to extend H_D to the whole surface S, so it is important to underline that the conclusion holds *inside the four-holed sphere that P cuts out.*

2.4. The finite rigid set

For the rest of the section let $b \geqslant 5$. We denote by X_b the full subgraph of $\mathcal{C}$ defined in [3, Section 3]. More precisely, we represent S_b as the double of a regular b-gon with vertices removed. An arc connecting non-adjacent sides of the b-gon doubles to a curve on S_b, and let X_b be the subgraph spanned by such curves, as in Figure 5 for the case $b = 8$.

A *copy* of X_b will be the image of an isometric embedding $X_b \to \mathcal{C}$ (respectively $\mathcal{C}/\mathcal{N}$). The reason we are interested in X_b is the following theorem, which is a somewhat refined version of Ivanov's and was proven in [3, Theorem 3.1]:

THEOREM 2.19. *For every $b \geqslant 5$, any locally injective simplicial map $\phi : X_b \to \mathcal{C}$ is induced by a mapping class $h \in MCG^{\pm}$, meaning $\phi = h|_{X_b}$. Moreover, any two such* h *differ by an element of the pointwise stabiliser $Pstab(X_b)$, generated by the reflection* r *that swaps the two copies of the* b*-gon.*

Another important fact about X_b is that every intersection is special, regardless of the ambient graph Γ because it can be detected using only vertices that belong to X_b (this fact should be compared with [3, Lemma 3.2]). For example, the special intersection between α and β in Figure 6 is detected by the special pentagon spanned by α, β, the simplex $R \subset X_b$ of codimension two and the three curves $x, y, z \in X_b$.

The next lemmas show how to construct some copy of X_b starting from a given copy of X_{b-1}, and describe how much freedom one has to do so.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

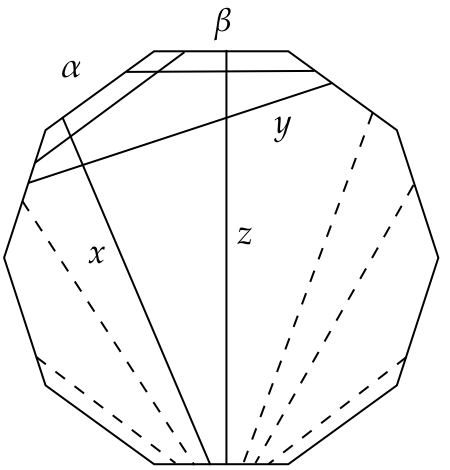


Figure 6. The dashed lines represent the codimension two simplex R. Any two intersecting curves of this "star" have special intersection, with respect to some facet that extends R.

DEFINITION 2.20. *If $b \geqslant 6$, we say that a vertex β in a copy of* $\mathrm{X_b}$ *is* minimal *if* $Lk_{X_b}(\beta) \cong X_{b-1}$.

LEMMA 2.21. (Existence of an extension of X_{b-1}) *Let $b \geqslant 6$. Every copy of X_{b-1} inside $\mathcal{C}$ which is in the link of a vertex β may be completed to a copy of* $\mathrm{X_b}$ *that contains β. Moreover, if α has special intersection with β, with respect to some facet $P \subseteq X_{b-1}$, then we can choose* $\mathrm{X_b}$ *to contain α.*

Proof. First we show that β bounds a twice-punctured disk. To see this, notice that, by construction of the graph X_{b-1}, if two curves $x, y \in X_{b-1}$ are at distance 1 then there exists a curve $z \in X_{b-1}$ which is at distance 2 from both. This fact still holds if we replace distances in X_{b-1} with distances in the curve graph, since X_{b-1} is isometrically embedded, and it translates to the fact that whenever two curves in X_{b-1} are disjoint there is a third curve which intersects both. Hence the union of all curves of X_{b-1} must lie on the same connected component of $S_b \setminus \beta$, which is a disk that we call S'. But then S' contains $b-4$ pairwise disjoint curves (other than the boundary curve), and therefore must contain at least $b-2$ punctures. Thus the other connected component of $S_b \setminus \beta$ must be a disk with at most 2 punctures, and since β is an essential curve it cannot bound disks with less than two punctures.

Now, since β bounds a twice-punctured disk it is easy to find some X'_b which contains β as a minimal curve. Let $X'_{b-1} = \mathrm{Lk}_{X'_b}(\beta)$, and consider the S_{b-1} obtained by shrinking β to a puncture B. By finite rigidity of X_{b-1} there exists a mapping class $f \in MCG(S_{b-1})$ mapping X'_{b-1} to X_{b-1}, which we may choose to be the identity in a neighbourhood of B, up to rotations of S_{b-1} and isotopy. Therefore, f extends to a mapping class $F \in MCG(S_b)$ (for example, by setting F to be the identity in a neighbourhood of the twice-punctured disk surrounded by β), and now $F(X'_b)$ is a copy of X_b that completes $\beta \star X_{b-1}$.

For the "moreover" part, choose a copy of X_b that completes $\beta \star X_{b-1}$ and let α' be the curve in this copy that corresponds to α (i.e., the curve that has special

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

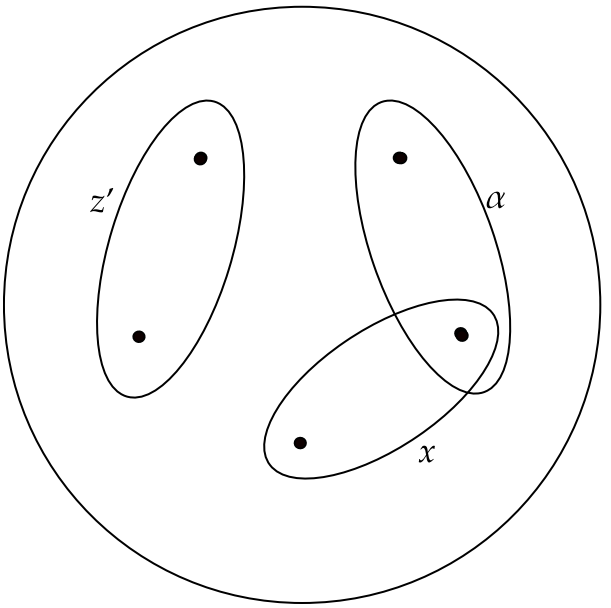


Figure 7. The five-punctured sphere that R cuts out.

intersection with β with respect to the same P). By Lemma 2.17 and the discussion in Remark 2.18 there is a suitable power H_β^k of the half twist around β that maps α to α'. Since H_β fixes $\beta \star X_{b-1}$ we may apply this mapping class to obtain the desired copy of X_b. □

LEMMA 2.22. (Uniqueness of the extension) *Let $b \geqslant 6$. In any copy of* $\mathrm{X_b}$ *inside $\mathcal{C}$, let X_{b-1} be the link of some minimal vertex β and let $\alpha \in X_b$ be a curve intersecting β. Then for every other $z \in X_b$ that intersects β there is a facet $P \subset X_{b-1}$ such that* z *is the unique curve in $\mathcal{C}$ that belongs to $Lk(P) \cap Lk(\alpha)$. In other words, a copy of* $\mathrm{X_b}$ *is uniquely determined by the star of a minimal vertex β and one of the curves that intersect β.*

Proof. Let $R \subset X_{b-1}$ be a simplex of codimension two such that $\alpha, z \in \mathrm{Lk}(R)$, and complete it to a codimension one simplex P by adding some curve $x \in \mathrm{Lk}(\alpha_i) \cap X_{b-1}$ that intersects α. For example, x and R can be chosen as in Figure 6. We need to show that if some curve z' lies in $\mathrm{Lk}(P) \cap \mathrm{Lk}(\alpha)$ then it is unique. First notice that z' must be an essential curve of the five-punctured sphere S_5 that R cuts out (that is, it cannot surround only one puncture of S_5, or it would coincide with some curve in R). Moreover, we know that α and x are both essential curves of S_5 and they have intersection number 2, so we are in the situation depicted in Figure 7. Now, α and x together cut S_5 in three once-punctured disk and a twice-punctured disk. Since z' is essential in S_5 it must be the boundary of the twice-punctured disk, and thus it is unique. □

REMARK 2.23. When $b = 5$ we cannot define minimal curves as in Definition 2.20. Thus, with a slight abuse of notation, we will say that every curve in X_5 is minimal,

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

because it bounds a twice-punctured disk. Moreover Lemma 2.21 is false: $\mathrm{Lk}_{X_5}(\beta)$ is made of two points, but if we take two vertices $\gamma, \gamma' \in \mathrm{Lk}_{\mathcal{C}_5}(\beta)$ with intersection number greater than 2 we cannot complete them to a copy of X_5. However, we can run the proof of Lemma 2.22, with $R = \emptyset$, to get the following:

LEMMA 2.24. *For* $\mathrm{b} = 5$*, if two copies* $X_5, X_5' \subset \mathcal{C}_5$ *share four curves then they must coincide.*

3. Projection lemmas

In this section, we gather various results about the nature of the projection map $\pi : \mathcal{C} \to \mathcal{C}/\mathcal{N}$. We start by showing that the restriction of π to the curve graph of a subsurface is what one would expect.

LEMMA 3.1. *Let* U *be a subsurface of* S*. Let* $\mathcal{N}$ *be a deep enough subgroup, and let* $\mathcal{N}(U)$ *be the subgroup of* $\mathcal{N}$ *generated by the elements with support in* U*. If* x, x' *are curves on* U *and there exists an element* $g \in \mathcal{N}$ *mapping* x *to* x′ *then we may find another element* $h \in \mathcal{N}(U)$ *mapping* x *to* x′*.*

Proof. Let Δ be a pants decomposition of the complement of U, including its boundary curves. We proceed by induction on the complexity of g. If g is the identity then we are done; otherwise let (s, γ_s) be as in Proposition 2.3.

If $d_{\mathcal{C}}(x, s) \leqslant 1$ then we can apply γ_s to both curves, in order to reduce the complexity of g. Then by inductive hypothesis we can find some $h \in \mathcal{N}(\gamma_s(U))$ mapping $\gamma_s(x)$ to $\gamma_s(x')$, and therefore $\gamma_s^{-1} h \gamma_s$ maps x to x' and belongs to $\mathcal{N}(U)$ by properties of Dehn twists under conjugation.

Otherwise $d_s(x, x') > \Theta$. Arguing exactly as in the proof of Lemma 2.8 we get that γ_s fixes Δ pointwise whenever $\mathcal{N}$ is deep enough. Now recall that, by construction, $\gamma_s \in \langle T_s^N \rangle$, where T_s is the Dehn twist around s. Then s must be disjoint from all curves in Δ, which means that either $s \in \Delta$ (which is impossible, since otherwise γ_s would fix x) or $s \subseteq U$. Therefore, $\gamma_s \in \mathcal{N}(U)$, and if we apply it to x' we can conclude by induction. □

COROLLARY 3.2. *Let* U *be a connected subsurface of* S *of complexity at least* 2*, and let* $\mathcal{N}$ *be a deep enough subgroup. If* $\pi : \mathcal{C}(S) \to \mathcal{C}(S)/\mathcal{N}$ *is the quotient projection then* $\pi(\mathcal{C}(U)) \cong \mathcal{C}(U)/\mathcal{N}(U)$*.*

Proof. By Lemma 3.1 these graphs have the same vertices, and since U has complexity at least 2 the adjacency relation is having disjoint lifts in both cases. □

3.1. Isometric projections

We move on to show that, whenever $\mathcal{N}$ is deep enough, the projection is an isometry on a variety of subgraphs of the curve graph. All the result of this subsection rely on the following refinement of [19, Lemma 4.4], which roughly speaking says that the projection $\pi : \mathcal{C} \to \mathcal{C}/\mathcal{N}$ preserves directions with small projections, in the link of some fixed simplex:

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

LEMMA 3.3. *Fix $\Theta > 0$. Let $\mathcal{N}$ be deep enough and let $\overline{\Delta} \subseteq \mathcal{C}/\mathcal{N}$ be a simplex and $\Delta \subseteq X$ be one of its lifts. Suppose that $x, y \in Lk(\Delta)$ project to $\overline{x}, \overline{y} \in Lk(\overline{\Delta})$ and have the property that $d_s(x, y) < \Theta$ for every $s \in \mathcal{C}$ for which the quantity is defined. Then $\pi|_{[x,y]}$ is an isometric embedding into $Lk(\overline{\Delta})$, for any geodesic $[x, y] \subseteq Lk(\Delta)$.*

Proof. The proof is very similar to that of [19, Lemma 4.4]. Suppose by contradiction that there is a shorter path from $\overline{x}$ to $\overline{y}$ inside $\mathrm{Lk}(\overline{\Delta})$. Lift this shorter path as a geodesic segment $[y, x'] \subseteq \mathrm{Lk}(\Delta)$, which can be done by Lemma 2.7. There exists $\gamma \in \mathcal{N}$ such that $\gamma x = x'$. We proceed by induction on the complexity of γ, to prove that $d_{\mathrm{Lk}(\Delta)}(y, x') = d_{\mathrm{Lk}(\Delta)}(y, x)$ (thus contradicting that $d_{\mathrm{Lk}(\overline{\Delta})}(\overline{x}, \overline{y}) < d_{\mathrm{Lk}(\Delta)}(x, y)$). If $\gamma = 1$ we are done, otherwise let (s, γ_s) be as in 2.3. If $d_{\mathcal{C}}(x, s) \leqslant 1$ then $\gamma_s x = x$, and we can apply γ_s to both geodesics and to Δ and conclude by the induction hypothesis. Otherwise $d_s(x, \gamma x) > \Theta$, and arguing as in Lemma 2.8 we see that γ_s must fix Δ pointwise. Moreover, either $d_{\mathcal{C}}(y, s) \leqslant 1$ or $d_s(y, x') \geqslant d_s(x, x') - d_s(x, y)$ must be large, since $d_s(x, y)$ is assumed to be small. In both cases there must be some $s' \in [y, x']$ in the star of s. Thus, one can change the lift of $[y, x]$ as $[y, \gamma_s x']$, while keeping it an isometric lift inside the link of Δ. One concludes by induction hypothesis, which applies to $\gamma_s \gamma$. □

The previous lemma is particularly useful in the following form:

COROLLARY 3.4. *For every finite set of vertices $V \subseteq \mathcal{C}$ there exists Θ such that, whenever $\mathcal{N}$ is deep enough with respect to Θ, the projection is an isometry on* V.

Proof. Since V is finite, we only need to show that for every $x, z \in \mathcal{C}$ there is some constant $M(x, z)$ such that $\sup_{s \in \mathcal{C}} d_s(x, z) < M(x, z)$, because then we can choose $\Theta > M = \max_{x,z \in V} M(x, z)$ to ensure that the hypothesis of Lemma 3.3 applies. One way to see this is to complete x, z to complete clean markings μ, ν, in the sense of [26, Section 2.5]. By the Distance formula [26, Theorem 6.12] there exists a constant $M'(S)$ such that the sum $\sum_{s \in \mathcal{C},\, d_s(\mu,\nu) > M'} d_s(\mu, \nu)$ is bounded above in terms of the distance between μ and ν in the marking graph. In particular the sum is finite, and since every term is greater than a constant there must be a finite number of terms. Moreover, for every $s \in \mathcal{C}$ for which the quantity $d_s(x, y)$ is defined (in particular, not a base curve of μ or ν), let $x' \in \mu$ and $y' \in \nu$ be curves which realise $d_s(\mu, \nu)$. Then by triangle inequality

$$d_s(x, y) \leqslant d_s(x, x') + d_s(x', y') + d_s(y', y) = d_s(x, x') + d_s(\mu, \nu) + d_s(y', y)$$

and since annular projections are 2-Lipschitz (see [26, Lemma 2.3]) we have that

$$d_s(x, y) \leqslant d_s(\mu, \nu) + 4$$

Thus it suffices to choose $M(x, z) > \max_{s \in \mathcal{C},\, d_s(\mu,\nu) > M'} d_s(\mu, \nu) + 4$. □

Now we want to apply Corollary 3.4 to the rigid set.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

DEFINITION 3.5. *Given a copy of $X_b \subset \mathcal{C}$ for $b \geqslant 5$, set*

$$Y_b = X_b \bigcup_{\beta \in X_b \text{ minimal}} H_\beta^{\pm 1}(X_b).$$

LEMMA 3.6. *For every $b \geqslant 5$ there exists $M \in \mathbb{R}^+$ with the following property. Given a copy $X_b \subset \mathcal{C}$ set $\mathrm{Y_b}$ as in Definition 3.5. Then for every $x, z \in Y_b$ and every $s \in \mathcal{C}$ we have that $d_s(x, z) \leqslant M$.*

Proof. Fix a copy X_b and let $M = \max_{x,z \in Y_b} M(x, z)$, where $M(x, z)$ is defined as in Corollary 3.4. We are left to prove that, if X_b' is another copy of the rigid set and Y_b' is the corresponding union of half twists, then M works also for Y_b'. Take an extended mapping class f that maps X_b to X_b', which exists by Theorem 2.19. In particular f maps minimal curves to minimal curves. Therefore,

$$\begin{aligned} Y_b' &= X_b' \bigcup_{\beta' \in X_b' \text{ minimal}} H_{\beta'}^{\pm 1}(X_b') = f(X_b) \bigcup_{\beta \in X_b \text{ minimal}} H_{f(\beta)}^{\pm 1} \circ f(X_b) \\ &= f(X_b) \bigcup_{\beta \in X_b \text{ minimal}} f \circ H_\beta^{\pm 1}(X_b) = f(Y_b) \end{aligned}$$

Now the conclusion follows, since for every $x, z \in Y_b'$ and $s \in \mathcal{C}$ we have that

$$d_s(x, z) = d_{f^{-1}(s)}(f^{-1}(x), f^{-1}(z)) \leqslant M.$$

□

This corollary is crucial for us, so we state it as a theorem:

THEOREM 3.7. *For every $b \geqslant 5$ the following holds whenever $\mathcal{N}$ is deep enough. Let $X_b \subset \mathcal{C}$ be a copy of the rigid set and let $\mathrm{Y_b}$ be as in Definition 3.5. Then the restriction $\pi|_{Y_b}$ is an isometry. In particular there exists an isometrically embedded copy of $\mathrm{X_b}$ inside $\mathcal{C}/\mathcal{N}$.*

Proof. Just combine Lemma 3.6 and Corollary 3.4. □

There are other two consequences of Lemma 3.3 that will be quite useful later. The first one deals with the cardinality of the quotient:

COROLLARY 3.8. *Whenever $\mathcal{N}$ is deep enough, the quotient $\mathcal{C}/\mathcal{N}$ is infinite.*

Proof. For every curve $x \in \mathcal{C}$ and every pseudo-Anosov mapping class $g \in MCG(S)$ there exists $\Theta > 0$ such that $\sup_{s \in \mathcal{C}, n \in \mathbb{Z}} d_s(x, g^n(x)) < \Theta$. This follows from an argument in the proof of [19, Theorem 2.1] which just uses the Bounded geodesic image Theorem 2.4 and the fact that for every $x, z \in \mathcal{C}$ then $\sup_{s \in \mathcal{C}} d_s(x, z) < \infty$, which we proved in Corollary 3.4. Hence, if $\mathcal{N}$ is deep enough with respect to Θ, by Lemma 3.3 the projection is an isometry on the orbit $\{g^n(x)\}_{n \in \mathbb{Z}}$, which is infinite since g is pseudo-Anosov. □

The second consequence deals with filling curves:

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

COROLLARY 3.9. *Fix $\Theta > 0$. Let $\mathcal{N}$ be deep enough and let $\overline{\Delta} \subseteq \mathcal{C}/\mathcal{N}$ be a simplex and $\Delta \subseteq X$ be one of its lift. Suppose that Δ cuts out a single subsurface Σ of complexity at least 1, and let $x, y \subset \Sigma$ be a pair of filling curves with the property that $d_s(x, y) < \Theta$ whenever the quantity is defined. If x, y project to $\overline{x}, \overline{y} \in Lk(\overline{\Delta})$ then any two lifts $x', y' \in Lk(\Delta)$ again fill Σ.*

Proof. Recall that x and y fill Σ if and only if $d_{\mathrm{Lk}(\Delta)}(x, y) \geqslant 3$. If x and y satisfy the hypothesis of Lemma 3.3 then the distance between $\overline{x}$ and $\overline{y}$ in $\mathrm{Lk}(\overline{\Delta})$ remains at least three. But then every pair of lifts must be at distance at least three, since the projection map is 1-Lipschitz. □

3.2. Puncture separations

Similarly to the previous subsection, we aim to establish some results of isometric projection, but this time they will arise from some observations on how curves separate punctures on a sphere. Choose an enumeration of the punctures of S_b. A curve α on S_b is separating, therefore it splits the punctures into two sets $\mathcal{B}^+(\alpha)$ and $\mathcal{B}^-(\alpha)$.

DEFINITION 3.10. *The* puncture separation *induced by α is the unordered couple $\{\mathcal{B}^+(\alpha), \mathcal{B}^-(\alpha)\}$.*

REMARK 3.11. Notice that a Dehn twist fixes each puncture, therefore preserves puncture separations. In fact if two punctures are on the same side of α they may be joined by an arc which is disjoint from α, and the image of this arc via a Dehn twist will again join the same punctures while being disjoint from the image of α. In other words, any two curves that induce different puncture separations cannot be identified in a quotient of the curve graph by a subgroup generated by powers of Dehn twists, such as all deep enough subgroups $\mathcal{N}$ in the light of Remark 2.6.

DEFINITION 3.12. *Two curves α and α' induce* nested *puncture separations if $\mathcal{B}^+(\alpha) \subseteq \mathcal{B}^{\pm}(\alpha')$. In other words, the splitting induced by $\alpha \cup \alpha'$ refines the splitting induced by α.*

Notice that two disjoint curves induce nested puncture separations, but the converse is not true (for example, choose some β intersecting α and set $\alpha' = T_\beta(\alpha)$, which induces the same puncture separation by Remark 3.11). The following lemma shows a peculiar behaviour of the projection $\mathcal{C} \to \mathcal{C}/DT_K$ which holds for all $K \in \mathbb{N}$:

LEMMA 3.13. *For every subgroup $H \leqslant MCG(S)$ generated by powers of Dehn Twists, the following hold for the projection $\mathcal{C} \to \mathcal{C}/H$:*

1. *For $\mathrm{b} = 4$, if α and α' are adjacent in the Farey complex then their projections remain at distance 1.*
2. *For $b \geqslant 5$, if α and α' are disjoint curves then their projections remain at distance 1.*
3. *For $b \geqslant 5$, if α and α' intersect and their puncture separations are not nested then their projections remain at distance at least 2.*

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

Proof. In all cases we will use that the projection is 1-Lipschitz, since the action of H over $\mathcal{C}$ is simplicial.

(1) Two adjacent curves in the Farey complex induce different puncture separations, therefore their projections must be at distance 1 since they cannot coincide.
(2) If α and α' are disjoint then they must induce different puncture separations, otherwise they would bound an unpunctured annulus and therefore they would be isotopic. Then their projections can only be at distance 1, for the same reasons as before.
(3) Suppose that α and α' intersect and induce non-nested puncture separations. Firstly, α and α' cannot have the same projection, since they induce different puncture separations. Furthermore, if by contradiction $\pi(\alpha)$ and $\pi(\alpha')$ are the endpoints of an edge then we may lift it, since the action is simplicial. In other words there exists some $g \in H$ such that α is disjoint from $g(\alpha')$. But then α and $g(\alpha')$ must induce nested puncture separations, and so must do α and α' since Dehn twists preserve puncture separations by Remark 3.11.

□

As an immediate consequence of Items (1) and (2) we get:

COROLLARY 3.14. *For every $b \geqslant 4$, the projection map is an isometry on every simplex $\Delta \subset \mathcal{C}$.*

REMARK 3.15. Item (1) directly implies Theorem 3.7 for $b = 4$. Moreover, Items (2) and (3) prove Theorem 3.7 for $b \geqslant 5$, since Y_b has diameter 2 and by construction each pair of intersecting curves induce non-nested puncture separations. This alternative proof works for every choice of K, though only if the surface is a punctured sphere.

4. Lifting the rigid set

The aim of this section is to prove the following lifting result:

THEOREM 4.1. *For every $b \geqslant 5$ and every $\mathcal{N}$ deep enough, every copy of X_b in $\mathcal{C}/\mathcal{N}$ has a lift inside $\mathcal{C}$, and any two lifts are conjugated by an element of $\mathcal{N}$.*

REMARK 4.2. At the moment, it is not known to the authors whether there exists a finite graph $\overline{G}$ such that, whenever $\overline{G}$ is immersed in some $\mathcal{C}/DT_K$, then $\overline{G}$ does not admit a lift in $\mathcal{C}$. However there are easy examples where a graph $\overline{G}$ admits two lifts G, G' which belong to different DT_K-orbits; for applications it is often crucial for us that there is only one orbit of lifts of specific graphs. For example, for every $K > 1$ let $\bar{l} = \{\overline{\alpha}, \overline{\beta}, \overline{\gamma}\}$ be a geodesic segment of length 2 in $\mathcal{C}_4/DT_K$, and let $l = \{\alpha, \beta, \gamma\}$ be one of its lifts (whose existence is granted by Lemma 2.7). Now let $\gamma' = T_\beta^K(\gamma)$, so that $l' = \{\alpha, \beta, \gamma'\}$ is another isometric lift of $\bar{l}$. If by contradiction $g \in DT_K \setminus \{1\}$ maps l to l' then it must fix α and β. Notice moreover that, if

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

an orientation preserving mapping class fixes an edge of the Farey complex then it fixes the whole complex. Thus g is a hyperelliptic involution, and since it must permute the punctures it cannot belong to any group generated by powers of Dehn twists, such as DT_K.

4.1. Existence of a lift

Proof of existence. We proceed by induction, using the case $b = 5$ as the base case since X_5 is a pentagon, which lifts by Theorem 2.7. Assume that $b \geqslant 6$ and that every copy of X_{b-1} has a lift inside $\mathcal{C}_{b-1}$, and let $\overline{X}_b$ be a copy of X_b inside $\mathcal{C}_b/\mathcal{N}$. Let $\overline{\beta} \in \overline{X}_b$ be a minimal vertex, and let $\overline{X}_{b-1} := \mathrm{Lk}_{\overline{X}_b}(\overline{\beta})$. In order to apply the inductive hypothesis we need to show that $\mathrm{Lk}_{\mathcal{C}_b}(\beta)$ is isomorphic to $\mathcal{C}_{b-1}$ and that its projection is isomorphic to $\mathcal{C}_{b-1}/\mathcal{N}(S_{b-1})$. We start with the first assertion.

LEMMA 4.3. *Any lift β of a minimal vertex $\overline{\beta} \in \overline{X}_b$ bounds a twice-punctured disk. In other words $Lk_{\mathcal{C}_b}(\beta) \cong \mathcal{C}_{b-1}$.*

Proof. Let $\overline{P}$ be a facet inside $\overline{X}_{b-1}$. Since $\overline{\beta} \star \overline{P}$ is a simplex we can find a lift $\beta' \star P$ by Lemma 2.7, and up to applying some elements of $\mathcal{N}$ we can assume that $\beta' = \beta$ and therefore $P \in \mathrm{Lk}_{\mathcal{C}_b}(\beta)$. We want to show that all curves in $\mathrm{Lk}_{\mathcal{C}_b}(\beta)$ lie in the same subsurface of $S \setminus \beta$, or in other words that, given any two curves $\gamma, \delta \in \mathrm{Lk}_{\mathcal{C}_b}(\beta)$, there is a chain $\gamma = \eta_0, \eta_1, \ldots, \eta_k = \delta \in \mathrm{Lk}_{\mathcal{C}_b}(\beta)$. Since every curve in $\mathrm{Lk}_{C_b}(\beta)$ either coincides with or intersects a curve in P, which has maximal dimension in $\mathrm{Lk}_{\mathcal{C}_b}(\beta)$, it is enough to show that for every γ, $\delta \in P$ there is a curve η that intersects both. Now, since $\overline{X}_{b-1}$ is isometrically embedded inside $\mathcal{C}_b/\mathcal{N}$ we may pick some $\overline{\eta} \in \overline{X}_{b-1}$ at distance 2 from both $\overline{\gamma}$ and $\overline{\delta}$, and let η be one of its lifts in $\mathrm{Lk}_{\mathcal{C}_b}(\beta)$. Since the projection map is 1-Lipschitz, η must be at distance at least 2 from (hence intersect) both γ and δ, and we are done. □

Now, Corollary 3.2 shows that the projection of $\mathrm{Lk}_{\mathcal{C}_b}(\beta)$ is isomorphic to $\mathcal{C}_{b-1}/\mathcal{N}(S_{b-1})$. Therefore, we can use the inductive hypothesis to lift $\overline{X}_{b-1}$ to some X_{b-1} inside $\mathrm{Lk}_{\mathcal{C}_b}(\beta)$. Now the key point of the following construction will be this uniqueness lemma:

LEMMA 4.4. Unique lift for "very" special intersections *For $b \geqslant 5$ the following holds whenever $\mathcal{N}$ is deep enough. Let $x, x' \in \mathcal{C}$ be two curves with special intersection with the same y, with respect to the same facet P, and let S be any graph in the link of y. Suppose that x, y, S can be completed to a copy of $\mathrm{X_b}$. Then if* x *and* x′ *project to the same element in the quotient there is some $g \in \mathcal{N}$ which maps x to x' and fixes y and S.*

Proof. Since both x and x' have special intersection with y, with respect to the same facet P, by Lemma 2.17 there exists $k \in \mathbb{Z}$ such that, in the four-holed sphere S_4 that P cuts out, $x' = H_D(x)$, where D is one of the disks bounded by y. Now we claim that $k = 2m$ must be even, and therefore $H_D^{2m} = T_y^m$ extends to the whole surface S. To see this first notice that since x and x' are in the same $\mathcal{N}$-orbit they must induce the same puncture separation by Remark 3.11. Moreover x and y intersect twice, thus they separate the punctures of S into four

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

sets $\mathcal{A}^{\pm}, \mathcal{B}^{\pm}$, each of which corresponds to one of the punctures of S_4. Suppose that x induces the separation $\{\mathcal{A}^+ \cup \mathcal{B}^+, \mathcal{A}^- \cup \mathcal{B}^-\}$ while y induces the separation $\{\mathcal{A}^+ \cup \mathcal{A}^-, \mathcal{B}^+ \cup \mathcal{B}^-\}$, and that D contains the punctures corresponding to $\mathcal{A}^{\pm}$. Then H_D swaps the two punctures corresponding to $\mathcal{A}^{\pm}$, which shows that $H_D^k(x)$ induces the same puncture separation as x if and only if k is even.

Now let $\gamma \in \mathcal{N}$ such that $x' = \gamma(x)$. If $\gamma = 1$ then we are done. Otherwise let (s, γ_s) be as in Proposition 2.3. If $d_{\mathcal{C}}(x, s) \leqslant 1$ then we may apply γ_s to the whole data and proceed by induction on the complexity of γ. Otherwise $d_s(x, x') \geqslant \Theta$. Now we claim that any other $z \in S \cup \{y\}$ is in the star of s, so that we can replace x' with $\gamma_s(x')$ and proceed by induction. If this is not the case then $\pi_s(z)$ is well-defined, and we have that

$$\begin{aligned} d_s(x, x') &\leqslant d_s(x, z) + d_s(x', z) = d_s(x, z) + d_s(T_y^m(x), T_y^m(z)) \\ &= d_s(x, z) + d_{T_y^{-m}(s)}(x, z) \end{aligned}$$

where we used that z belongs to the star of y and is therefore fixed by T_y. But now

$$\Theta \leqslant d_s(x, x') \leqslant d_s(x, z) + d_{T_y^{-m}(s)}(x, z) \leqslant 2 \max_{p,q \in X_b,\, s \in \mathcal{C}} d_s(p, q) \leqslant 2M$$

where M is the constant from Lemma 3.6, which does not depend on Θ. This is a contradiction if $\mathcal{N}$ is deep enough. □

Now, in order to complete our lift of $\overline{X}_b$ we still need to lift the vertices $\overline{\alpha}_1, \dots, \overline{\alpha}_{b-3} \in \overline{X}_b$ outside the link of $\overline{\beta}$. For every $i = 1, \dots, b-3$ we can find vertices $\overline{x}, \overline{y} \in \overline{X}_{b-1}$ and a codimension two simplex $\overline{R} \subset \overline{X}_{b-1}$ such that $\overline{\beta}$, $\overline{\alpha}_1$, $\overline{\alpha}_i$, $\overline{x}$, $\overline{y}$ and $\overline{R}$ span a special pentagon $\overline{\mathfrak{P}}$. For example, one could look at Figure 6, replace z with α_i and then consider the corresponding vertices inside $\overline{X}_b$.

Lift $\overline{\mathfrak{P}}$ to a special pentagon $\mathfrak{P}'$, let $Q' = \beta' \star y' \star R'$ be the lift of the simplex $\overline{Q} = \overline{\beta} \star \overline{y} \star \overline{R}$ inside $\mathfrak{P}'$, and let $Q = \beta \star y \star R$ be the corresponding simplex inside $\beta \star X_{b-1}$. Since there is a unique $\mathcal{N}$-orbit of lifts of $\overline{Q}$ there exists $g \in \mathcal{N}$ mapping Q' to Q. Thus we can apply g to $\mathfrak{P}'$ and assume that $Q = Q'$. Moreover, let $x' \in \mathfrak{P}'$ be the lift of $\overline{x}$, and x the corresponding lift inside X_{b-1}. If we apply Lemma 4.4 to x, x', y and $S = \beta \star R$ we find some element $g \in \mathcal{N}$ that maps x' to x and fixes Q. Thus, we can apply g to $\mathfrak{P}'$ and assume that also $x = x'$.

Now, let $G_i = \mathfrak{P}' \cup X_{b-1}$. Let $\alpha_1^i \in \mathfrak{P}'$ be the lift of $\overline{\alpha}_1$ inside the special pentagon. Lemma 2.21 provides a copy X_b^i which extends $\beta \star X_{b-1}$ and contains α_1^i. By Lemma 2.22, X_b^i also contains α_i, therefore $G_i \subset X_b^i$. Now, consider the copies X_b^2 and X_b^j, for $3 \leqslant j \leqslant b-3$. These copies coincide on $\beta \star X_{b-1}$, and by construction ${\alpha_1}^2$ and α_1^j project to the same element $\overline{\alpha}_1$. Therefore, if we apply Lemma 4.4 with $x = \alpha_1^2$, $x' = \alpha_1^j$, $y = \beta$ and $S = X_{b-1}$, we find an element $g \in \mathcal{N}$ that maps ${\alpha_1}^2$ to α_1^j and fixes $\beta \star X_{b-1}$. Then we conclude that $g(X_b^2) = X_b^j$ by Lemma 2.22, which implies that the projection of X_b^2 coincides with $\overline{X}_b$ also on $\overline{\alpha_j}$. Since we may do this for every $3 \leqslant j \leqslant b-3$ we see that X_b^2 is the lift we were looking for. □

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

4.2. Uniqueness of the lift

Proof. Proof of uniqueness. Again we proceed by induction. First we discuss the base case $b = 5$. □

LEMMA 4.5. *If $\mathcal{N}$ is deep enough, any two lifts of $\overline{X}_5$ inside $\mathcal{C}_5$ differ by an element of $\mathcal{N}$.*

Proof. The proof is just a sequence of iterations of Lemma 4.4. More precisely, let $\{\overline{\gamma}_i\}_{i=1,\dots,5}$, be the vertices of $\overline{X}_5$, and let γ_i and γ_i' be lifts that form two pentagons X_5 and X_5'. Up to applying some element of $\mathcal{N}$ to X_5' we may assume that $\gamma_1 = \gamma_1'$ and $\gamma_2 = \gamma_2'$, because the edge $\overline{\gamma}_1, \overline{\gamma}_2$, which is a simplex of dimension 1, admits a unique orbit of lifts by Lemma 2.7. Now, if we apply Lemma 4.4 with $x = \gamma_3$, $x' = \gamma_3'$, $y = \gamma_1$ and $S = P = \gamma_2$ we find some element $g \in \mathcal{N}$ mapping γ_3 to γ_3' and fixing γ_1, γ_2, so we can apply g to X_5 and assume that $\gamma_3 = \gamma_3'$. Moreover, if we repeat the argument with $x = \gamma_4$, $x' = \gamma_4'$, $y = \gamma_2$, $P = \gamma_3$ and $S = \{\gamma_1, \gamma_3\}$ we may also assume that $\gamma_4 = \gamma_4'$. Now X_5 and X_5' coincide on four curves, and therefore they must coincide because of Lemma 2.24. □

Now assume that $b \geqslant 6$ and every copy of X_{b-1} inside $\mathcal{C}_{b-1}/\mathcal{N}(S_{b-1})$ has a unique lift inside $\mathcal{C}_{b-1}$. The next step is the following:

LEMMA 4.6. *If $\overline{\beta} \in \overline{X}_b$ is a minimal vertex, then $\overline{G} = \overline{\beta} \star \overline{X}_{b-1}$ has a unique lift inside $\mathcal{C}_b$, up to elements in $\mathcal{N}$.*

Proof. Let $G = \beta \star X_{b-1}$ and $G' = \beta' \star X_{b-1}'$ be two lifts of $\overline{G}$. Up to an element of $\mathcal{N}$ we may assume that $\beta = \beta'$. Now X_{b-1} and X_{b-1}' are in $\mathrm{Lk}_{C_b}(\beta) \cong \mathcal{C}_{b-1}$, where again we see S_{b-1} as the surface obtained by shrinking β to a puncture. Moreover, as already noticed the projection of $\mathrm{Lk}_{\mathcal{C}_b}(\beta)$ is isomorphic to $\mathcal{C}_{b-1}/\mathcal{N}(S_{b-1})$, thus by induction there is an element $h \in \mathcal{N}(S_{b-1})$ that maps X_{b-1} to X_{b-1}'. If we extend h to be the identity on the twice-punctured disk bounded by β we get a mapping class $H \in \mathcal{N}(S_b)$ that maps G to G', as required. □

Now we are finally able to prove the uniqueness part of Theorem 4.1. Choose two lifts X_b and X_b' of $\overline{X}_b$. By the previous lemma we may assume that they coincide on $\beta \star X_{b-1}$. Now apply Lemma 4.4 with $x = \alpha_1$, $x' = \alpha_1'$, $y = \beta$ and $S = X_{b-1}$, so that we may assume that X_b and X_b' coincide also on α_1. Finally, by Lemma 2.22 we see that the two lifts now coincide.

5. Combinatorial rigidity

The aim of this section is to prove Theorem 3 (combinatorial rigidity).

5.1. The four-punctured case

Firstly, we show an analogue of Theorem 3 for the special case $b = 4$. The proof will highlight the core of the general case, though the latter will require some more machinery. The subgraph playing the role of the rigid set will be any triangle T

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

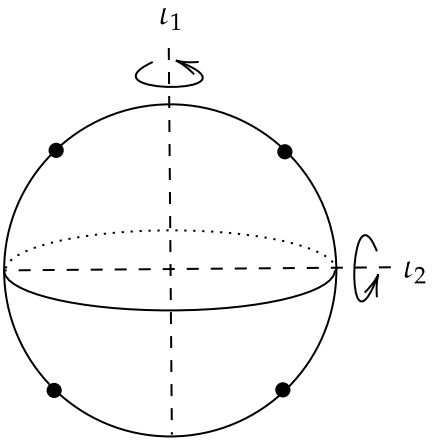


Figure 8. The hyperelliptic involutions ι_1, ι_2 generate a Klein four-group $\mathcal{K}$ which is the kernel of the action of $MCG^{\pm}$ over the curve graph.

inside the Farey complex, since the following well-known result holds (for example, it follows from the discussions in [3, Section 3] and in [21, Section 3.4]):

THEOREM 5.1. *Given two triangles $T, T' \subset \mathcal{C}_4$ there exists an extended mapping class $h \in MCG^{\pm}(S_{0,4})$ mapping T to T. Any two such* h *differ by an element of the Klein four-group $\mathcal{K}$, generated by the classes of the two hyperelliptic involutions in Figure 8.*

Now we want to prove the following:

THEOREM 5.2. *There exists $K_0 \in \mathbb{N}_{>0}$ such that if* K *is a non-trivial multiple of* K_0 *the action $MCG^{\pm}(S_{0,4})/DT_K \to Aut(\mathcal{C}(S_{0,4})/DT_K)$ is an epimorphism, with kernel the projection of the Klein four-group $\mathcal{K}$, generated by the classes of the two hyperelliptic involutions in Figure 8.*

We split the proof in two propositions below.

PROPOSITION 5.3. *For any normal subgroup $\mathcal{N}$ deep enough, any automorphism $\overline{\phi} \in Aut(\mathcal{C}_4/\mathcal{N})$ is induced by an element of $MCG^{\pm}(S_4)$.*

The proof relies on the following:

LEMMA 5.4. *The following facts hold:*

(1) *Any two adjacent curves $\alpha, \beta \in \mathcal{C}_4$ belong to exactly two triangles.*
(2) *Any two adjacent vertices $\overline{\alpha}, \overline{\beta} \in \mathcal{C}_4/\mathcal{N}$ belong to at most two triangles (namely, the projections of the two triangles that complete some of their lifts).*

Proof. Item (1) is true by construction of the Farey complex. Regarding Item (2) fix a pair α, β of adjacent lifts of $\overline{\alpha}, \overline{\beta}$. Any triangle $\overline{T} = \{\overline{\alpha}, \overline{\beta}, \overline{\gamma}\}$ has a lift

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

$\{\alpha', \beta', \gamma'\}$, and by uniqueness of the orbit of lifts of an edge (which is a particular case of Theorem 2.7) we may assume that $\alpha = \alpha'$ and $\beta = \beta'$. Therefore, $\overline{T}$ is the projection of one of the two triangles containing α, β. □

Proof of Proposition 5.3. Let $T \subseteq \mathcal{C}_4$ be a triangle, and let $\overline{T}$ be its projection, which is still an isometrically embedded triangle by Item (1) of Lemma 3.13. Let $\widetilde{T}$ be a lift of $\overline{\phi}(\overline{T})$, which exists by Theorem 2.7. By Theorem 5.1 there is a mapping class $h \in MCG^{\pm}(S_4)$ mapping T to $\widetilde{T}$, which means that $\overline{\phi}$ coincides with the induced map $\overline{h}$ on $\overline{T}$. In other words, we showed that the following diagram commutes on $T \subset \mathcal{C}_4$:

$$\begin{array}{ccc} \mathcal{C}_4 & \xrightarrow{h} & \mathcal{C}_4 \\ \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle \pi} \\ \mathcal{C}_4/\mathcal{N} & \xrightarrow{\overline{\phi}} & \mathcal{C}_4/\mathcal{N} \end{array} \tag{1}$$

Our next goal is to show that, if the diagram commutes on some triangle T, then it commutes on any triangle T' that shares an edge with T. If we prove the claim we are done, since any triangle in the Farey complex maps to a fixed triangle via a finite sequence of reflections along sides. Let $\alpha, \beta \in T$ be the vertices of the common edge. Notice that h, which is a graph automorphism, must map T' to the triangle $\widetilde{T}$ that contains $h(\alpha)$ and $h(\beta)$ and is not $\widetilde{T}$, and similarly $\overline{\phi}$ must map $\pi(T')$ to $\pi(\widetilde{T}')$ since there are no other triangles containing $\pi(\alpha)$ and $\pi(\beta)$. Therefore Diagram (5.1) commutes also on T', as required. □

PROPOSITION 5.5. *For* $\mathrm{b} = 4$ *and any* $\mathcal{N}$ *deep enough, if* $g \in MCG^{\pm}$ *induces the identity on* $\mathcal{C}_4/\mathcal{N}$*, then* $g \in \mathcal{N}\mathcal{K}$.

Proof. Choose a copy $\overline{T} \subset \mathcal{C}/\mathcal{N}$ of the rigid set, and let T be one of its lifts. Then $g(T)$ is also a lift of $\overline{T}$, and by Lemma 2.9 there is an element $h \in \mathcal{N}$ mapping T to $g(T)$. Hence $h^{-1} \circ g$ is the identity on T, and therefore belongs to $\mathcal{K}$ by the "moreover" part of Theorem 5.1. □

While the injectivity part of the case $b \geqslant 5$ will be very similar to the proof of Proposition 5.5, for the surjectivity part we will need a more general version of Lemma 5.4. The following subsections are devoted to establish such a result, the main tool being half twists around minimal vertices.

5.2. Characterisation of half twists

From now on let $b \geqslant 5$. Let $\beta \in X_b \subseteq \mathcal{C}$ be a minimal curve, and recall that we denote by H_β the half twist around β. Our first goal is to show that for any $\alpha \in X_b$ its image $H_\beta(\alpha)$ admits a graph-theoretic characterisation, which therefore is preserved by any automorphism of the curve complex. More precisely we want to prove the following:

LEMMA 5.6. *For any* $\alpha \in X_b$*:*

- $H_\beta(\alpha) = \alpha$ *iff* $d_{\mathcal{C}}(\alpha, \beta) \leqslant 1$*;*

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

- *otherwise there exists a facet $P \subseteq X_b$ such that $H_\beta(\alpha)$ is one of the two curves with special intersection with both α and β, with respect to P.*

Notice that we did not uniquely characterise $H_\beta(\alpha)$, as it is impossible to distinguish it from $H_\beta^{-1}(\alpha)$ just from the curve graph (indeed, the orientation-reversing reflection r fixes X_b pointwise but swaps the half-twists).

Proof. The first bullet is clearly true, so we focus on the second. Let α be a curve in X_b that intersects β, and let P be a facet that both α and β complete, which cuts out a four-holed sphere S_4. Now, clearly $H_\beta(\alpha)$ has special intersection with both α and β with respect to P (the curves γ and δ are easy to find in both cases); therefore, it suffices to notice that there are only two curves on S_4 with intersection 2 with both α and β, since in the curve graph of S_4 the edge with endpoints α and β belongs to only two triangles. □

5.3. Definition of half twists in the quotient

For the rest of the section, let $\mathcal{N}$ be a deep enough subgroup.

Definition 5.7. *Let $\overline{X}_b$ be a copy of $\mathrm{X_b}$ inside $\mathcal{C}/\mathcal{N}$, and let $\overline{\beta} \in \overline{X}_b$ be a minimal vertex. Let $s : \overline{X}_b \to \mathcal{C}$ be a lift. Then we define the half twist around $\overline{\beta}$ as the map $H_{\overline{\beta}} \colon \overline{X}_b \to \mathcal{C}/\mathcal{N}$, defined as the composition $\pi \circ H_{s(\overline{\beta})} \circ s$.*

In other words, we define the half twist by lifting $\overline{X}_b$, applying a half twist and projecting, so that the half twist commutes with the quotient projection. The definition is well-posed, since any two lifts differ by some $g \in \mathcal{N}$. Notice moreover that $H_{\overline{\beta}}(\overline{X}_b)$ is still a copy of X_b by Theorem 3.7, since it is the projection of $H_{s(\overline{\beta})}\left(s(\overline{X}_b)\right)$.

Our next goal is to show that the combinatorial characterisation of the half twist described in Theorem 5.6 still works in the quotient. In order to do so we need to show that:

(1) $H_{\overline{\beta}}(\overline{\alpha})$ still satisfies the properties described in Theorem 5.6;
(2) Any $\overline{\eta}$ that satisfies the properties with $\overline{\alpha}$, $\overline{\beta}$ and $\overline{P}$ lifts to a curve that satisfies the properties with some lifts α, β, P. Therefore, $\overline{\eta}$ must be the projection of $H_\beta^{\pm 1}(\alpha)$.

The following subsections will be dedicated to the proofs of these items.

5.3.1. The projection of a twist looks like a twist

Let $\overline{X}_b$, $\overline{\alpha}$, $\overline{\beta}$ as above. For short we set $\alpha = s(\overline{\alpha})$ and similarly for any other vertex and subset of $s(\overline{X}_b)$. We subdivide the proof into two lemmas:

Lemma 5.8. *$H_{\overline{\beta}}(\overline{\alpha}) = \overline{\alpha}$ iff $d_{\mathcal{C}/DT_K}(\overline{\alpha}, \overline{\beta}) \leqslant 1$.*

Proof. Both $\overline{\alpha}$ and $\overline{\beta}$ belong to $\overline{X}_b$, which has diameter 2. If $d_{\mathcal{C}/DT_K}(\overline{\alpha}, \overline{\beta}) \leqslant 1$ then $H_{\overline{\beta}}(\overline{\alpha}) = \overline{\alpha}$, as by construction $H_{\overline{\beta}}$ is the identity on the link of $\overline{\beta}$. If otherwise

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

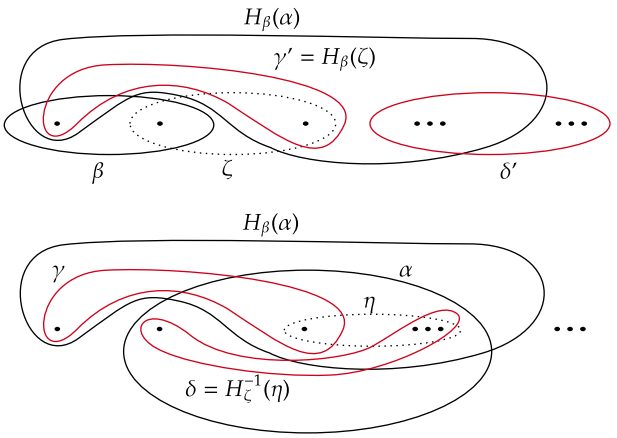


Figure 9. An example of how the auxiliary curves (here in red) may be chosen only using half twists around β and ζ, which are both in X_b. In our example $\gamma = \gamma'$, but this is not necessary. The case where α bounds a twice-punctured disk can be dealt with similarly.

$\overline{\alpha}$ is at distance 2 from $\overline{\beta}$, then α and $H_\beta(\alpha)$ are different curves in Y_b, and by Theorem 3.7 they project to different points. □

LEMMA 5.9. *If $d_{\mathcal{C}/DT_K}(\overline{\alpha}, \overline{\beta}) = 2$ then there exists a facet $\overline{P}$ such that $H_{\overline{\beta}}(\overline{\alpha})$ has special intersection with both $\overline{\alpha}$ and $\overline{\beta}$, with respect to $\overline{P}$.*

Proof. Choose a facet $P \subseteq X_b$ that both α and β complete, and let $\gamma, \delta \in X_b$ be some auxiliary curves detecting their special intersection, as in Figure 6. The images of these curves under H_β are auxiliary curves for β and $H_\beta(\alpha)$. Moreover we can find $\gamma', \delta' \in Y_b$ that detect the special intersection between α and $H_\beta(\alpha)$ (see Figure 9 for an example). The graph G spanned by these curves is the union of two special pentagons patched along the simplex $H_\beta(\alpha) \star P$, as in Figure 10. Since G it is a subgraph of the graph Y_b, defined in Definition 3.5, by Theorem 3.7 its projection is an isometric embedding of two special pentagons, detecting that $H_{\overline{\beta}}(\overline{\alpha}) = \pi(H_\beta(\alpha))$ has special intersection with both $\overline{\alpha}$ and $\overline{\beta}$. □

5.3.2. *If it looks like a twist it is the projection of a twist*

LEMMA 5.10. *Let $\beta \in X_b \subset \mathcal{C}$ be a minimal curve, and let $\alpha \in X_b$ have special intersection with β with respect to the facet P. Let $\overline{\alpha}$, $\overline{\beta}$, $\overline{P}$ be their projections, and let $\overline{\eta} \in \mathcal{C}/\mathcal{N}$ be a vertex with special intersection with both $\overline{\alpha}$ and $\overline{\beta}$ with respect to $\overline{P}$. Then $\overline{\eta}$ lifts to a curve η with special intersection with both α and β with respect to P.*

Proof. Consider the subgraph $\overline{G} \subseteq \mathcal{C}/\mathcal{N}$ made of two special pentagons that detect the special intersections, as in Figure 10. This subgraph has a lift $G \subseteq \mathcal{C}$ (more precisely, each one of the special pentagons lifts, and we may arrange that the lifts coincide on $\overline{\eta} \star \overline{P}$ since there is a unique $\mathcal{N}$-orbit of lifts of simplices). Now, as a corollary of Lemma 4.4 the union of the two simplices $\overline{\alpha} \star \overline{P} \cup \overline{\beta} \star \overline{P}$ admits a unique

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

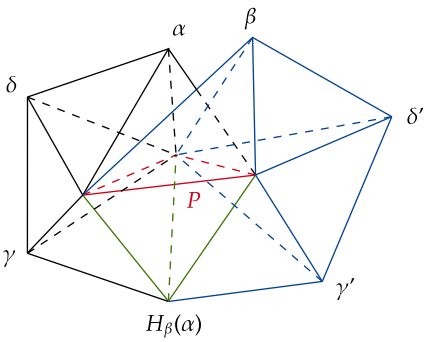


Figure 10. The configuration G that detects the special intersection of $H_\beta(\alpha)$ with both α and β, that consists in two special pentagons (here, in black and blue) which overlap over a simplex (here, in green and red). Possibly $\gamma = \gamma'$ and $\delta = \delta'$, but each of the pentagons is isometrically embedded.

lift, up to elements in $\mathcal{N}$, and therefore, we may assume that G is glued to X_b in such a way that the corresponding copies of α, β, P coincide. But now G detects that $\overline{\eta}$ lifts to a curve with special intersection with the copies of α and β inside X_b, as required. □

5.4. The general case

We are finally ready to prove the following combinatorial rigidity statement, which implies Theorem 3 when $\mathcal{N} = DT_K$ for suitable K.

THEOREM 5.11. *For every $b \geqslant 5$ and for every deep enough subgroup $\mathcal{N}$, the natural map $MCG^{\pm}(S_b)/\mathcal{N} \to Aut(\mathcal{C}(S_b)/\mathcal{N})$ is an isomorphism.*

As before, we split the proof into two theorems.

THEOREM 5.12. *For $b \geqslant 5$ and any $\mathcal{N}$ deep enough, any automorphism $\overline{\phi} \in Aut(\mathcal{C}/\mathcal{N})$ is induced by an element of $MCG^{\pm}(S_b)$.*

Proof. Consider X_b inside $\mathcal{C}$ and let $\overline{X}_b$ be its projection, which is still a copy of X_b by Theorem 3.7. Moreover, let $\widetilde{X}_b$ be a lift of $\overline{\phi}(\overline{X}_b)$. By Theorem 2.19 there is an extended mapping class $h \in MCG^{\pm}(S_b)$ mapping X_b to $\widetilde{X}_b$, which means that the following diagram commutes on $X_b \in \mathcal{C}$:

$$\begin{array}{ccc} \mathcal{C} & \xrightarrow{h} & \mathcal{C} \\ \downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle \pi} \\ \mathcal{C}/\mathcal{N} & \xrightarrow{\overline{\phi}} & \mathcal{C}/\mathcal{N} \end{array} \tag{2}$$

Now, choose a minimal vertex $\beta \in X_b$, and let r be the reflection fixing X_b. We know that, if we fix a curve $\alpha_1 \in X_b$ that intersects β, h must map $H_\beta(\alpha_1)$ to either $H_{h(\beta)}(h(\alpha_1))$ or $H^{-1}_{h(\beta)}(h(\alpha_1))$, since these are the only two curves that satisfy the graph-theoretic characterisation of the half twist, Theorem 5.6. Section 5.3

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

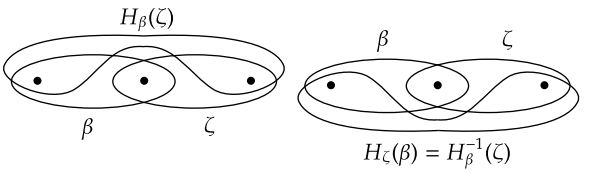


Figure 11. The two minimal curves from the proof of theorem 5.12 and their respective half twists.

shows that the same argument works in the quotient, hence there are only two possibilities for $\overline{\phi} \circ H_{\overline{\beta}}(\overline{\alpha}_1)$. Up to replacing h with $h \circ r$ we may therefore assume that Diagram 5.2 commutes also on $H_\beta(\alpha_1)$, and hence on the whole $H_\beta(X_b)$, which by Lemma 2.22 is uniquely determined by $H_\beta(\alpha_1)$ (and the copy of X_{b-1} which is the link of β in X_b; this is fixed by H_β).

We now claim that, if the diagram commutes on $X_b \cup H_\beta(X_b)$ and ζ is a minimal curve in X_b that intersects β then the diagram commutes also on $H_\zeta(X_b)$. In fact $H_\zeta^{-1}(\beta) = H_\beta(\zeta)$, as shown in Figure 11, therefore we already know that the diagram commutes on $H_\zeta^{-1}(\beta)$ (and therefore on $H_\zeta(\beta)$ which is the only other curve that satisfies the same characterisation as $H_\zeta^{-1}(\beta)$). Therefore, the diagram commutes also on the whole $H_\zeta(X_b)$, which is uniquely determined by $H_\zeta(\beta)$.

We may repeat the previous argument to show that Diagram 5.2 commutes on X_b and on any half twist around one of its minimal curves. Now it suffices to prove that the orbit of X_b by successive half twists covers $\mathcal{C}$. That is, we are left to prove:

Claim: Given two copies X_b and X_b', there is a sequence $X_b = X^0, X^1, \dots, X^k = X_b'$ of copies of X_b such that, for every $i = 0, \dots, k-1$, $X^{i+1} = H_\theta^\pm(X^i)$ for some minimal curve $\theta \in X^i$.

Proof of Claim. To show this let h be a (orientation-preserving) mapping class that maps X_b to X_b'. Now, if $\theta_1, \dots, \theta_b$ are the minimal curves in X_b we know that $H_{\theta_1}, \dots, H_{\theta_b}$ generate the mapping class group. Therefore, let $h = H_{\theta_{i_1}}^\pm \dots H_{\theta_{i_k}}^\pm$, for some possibly repeated indices $i_1, \dots, i_k \in \{1, \dots, b\}$. Now we proceed by induction on k. If $k = 1$ we are done. Otherwise notice that, by the properties of half Dehn twists,

$$h = H_{\theta_{i_1}}^\pm \dots H_{\theta_{i_k}}^\pm = H_{H_{\theta_{i_1}}^\pm(\theta_{i_2})}^\pm \dots H_{H_{\theta_{i_1}}^\pm(\theta_{i_k})}^\pm H_{\theta_{i_1}}^\pm$$

Now we have expressed h in terms of $H_{\theta_{i_1}}^\pm$ and $k-1$ half twists around some minimal curves of $X^1 := H_{\theta_{i_1}}^\pm(X_b)$. Therefore, we conclude by the inductive hypothesis. □

This concludes the proof of Theorem 5.12. □

Now we turn to the injectivity part:

THEOREM 5.13. *For any $b \geqslant 5$ and any $\mathcal{N}$ deep enough, if $g \in MCG^\pm$ induces the identity on $\mathcal{C}/\mathcal{N}$, then $g \in \mathcal{N}$.*

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

Proof. Choose a copy of $\overline{X}_b \subseteq \mathcal{C}/\mathcal{N}$, and let X_b be one of its lifts. Then $g(X_b)$ is also a lift of $\overline{X}_b$, and by the uniqueness part of Theorem 4.1 there is an element $h \in \mathcal{N}$ mapping X_b to $g(X_b)$. Therefore, $h^{-1} \circ g$ is the identity on X_b, and by the "moreover" part of Theorem 2.19 we must have that $h^{-1} \circ g \in \langle r \rangle$. Now, $h^{-1}g$ induces the identity map on $\mathcal{C}/\mathcal{N}$, as g does so by hypothesis and h is an element of $\mathcal{N}$. Therefore, in order to prove that $h^{-1}g$ must be the identity, we must prove that r does not induce the identity on $\mathcal{C}/\mathcal{N}$. If this is not the case then, for any minimal curve $\beta \in X_b$, we have that $H_\beta(X_b)$ and $r \circ H_\beta(X_b) = r \circ H_\beta \circ r(X_b) = H_\beta^{-1}(X_b)$ coincide in the quotient (here we used that $r(X_b) = X_b$ and the properties of a half Dehn twist under conjugation). Therefore, by uniqueness of the orbit of lifts, there is an element $k \in \mathcal{N}$ which, for every curve $\gamma \in X_b$, maps $H_\beta^{-1}(\gamma) = r \circ H_\beta(\gamma)$ to $H_\beta(\gamma)$. By the "moreover" part of Theorem 2.19 there is only one orientation-preserving mapping class with this property, and therefore we must have that $T_\beta = k \in \mathcal{N}$ since the Dehn twist T_β clearly maps $H_\beta^{-1}(\gamma)$ to $H_\beta(\gamma)$. But then since $\mathcal{N}$ is a normal subgroup it must contain all conjugates of T_β, that is, all Dehn twists around curves that bound twice-punctured disks. These mapping classes generate the *pure* mapping class group $PMCG(S_b)$, which is the finite index subgroup of $MCG(S_b)$ consisting of elements that fix each puncture individually. But this subgroup is not deep enough if we choose a large enough constant Θ (for example, since the quotient $\mathcal{C}/PMCG$ is finite, which contradicts Corollary 3.8). □

5.5. Finite rigidity of the quotient

As a by-product of the proofs of Theorems 5.11 and 5.2 we get an analogue of Theorem 2.19 for the quotient.

COROLLARY 5.14. (Finite rigidity of $\mathcal{C}/\mathcal{N}$) *For $b \geqslant 4$ and $\mathcal{N}$ deep enough, for every fixed copy $\overline{X}_b \subset \mathcal{C}/\mathcal{N}$ (where by $\overline{X}_4$ we mean a triangle), any isometric embedding $\phi : \overline{X}_b \to \mathcal{C}/\mathcal{N}$ is induced by a mapping class $\overline{h} \in MCG/\mathcal{N}$. Moreover, two such $\overline{h}$ differ by an element of the pointwise stabiliser $PStab(\overline{X}_b)$, which is the projection of $Pstab(X_b)$.*

Proof. For the existence part let X_b and X_b' be lifts of $\overline{X}_b$ and $\phi(\overline{X}_b)$, respectively. Then by finite rigidity there is a mapping class $h \in MCG$ mapping X_b to X_b', and the induced map $\overline{h} \in MCG/\mathcal{N}$ maps $\overline{X}_b$ to $\phi(\overline{X}_b)$.

For the uniqueness part, if h and h' both induce maps that map $\overline{X}_b$ to $\phi(\overline{X}_b)$ then $h(X_b)$ and $h'(X_b)$ are both lifts of $\phi(\overline{X}_b)$, hence by uniqueness of the orbit of lifts there exists a $g \in \mathcal{N}$ such that $g \circ h|_{X_b} = h'|_{X_b}$. But then $g \circ h$ and h' must differ by an element of $\mathrm{Pstab}(X_b)$. □

6. From 1-separating to strongly separating

DEFINITION 6.1. *We say that a curve $\gamma \in \mathcal{C}$ is* 1-separating *if $S - \gamma$ has a connected component of complexity 1. If S is a punctured sphere, γ is 1-separating if and only if it bounds a disk with three punctures. Let $\mathcal{C}^1$ be the subgraph of $\mathcal{C}$ spanned by 1-separating curves.*

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

Moreover, we recall the following definition from [12]:

DEFINITION 6.2. *Let $\mathcal{C}^{ss}$ the full subgraph of the curve graph spanned by all* strongly separating *curves, i.e. those separating curves that do not bound a twice-punctured disk.*

Notice that $\mathcal{C}^1 \subseteq \mathcal{C}^{ss}$ whenever $b \geqslant 6$. Moreover, the action of DT_K restricts to both $\mathcal{C}^1$ and $\mathcal{C}^{ss}$, therefore the quotient $\mathcal{C}^1/DT_K$ is a subgraph of $\mathcal{C}^{ss}/DT_K$, which is in turn a subgraph of $\mathcal{C}/DT_K$.

As we will see in § 9, every self-quasi-isometry of MCG/DT_K induces an automorphism of $\mathcal{C}^1/DT_K$. Therefore, we must find a way to relate an automorphism of $\mathcal{C}^1/DT_K$ with an automorphism of $\mathcal{C}/DT_K$, which we understand in view of the combinatorial rigidity Theorem 5.11. For short, in the rest of the paper we will say that a proposition holds *for all large multiples* K if there exists $K_0 \in \mathbb{N}_{>0}$ such that the proposition holds whenever K is a non-trivial multiple of K_0.

THEOREM 6.3. *For all $b \geqslant 7$ and for all large multiples K, any automorphism of $\mathcal{C}^1/DT_K$ is the restriction of an automorphism of $\mathcal{C}/DT_K$.*

Combining this theorem with the combinatorial rigidity Theorem 5.11 we get the following:

THEOREM 6.4. *For all $b \geqslant 7$ and for all large multiples K, any automorphism of $\mathcal{C}^1/DT_K$ is the restriction of an element of $MCG^{\pm}(S_b)/DT_K$.*

Following the roadmap of [12, 13] we split the proof of Theorem 6.3 into two intermediate extensions, one from $\mathcal{C}^1/DT_K$ to $\mathcal{C}^{ss}/DT_K$ and one from $\mathcal{C}^{ss}/DT_K$ to $\mathcal{C}/DT_K$.

For the rest of the section, we assume $b \geqslant 7$ and we take a multiple K which is large enough that $\mathcal{C}/DT_K$ satisfies the Combinatorial Rigidity Theorem 3. We first recall some definitions:

DEFINITION 6.5. *Let* G *be a graph. The* dual graph G^* *of* G *is the graph with the same vertices of* G *and with an edge between two vertices if and only if they are not adjacent in* G.

Notice that if $G \subseteq \mathcal{C}$ then G^* is connected iff any two curves can be joined by a chain as in Definition 2.11.

The following definition is from [13, Section 7]:

DEFINITION 6.6. *Let Γ be either $\mathcal{C}^1$ or $\mathcal{C}^1/DT_K$. A* division *of Γ is an unordered pair $\{P^+, P^-\}$ of disjoint infinite subsets such that $P^+ \star P^- \subset \Gamma$ is a maximal join and both $(P^+)^*$ and $(P^-)^*$ are connected. Two divisions $P^{\pm}, Q^{\pm}$ are nested if either $P^+ \subseteq Q^+$ or $P^+ \subseteq Q^-$.*

By [13, Lemma 7.2] and the discussion following [13, Lemma 7.5], any division of $\mathcal{C}^1$ is induced by a unique strongly separating curve $\delta \in \mathcal{C}^{ss} \setminus \mathcal{C}^1$: namely, $P^{\pm}(\delta)$ correspond to the 1-separating curves that fill the two sides of δ. Moreover, nesting is equivalent to the corresponding δs being disjoint. Our goal is to give a similar description for the classes of $(\mathcal{C}^{ss}/DT_K) \setminus (\mathcal{C}^1/DT_K)$.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

DEFINITION 6.7. *Let* Γ *be either* $\mathcal{C}^1$ *or* $\mathcal{C}^1/DT_K$. *A* filled division *of* Γ *is a triple* $(P^\pm, \alpha^\pm, \beta^\pm)$ *such that:*

(1) $P^\pm$ *is a division of* Γ, *called the* underlying *division;*
(2) $\alpha^\pm, \beta^\pm \in P^\pm$, *respectively;*
(3) $d_\Gamma(\alpha^+, \beta^+) = d_\Gamma(\alpha^-, \beta^-) = 2$
(4) $P^+ = Lk_\Gamma(\alpha^-) \cap Lk_\Gamma(\beta^-)$, *and similarly for* P^-.

Two filled divisions are said to be equivalent if they have the same underlying $P^\pm$. *From now on, by division we will always mean an equivalence class of filled divisions (i.e., we won't consider divisions which do not allow a filling).*

REMARK 6.8. When $\Gamma = \mathcal{C}^1$, any division $P^\pm$ in the sense of Bowditch admits a filling. For example, since the division is induced by some strongly-separating curve δ, we may choose $\alpha^\pm, \beta^\pm$ to be two pairs of curves that respectively fill the two surfaces cut out by δ (hence the name "filled division"). Notice, however, that this is not necessarily the case: the subsurfaces filled by $\alpha^\pm, \beta^\pm$ could cut out some twice punctured disks, which are irrelevant to our argument since they cannot contain any 1-separating curve. We will elaborate on this Remark in the proof of Lemma 6.14, and the situation will be depicted in Figure 13.

REMARK 6.9. We could state Definition 6.7 just in terms of four vertices, by saying that a filled division is given by two pairs of vertices $(\alpha^\pm, \beta^\pm) \in \Gamma$ such that:

(1) $\alpha^+, \alpha^-, \beta^+, \beta^-$ is an isometrically embedded square;
(2) $P^\pm := \mathrm{Lk}_\Gamma(\alpha^\mp) \cap \mathrm{Lk}_\Gamma(\beta^\mp)$ are infinite;
(3) if $\gamma \in P^+$ and $\gamma' \in P^-$ then $d_\Gamma(\gamma, \gamma') = 1$.

This is because by Conditions 2 and 3 of Definition 6.9 we have that P^+ and P^- are disjoint infinite subsets that form a join $P^+ \star P^-$, which is maximal since any subgraph $Q^+ \subset \Gamma$ which forms a join with α^- and β^- must be contained in P^+. Moreover, Condition 1 says that α^+ and β^+ are connected in $(P^+)^*$. Finally, since P^+ and P^- are disjoint we get that there is no $\gamma \in \Gamma$ such that $\alpha^\pm, \beta^\pm \in \mathrm{Lk}_\Gamma(\gamma)$; therefore a vertex in $P^+ = \mathrm{Lk}_\Gamma(\alpha^-) \cap \mathrm{Lk}_\Gamma(\beta^-)$ which is not α^+ nor β^+ must be at distance at least 2 from either α^+ or β^+, so $(P^+)^*$ is connected.

The following definition is needed to take into account also classes in $\mathcal{C}^1/DT_K$:

DEFINITION 6.10. *A* slice *(respectively,* filled slice*) is either a division (resp. filled division) or a vertex* $\delta \in \Gamma$. *A vertex is nested into a division* $P^\pm$ *if it belongs to either* P^+ *or* P^-. *Two vertices are nested if they are disjoint.*

Now we want to assign to every $\overline{\delta} \in (\mathcal{C}^{ss}/DT_K) \setminus (\mathcal{C}^1/DT_K)$ some filled division $(\overline{P}^\pm, \overline{\alpha}^\pm, \overline{\beta}^\pm)$. First, choose representatives $\delta_1, \dots, \delta_k \in \mathcal{C}^{ss} \setminus \mathcal{C}^1$ for every homeomorphism type of strongly separating, non-1-separating curves. For every $i = 1 \dots, k$ choose $\alpha_i^\pm, \beta_i^\pm$ as two pairs of curves in $\mathcal{C}^1$ that fill the sides of δ_i. Then for every $\delta \in \mathcal{C}^{ss} \setminus \mathcal{C}^1$ choose some mapping class f such that $\delta = f(\delta_i)$ for some $i \in \{1, \dots, k\}$, and set $(\alpha^\pm(\delta), \beta^\pm(\delta)) = f(\alpha_i^\pm, \beta_i^\pm)$, which are again two pairs of filling curves for

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

the sides of δ. Now, choose Θ big enough that, for every $i = 1, \dots, k$ and for every $s \in \mathcal{C}$, we have that $\max\{d_s(\alpha_i^{\pm}, \beta_i^{\pm}), d_s(\alpha_i^{\pm}, \beta_i^{\mp})\} \leqslant \Theta$. Since $(\alpha^{\pm}(\delta), \beta^{\pm}(\delta))$ are defined as images via mapping classes of some $(\alpha_i^{\pm}, \beta_i^{\pm})$ we also get that, for every $\delta \in \mathcal{C}^{ss} \setminus \mathcal{C}^1$ and for every $s \in \mathcal{C}$,

$$\max\{d_s(\alpha^{\pm}(\delta), \beta^{\pm}(\delta)), d_s(\alpha^{\pm}(\delta), \beta^{\mp}(\delta))\} \leqslant \Theta. \tag{6.1}$$

This way, if K is a large enough multiple with respect to this constant Θ we will be able to invoke Lemma 3.3 and its corollaries, which will tell us that the projection map is isometric when restricted to these curves. The key point is that, since there are only finitely many MCG-orbits of curves δ, we could fix once and for all a MCG-equivariant choice of four curves $\alpha^{\pm}(\delta), \beta^{\pm}(\delta)$ in every orbit, and therefore, we only need to bound the annular projections of finitely many curves. This kind of argument will be recurrent throughout the paper.

Now let $(\overline{\alpha}^{\pm}, \overline{\beta}^{\pm}) := \pi(\alpha^{\pm}(\delta), \beta^{\pm}(\delta))$.

LEMMA 6.11. *The two pairs $(\overline{\alpha}^{\pm}, \overline{\beta}^{\pm})$ give a filled division.*

Proof. We show that the conditions of Remark 6.9 are satisfied. First notice that, since by construction all projections between $\alpha^{\pm}$ and $\beta^{\pm}$ are short (see Equation 6.1), Lemma 3.3 tells us that the quotient map π is an isometry on the square spanned by these curves. This proves Condition 1.

For Condition 2 let $P^{\pm} \subset \mathcal{C}^1$ be the division induced by δ, and let $\overline{P}^{\pm}$ be defined as in Remark 6.9. Clearly $\pi(P^{\pm}) \subset \overline{P}^{\pm}$, since any $\gamma \in P^-$ is disjoint from both α^+ and β^+ and therefore $\pi(\gamma) \in \mathrm{Lk}(\overline{\alpha}^+) \cap \mathrm{Lk}(\overline{\beta}^+)$, again by Lemma 3.13. Thus it is enough to show that $\pi(P^{\pm})$ is infinite. Let $\Sigma^{\pm}$ be the two subsurfaces cut out by δ, so that $P^{\pm} \subset \mathcal{C}(\Sigma^{\pm})$. Since $\delta \notin \mathcal{C}^1$ both subsurfaces have complexity at least 2, therefore $\pi(\mathcal{C}(\Sigma^{\pm})) \cong \mathcal{C}(\Sigma^{\pm})/DT_K(\Sigma^{\pm})$ by Corollary 3.2. Now we claim that we can find a large multiple K such that $\mathcal{C}(\Sigma^{\pm})/DT_K(\Sigma^{\pm})$ is infinite. In order to do so, for every topological type of Σ^+ we can choose a curve $x \in P^+$ and a pseudo-Anosov element $g \in MCG(\Sigma^+)$ that fixes the boundary δ pointwise. Notice that for every $n \in \mathbb{Z}$ we have that $g^n(x)$ is again 1-separating for S, thus $g^n(x) \in P^+$. Now we can proceed as in Corollary 3.8 to show that, whenever K is a large multiple, the projection π is an isometry on the axis $\{g^n(x)\}_{n\in\mathbb{Z}}$, and in particular $\pi(P^+)$ is infinite. Since there are only finitely many topological types of Σ^+ we can choose a large multiple K that works for any Σ^+. Moreover the whole argument can be repeated to show that $\pi(P^-)$ is infinite.

For Condition 3, first consider a simplex $\Delta^+ \in \mathcal{C}$ that contains δ, is disjoint from α^+ and β^+ and is maximal with these properties, and let $\overline{\Delta}^+ \in \mathcal{C}/DT_K$ be its projection, which is a simplex of the same dimension by Corollary 3.14. Let $\overline{\gamma} \in \mathcal{C}^1/DT_K$ be a vertex in $\mathrm{Lk}_{\mathcal{C}^1/DT_K}(\overline{\alpha}^+) \cap \mathrm{Lk}_{\mathcal{C}^1/DT_K}(\overline{\beta}^+)$. If we look at these three vertices inside $\mathcal{C}/DT_K$ we see that $\overline{\alpha}^+, \overline{\beta}^+ \in \mathrm{Lk}_{\mathcal{C}/DT_K}(\overline{\Delta}^+)$, i.e. we have a generalised square whose "vertices" are $\overline{\gamma}$, $\overline{\alpha}^+$, $\overline{\Delta}^+$, $\overline{\beta}^+$ as in Figure 12.

We may lift this square and assume that the lift of $\overline{\Delta}^+$ is Δ^+, by uniqueness of the orbit of lifts of a simplex, which is Theorem 2.7. In particular, we may assume

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

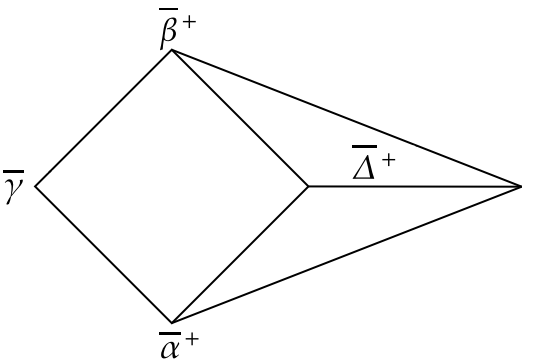


Figure 12. The generalised square from Lemma 6.11. The simplex $\overline{\Delta}^+$ is represented as a segment.

that the lift of $\overline{\delta}$ is still δ. Let α', β' be the lifts of $\overline{\alpha}^+$, $\overline{\beta}^+$, which again fill one of the sides of δ by Corollary 3.9. But now the lift γ of $\overline{\gamma}$ is disjoint from α' and β' and cannot coincide with δ, since $\delta \notin \mathcal{C}^1$. Therefore, γ must lie on the other side of δ with respect to α' and β', i.e. $\gamma \in P^-$. The same argument applies to any vertex $\overline{\eta} \in \mathrm{Lk}(\overline{\alpha}^-) \cap \mathrm{Lk}(\overline{\beta}^-)$ and produces a lift $\eta \in P^+$. But then γ and η lie on different sides of δ, and in particular they must be disjoint. Therefore, by Item 2 of Lemma 3.13, we get that $d_{\mathcal{C}^1/DT_K}(\overline{\gamma}, \overline{\eta}) = 1$, as required. □

Now we want to associate to every $\overline{\delta}$ in $\mathcal{C}^{ss}/DT_K$ a slice of $\mathcal{C}^1/DT_K$. If $\overline{\delta} \in \mathcal{C}^1/DT_K$ set $S(\overline{\delta}) := \overline{\delta}$. Otherwise choose a lift δ, take the corresponding filled division $(\alpha^{\pm}(\delta), \beta^{\pm}(\delta))$, constructed as above, and let $S(\overline{\delta}) := P^{\pm}(\overline{\delta})$ be the underlying division of $(\overline{\alpha}^{\pm}, \overline{\beta}^{\pm}) := \pi(\alpha^{\pm}(\delta), \beta^{\pm}(\delta))$.

THEOREM 6.12. *For all large multiples* K*, the map* S *described above is a well-defined bijection between vertices in* $\mathcal{C}^{ss}/DT_K$ *and slices of* $\mathcal{C}^1/DT_K$*, that translates adjacency into nesting. Hence any automorphism of* $\mathcal{C}^1/DT_K$ *extends to an automorphism of* $\mathcal{C}^{ss}/DT_K$*.*

We need to show that S is a well-defined bijection when restricted to $(\mathcal{C}^{ss}/DT_K)\setminus(\mathcal{C}^1/DT_K)$. If this is the case then clearly it maps adjacent vertices to nested slices and vice versa, and the proof of Theorem 6.12 is complete. We break the proof into a series of lemmas.

LEMMA 6.13. *The map S is well-defined. In other words, the corresponding division* $\overline{P}^{\pm}$ *of* $(\overline{\alpha}^{\pm}, \overline{\beta}^{\pm})$ *is independent on the choice of the lift* δ*.*

Proof. When we verified Conditions 2 and 3 in the proof of Lemma 6.11 we actually showed that $\overline{P}^{\pm} = \pi(P^{\pm})$, and the latter depends only on δ. Now it suffices to notice that, if $g \in DT_K$, then the division induced by $g(\delta)$ is clearly $g(P^{\pm})$. □

Now we look for an inverse of S.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

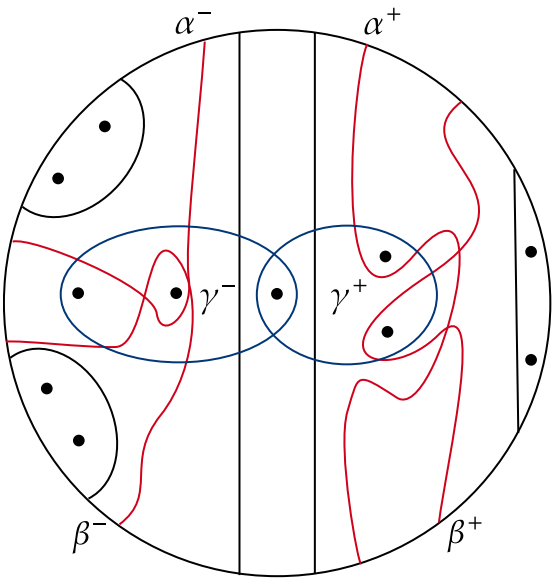


Figure 13. The red curves fill the subsurfaces whose boundaries are the black curves, and may cut out some punctured disks. If the annulus in between contains some punctures then the projections of the blue curves remain at distance at least 2.

LEMMA 6.14. *Let $\overline{P}^{\pm}$ be a division that admits a filling. There exists a vertex $\overline{\delta}$ such that $S(\overline{\delta}) = \overline{P}^{\pm}$.*

Proof. Let $(\overline{\alpha}^{\pm}, \overline{\beta}^{\pm})$ be two pairs of vertices whose corresponding division is $\overline{P}^{\pm}$. Lift them to some curves $(\alpha^{\pm}, \beta^{\pm})$ that form a square in $\mathcal{C}^1$. Let Σ^+, Σ^- be the subsurfaces filled by $\alpha^+ \cup \beta^+$ and $\alpha^- \cup \beta^-$, respectively. The boundaries of these subsurfaces are separating, as they are curves on a sphere, hence they cut out some peripheral punctured disks and a (possibly punctured) annulus between them. Moreover there is no curve $\gamma \in \mathcal{C}^1$ that does not cross any of the boundaries, because otherwise $\mathrm{Lk}(\pi(\gamma))$ would contain $(\overline{\alpha}^{\pm}, \overline{\beta}^{\pm})$. Therefore, we are in the situation of Figure 13: all peripheral disks contain two punctures each.

If we show that the annulus cannot contain any puncture then the core curve δ of this annulus is the only strongly separating curve which is disjoint from $\alpha^{\pm}$ and $\beta^{\pm}$, which can be also characterised as the only boundary component of the surface filled by $\alpha^+ \cup \beta^+$ which lies in $\mathcal{C}^{ss}$. If otherwise the annulus contains some puncture we can find two intersecting 1-separating curves $\gamma^{\pm}$ such that $\gamma^+ \in \mathrm{Lk}(\alpha^-) \cap \mathrm{Lk}(\beta^-)$ and $\gamma^- \in \mathrm{Lk}(\alpha^+) \cap \mathrm{Lk}(\beta^+)$. This is because the annulus must cut the surface into two subsurfaces, each of which contains the disk with three punctures bounded by α^+ and α^-, respectively. Hence every side of the annulus contains at least three punctures, and we can choose some curve γ^+ that bounds the puncture inside the annulus and two other punctures on the side of α^+. The same holds for γ^-, as depicted in Figure 13. Now, since we can choose $\gamma^{\pm}$ to induce different puncture

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

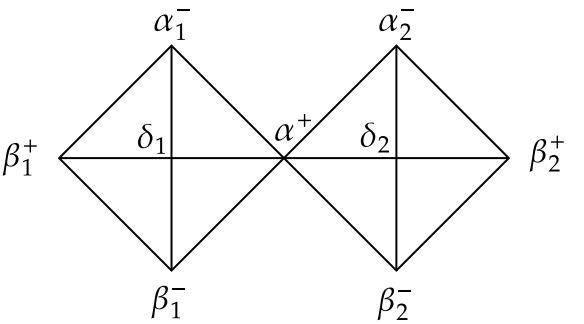


Figure 14. The two squares inside $\mathcal{C}^{ss}$ described in the proof of Lemma 6.15. We may independently choose a point for every column and find a path from β_1^+ to β_2^+ passing through those points. Every such path must pass through the star of s, i.e. must contain a vertex fixed by γ_s.

separations, their projections remain at distance at least 2 by Lemma 3.13, hence they contradict Condition 3 for $(\overline{\alpha}^{\pm}, \overline{\beta}^{\pm})$.

Now, let $P^{\pm}$ be the division induced by δ and filled by $(\alpha^{\pm}, \beta^{\pm})$. Clearly $\pi(P^{\pm}) \subseteq \overline{P}^{\pm}$, since a curve in $\mathrm{Lk}(\alpha^+) \cap \mathrm{Lk}(\beta^+)$ projects inside $\mathrm{Lk}(\overline{\alpha}^+) \cap \mathrm{Lk}(\overline{\beta}^+)$. Moreover by construction $\pi(P^{\pm}) = S(\pi(\delta))$, and in particular it is a maximal join. Hence $S(\pi(\delta)) = \overline{P}^{\pm}$ by maximality. □

Lemma 6.15. *For all large multiples* K*, the vertex* $\overline{\delta}$ *from Lemma 6.14 does not depend on the choice of the filling vertices* $(\overline{\alpha}^{\pm}, \overline{\beta}^{\pm})$*, nor on the choice of their lifts.*

The proof of this lemma will be prototypical of many arguments throughout the paper.

Proof. Choose two pairs of filling vertices $(\overline{\alpha}_1^{\pm}, \overline{\beta}_1^{\pm})$ and $(\overline{\alpha}_2^{\pm}, \overline{\beta}_2^{\pm})$. Without loss of generality we may assume that $(\overline{\alpha}_1^+, \overline{\beta}_1^+) = (\overline{\alpha}_2^+, \overline{\beta}_2^+)$, since we can replace one pair at a time. Let $(\alpha_1^{\pm}, \beta_1^{\pm})$ and $(\alpha_2^{\pm}, \beta_2^{\pm})$ two lifts forming two squares. The argument of Lemma 6.14, whose construction works for any choice of lifts of $(\overline{\alpha}_1^{\pm}, \overline{\beta}_1^{\pm})$, shows that there exists a unique curve $\delta_1 \in \mathcal{C}^{ss}$ which is disjoint from $\alpha_1^{\pm}$ and $\beta_1^{\pm}$, and we can similarly find δ_2. Up to elements in DT_K we may assume that $\alpha_1^+ = \alpha_2^+ = \alpha^+$, thus we are in the situation depicted in Figure 14.

Now, there exists $g \in DT_K$ such that $g(\beta_1^+) = \beta_2^+$. We want to prove that, up to changing the lift, we may "glue" β_1^+ to β_2^+, that is, we can find a lift of the two squares in Figure 14 such that $\beta_1^+ = \beta_2^+$. If g is not the identity let (s, γ_s) be as in Proposition 2.3. If $d_{\mathcal{C}}(s, \beta_1^+) \leqslant 1$ we may apply γ_s to all data and proceed by induction on the complexity of g. Otherwise $d_s(\beta_1^+, \beta_2^+) \geqslant \Theta$. Now, at least one of the cut sets $\{\alpha_1^-, \beta_1^-, \delta_1\}$, $\{\alpha^+\}$, $\{\alpha_2^-, \beta_2^-, \delta_2\}$ must be fixed pointwise by γ_s, because if DT_K is deep enough there cannot be a path from β_1^+ to β_2^+ of length 4 with no points in the star of s. In fact, for every fixed $n \in \mathbb{N}$ we can choose $\Theta \geqslant nB$, where B is the constant from the Theorem 2.4 (BGI). Therefore, whenever $p, q \in \mathcal{C}$ have large projections on some s, any piecewise geodesic path between p and q which is made of at most n geodesic subpaths must intersect the star of s, because

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

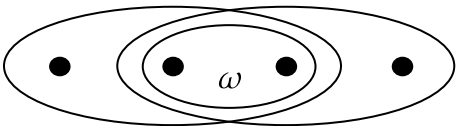


Figure 15. A surrounding pair and the minimal curve ω it surrounds.

otherwise we could find a geodesic subpath whose endpoints have projections at least B on s, thus violating BGI. In our case we can choose n to be 4. Therefore, we may apply γ_s to the lift "beyond" the points in the star of s (that is, to the connected component that contains β_2^+ of the complement of the cut set), and again proceed by induction on the complexity of g.

At the end of this procedure we have replaced δ_1 and δ_2 by some of their images via elements of DT_K. But now $(\alpha_1^+, \beta_1^+) = (\alpha_2^+, \beta_2^+)$, and therefore $\delta_1 = \delta_2$ since they are both characterised as the only boundary component of the surface filled by $\alpha^+ \cup \beta^+$ which lies in $\mathcal{C}^{ss}$. □

Now the proof of Theorem 6.12 is complete if we show that:

COROLLARY 6.16. *The map* S *is injective.*

Proof. We want to show that if $S(\overline{\delta}_1) = S(\overline{\delta}_2)$ then $\overline{\delta}_1 = \overline{\delta}_2$. Let δ_1 and δ_2 be two lifts of these vertices, let $(\alpha_1^\pm, \beta_1^\pm)$ and $(\alpha_2^\pm, \beta_2^\pm)$ be some filling curves for the lifts and let $(\overline{\alpha}_1^\pm, \overline{\beta}_1^\pm)$ and $(\overline{\alpha}_2^\pm, \overline{\beta}_2^\pm)$ be their projections, which induce the same division $\overline{P}^\pm$. Notice that, by construction, δ_1 is the curve obtained by applying the machinery of Lemma 6.14 to $(\overline{\alpha}_1^\pm, \overline{\beta}_1^\pm)$ with lifts $(\alpha_1^\pm, \beta_1^\pm)$, and similarly for δ_2. But the previous lemma shows that $\overline{\delta}_1 = \overline{\delta}_2$, since the class of the curve δ obtained by Lemma 6.14 does not depend on the choice of filled divisions for $\overline{P}^\pm$. □

7. From strongly separating to all curves

For this whole section, we follow the footsteps of [12], except we focus on the case of punctured spheres. We recall that, as in Definition 6.2, we denote by $\mathcal{C}^{ss}$ the full subgraph of $\mathcal{C}$ spanned by strongly separating curves, that is, all separating curves that do not bound twice-punctured disks. We want to prove that:

THEOREM 7.1. *For $b \geqslant 7$ and for all large multiples K, any automorphism of $\mathcal{C}^{ss}/DT_K$ extends to an automorphism of $\mathcal{C}/DT_K$.*

In [12], Bowditch identifies every *minimal curve* ω (meaning, a curve that bounds a twice-punctured disk) with a certain equivalence class of pairs of curves in $\mathcal{C}^{ss}$. We recall the following definitions from that paper:

DEFINITION 7.2. *Two curves $\alpha, \beta \in \mathcal{C}^{ss}$ form a* surrounding pair *if they bound three-punctured disks whose intersection is a twice-punctured disk. The boundary of the latter is the minimal curve* surrounded *by α and β. See Figure 15.*

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

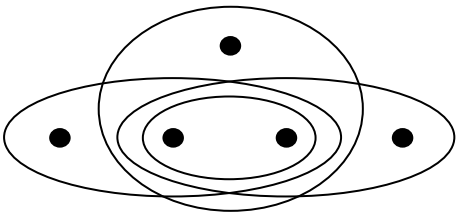


Figure 16. A surrounding triple.

DEFINITION 7.3. *Three curves $\alpha, \beta, \gamma \in \mathcal{C}^{ss}$ form a* surrounding triple *if any two of them form a surrounding pair that surrounds the same minimal curve, as in Figure 16.*

The following is a restatement of [12, Lemmas 4.2 and 4.3], which are core results of that paper:

LEMMA 7.4. *Any two surrounding pairs that surround the same ω are connected by a finite sequence of surrounding triples (i.e., by successively replacing one of the two curves with a third curve that forms a surrounding triple with the other two). Thus minimal curves correspond to equivalence classes of surrounding pairs, up to surrounding triples.*

Moreover, two minimal curves ω, ω' are disjoint if and only if there are two disjoint curves α surrounding ω and α' surrounding ω'. A minimal curve ω is disjoint from a curve β if and only if either β surrounds ω or there is some α surrounding ω and disjoint from β.

This implies that, whenever surrounding pairs and surrounding triples can be recognised inside $\mathcal{C}^{ss}$, any automorphism of $\mathcal{C}^{ss}$ extends to an automorphism of $\mathcal{C}$. Our goal is to repeat the same kind of argument: we want to define surrounding pairs and triples inside $\mathcal{C}^{ss}/DT_K$ just using combinatorial properties and then show that they actually correspond to projections of surrounding pairs and triples. Then we can rely on Lemma 7.4 to show that two surrounding pairs in $\mathcal{C}^{ss}/DT_K$ are connected by a finite sequence of surrounding triples, and we are almost done. As in [12] we distinguish three cases, whether the number of punctures is 7, 8 or at least 9.

For the rest of the section, K will always denote a large enough multiple, depending on the constants that will successively appear.

7.1. 7 punctures

We will say that an n-cycle inside a graph Γ is a cyclically ordered sequence of n vertices, where consecutive vertices are adjacent. The following Lemma, which summarises [12, Sections 3 and 4], shows that surrounding pairs and triples can be recognised inside $\mathcal{C}^{ss}(S_7)$ (which actually coincides with $\mathcal{C}^1(S_7)$) by using 7-cycles:

LEMMA 7.5. *In a seven-punctured sphere:*

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

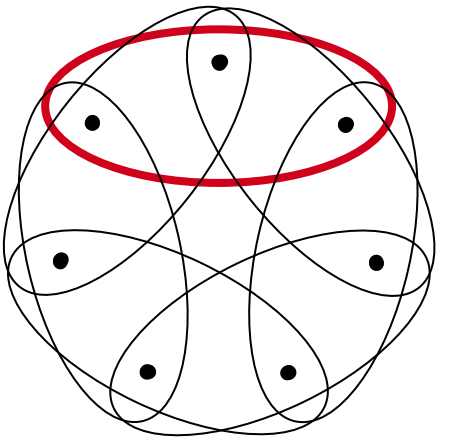


Figure 17. These curves form an isometrically embedded heptagon in $\mathcal{C}^1(S_7)$. The curves are all obtained from the same curve, say the red one at the top, by rotating. Notice that all curves induce different puncture separations.

- *Inside $\mathcal{C}^1(S_7)$ there is only one 7-cycle, or heptagon, up to the action of the mapping class group, and this heptagon is isometrically embedded;*
- *Two curves $\alpha, \beta \in \mathcal{C}^1$ form a surrounding pair if and only if they are at distance 2 in some heptagon;*
- *Three curves α, β, γ form a surrounding triple if any two of them form a surrounding pair and there is no curve in $\mathcal{C}^1$ that is disjoint from each of α, β, γ.*

Thus, a good definition for surrounding pairs inside $\mathcal{C}^{ss}/DT_K$ is the following:

DEFINITION 7.6. *Two vertices $\overline{\alpha}, \overline{\beta} \in \mathcal{C}^{ss}/DT_K$ form a surrounding pair if they are at distance 2 in some heptagon inside $\mathcal{C}^{ss}/DT_K$.*

It is clear that if $\overline{\alpha}, \overline{\beta}$ are a surrounding pair then we can find lifts α, β that are a surrounding pair in $\mathcal{C}^{ss}$, just by lifting the corresponding heptagon. Thus α, β surround some minimal curve ω, and we say that $\overline{\alpha}, \overline{\beta}$ surround its projection $\overline{\omega}$.

LEMMA 7.7. *The vertex $\overline{\omega}$ surrounded by $\overline{\alpha}, \overline{\beta}$ is well-defined, meaning it does not depend on the chosen heptagon, nor on its lift.*

Proof. Let $\overline{T}, \overline{T}' \subset \mathcal{C}^{ss}/DT_K$ be two heptagons in which $\overline{\alpha}, \overline{\beta}$ are at distance 2, and let T, T' be two lifts of $\overline{T}, \overline{T}'$. Up to elements of DT_K we can assume that the lifts of $\overline{\alpha}$ coincide, and let β, β' be the lifts of $\overline{\beta}$. Let ω, ω' be the minimal curves surrounded by the two pairs. Let $g \in DT_K$ be such that $g(\beta) = \beta'$. Then the picture in $\mathcal{C}$ is as in Figure 18.

Let (s, γ_s) be as in Proposition 2.3. We now show that we can apply γ_s to part of our diagram without breaking the heptagons, and ensuring that ω and ω' remain surrounded by the corresponding pairs. If this is true then we can proceed

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

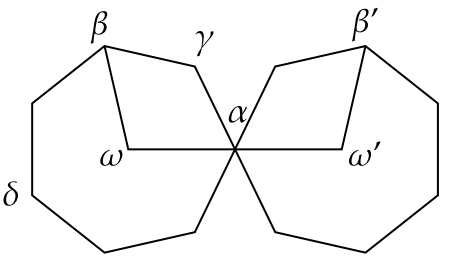


Figure 18. The two heptagons from the proof of Lemma 7.7.

by induction on the complexity of g to glue β to β', and in the end ω and ω' will both be the unique curve surrounded by α and β. We can argue as in the proof of Lemma 6.15 to show that γ_s must fix one of the following:

- β;
- ω and some curves $\gamma, \delta \in T$ on the two paths in T from β to α;
- α;
- ω' and some curves $\gamma', \delta' \in T'$ on the two paths in T' from β' to α'.

In all cases, we can apply γ_s beyond the cut set and proceed by induction. We just need to be careful that in the second case ω remains the curve surrounded by β and $\gamma_s(\alpha)$. This is true, since the curve surrounded by a pair is characterised as the only minimal curve inside both of the three-punctured disks defined by the surrounding curves. Thus it suffices to notice that $\omega = \gamma_s(\omega)$ is still inside the disk surrounded by $\gamma_s(\alpha)$. The exact same argument shows that, in the fourth case, ω' remains the curve surrounded by α' and $\gamma_s(\beta')$. □

This shows that every surrounding pair in $\mathcal{C}^{ss}/DT_K$ corresponds to some minimal $\overline{\omega}$. Conversely, given some $\overline{\omega}$ it is easy to find some $\overline{\alpha}, \overline{\beta}$ that surround it: just pick a lift ω, find a surrounding pair and let T be the corresponding heptagon. Since by Lemma 7.5 there is a unique heptagon in $\mathcal{C}^1(S_7)$ up to the action of the mapping class group, T looks like in Figure 17, and therefore all curves of T induce different puncture separations. Thus, by Lemma 3.11, the projection is injective when restricted to T, and therefore $\pi(T)$ is a 7-cycle.

Now we want to define surrounding triples, in order to characterise the class of pairs that surround the same $\overline{\omega}$. Again, our definition is inspired by Lemma 7.5.

Definition 7.8. *Three vertices $\overline{\alpha}, \overline{\beta}, \overline{\gamma} \in \mathcal{C}^{ss}/DT_K$ form a surrounding triple if they are pairwise surrounding pairs and there is no $\overline{\delta} \in \mathcal{C}^{ss}/DT_K$ which is adjacent to each of them.*

Lemma 7.9. *A surrounding triple $\alpha, \beta, \gamma \in \mathcal{C}^{ss}$ projects to a surrounding triple $\overline{\alpha}, \overline{\beta}, \overline{\gamma} \in \mathcal{C}^{ss}/DT_K$.*

Proof. We already know that surrounding pairs project to surrounding pairs. Now let $\overline{\delta} \in \mathcal{C}^{ss}/DT_K$ and let δ be one of its lifts. We want to show that $\overline{\delta}$ is not adjacent to each of $\overline{\alpha}, \overline{\beta}, \overline{\gamma}$. Since α, β, γ is a surrounding triple we have that δ must

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

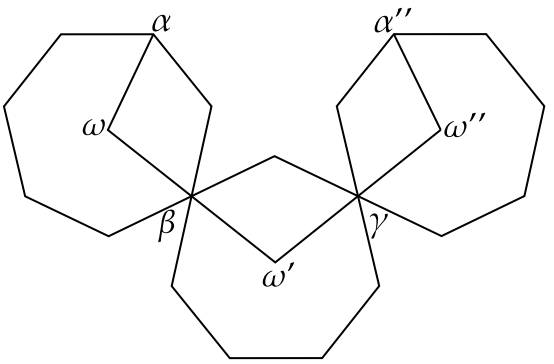


Figure 19. The three heptagons from the proof.

intersect at least one of the three curves, say α. If δ and α induce different puncture separations then $d_{\mathcal{C}^{ss}/DT_K}(\overline{\alpha}, \overline{\delta}) \geqslant 2$ by Lemma 3.13. Otherwise δ and α surround the same three punctures, and in particular δ and β must intersect with different puncture separations. Hence $d_{\mathcal{C}^{ss}/DT_K}(\overline{\beta}, \overline{\delta}) \geqslant 2$ for the same reason. □

Corollary 7.10. *Any two surrounding pairs for the same $\overline{\omega}$ are connected by a finite number of surrounding triples.*

Proof. Just lift the two surrounding pairs in such a way that they surround the same lift ω. Then the conclusion follows from Lemmas 7.5 and 7.9. □

Moving on, we need to show that the vertex $\overline{\omega}$ surrounded by a surrounding triple is well-defined, in order to identify $\overline{\omega}$ with an equivalence class of surrounding pairs up to surrounding triples.

Lemma 7.11. *If $\overline{\alpha}, \overline{\beta}, \overline{\gamma} \in \mathcal{C}^{ss}/DT_K$ form a surrounding triple then they pairwise surround the same $\overline{\omega}$. Therefore, if two pairs are connected by a sequence of surrounding triples, then they surround the same $\overline{\omega}$.*

Proof. Let $\overline{T}, \overline{T}', \overline{T}''$ be three heptagons for the three pairs. Lift them to three heptagons T, T', T'' in such a way that T and T' share some lift β and T' and T'' share some lift γ. Let $\alpha \in T$ and $\alpha'' \in T''$ be the lift of $\overline{\alpha}$, and let $\omega, \omega' \omega''$ be the three curves surrounded by the three pairs, as in Figure 19.

Arguing as in Lemma 7.7 one can glue α to α'' while preserving the fact that ω is the curve surrounded by α and β, and similarly for ω' and ω''. Now α, β, γ are pairwise surrounding pairs. Moreover they cannot lie in the link of some other curve $\delta \in \mathcal{C}^{ss}$, because otherwise their projections would be at distance 1 from $\overline{\delta}$. This implies that α, β, γ form a surrounding triple, and therefore $\omega = \omega' = \omega''$. □

Corollary 7.12. *There is a bijective correspondence between minimal vertices $\overline{\omega} \in \mathcal{C}/DT_K$ and equivalence classes of surrounding pairs, up to surrounding triples.*

The only thing left to do is the following straightforward observation, that shows that automorphisms of $\mathcal{C}^{ss}/DT_K$ preserve disjointness between equivalence classes of surrounding pairs:

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

LEMMA 7.13. *Two minimal vertices $\overline{\omega}, \overline{\omega}'$ are disjoint if and only if there are two disjoint vertices $\overline{\alpha}$ surrounding $\overline{\omega}$ and $\overline{\alpha}'$ surrounding $\overline{\omega}'$. A minimal vertex $\overline{\omega}$ is disjoint from a vertex $\overline{\beta}$ if and only if either $\overline{\beta}$ surrounds $\overline{\omega}$ or there is some $\overline{\alpha}$ surrounding $\overline{\omega}$ and disjoint from $\overline{\beta}$.*

We now know how to intrinsically recover $\mathcal{C}/DT_K$ from $\mathcal{C}^{ss}/DT_K$. Thus we get the following, which is the $b=7$ case of Theorem 7.1:

COROLLARY 7.14. *If $b = 7$, for all large multiples K every automorphism of $\mathcal{C}^{ss}/DT_K$ extends to an automorphism of $\mathcal{C}/DT_K$.*

7.2. At least 9 punctures

We move on to the general case in which our sphere has at least 9 punctures, temporarily ignoring the case S_8 which, as will become clear, is different in nature. First we need to identify the vertices corresponding to classes of 1-separating curves inside $\mathcal{C}^{ss}/DT_K$.

Recall that we denote by G^* the dual of a graph G.

DEFINITION 7.15. *Let Γ be (a subgraph of) either $\mathcal{C}$ or $\mathcal{C}/DT_K$. A vertex x* splits *Γ if $(Lk_\Gamma(x))^*$ is connected.*

Notice that we can recognise $\mathcal{C}^1$ inside $\mathcal{C}^{ss}$. More precisely, if $\delta \in \mathcal{C}^{ss}$ then $\delta \in \mathcal{C}^1$ if and only if δ does not split $\mathcal{C}^{ss}$. Now we want to establish the same result for $\mathcal{C}^{ss}/DT_K$.

LEMMA 7.16. *A vertex $\overline{\delta} \in \mathcal{C}^{ss}/DT_K$ belongs to $\mathcal{C}^1/DT_K$ if and only if it does not split $\mathcal{C}^{ss}/DT_K$.*

Lemma 7.16 is implied by the following:

LEMMA 7.17. *Let $\delta \in \mathcal{C}^{ss}$ be a curve and let $\alpha, \beta \in Lk(\delta)$. Let $\overline{\alpha}, \overline{\beta}, \overline{\delta}$ be their projections. Then δ separates α and β if and only if $\overline{\delta}$ separates $\overline{\alpha}$ and $\overline{\beta}$.*

Proof. Suppose that $\overline{\alpha}, \overline{\beta}$ are not separated by $\overline{\delta}$. Pick a chain $\overline{\gamma}_0 = \overline{\alpha}, \overline{\gamma}_1, \dots, \overline{\gamma}_k = \overline{\beta} \in \mathrm{Lk}(\overline{\delta})$ and choose some lifts $\gamma_0 = \alpha, \gamma_1, \dots, \gamma_k = \beta \in \mathrm{Lk}(\delta)$. Notice that, since the projection map is 1-Lipschitz, these curves form a chain themselves, which means that δ does not separate α and β.

Conversely, suppose that α, β lie on the same subsurface Σ cut out by δ. This subsurface must contain one of the disks that α cuts out, call it D_α, but cannot coincide with it. Therefore, there must be some puncture that belongs to Σ but not to D_α. The same argument works for β and some disk D_β. Hence it is always possible to find a strongly separating curve $\gamma \in \mathrm{Lk}(\delta)$ which intersects both α and β and induces a different puncture separation. More precisely, one can always find some disk $D_\gamma \subset \Sigma$ with at least three punctures and which contains (at least) one of the punctures not in D_α and (at least) one of the punctures not in D_β, as in Figure 20. Then $\overline{\gamma}$ must be at distance at least 2 from both $\overline{\alpha}$ and $\overline{\beta}$ by Lemma 3.11, and therefore $\overline{\delta}$ does not separate $\overline{\alpha}$ and $\overline{\beta}$. □

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

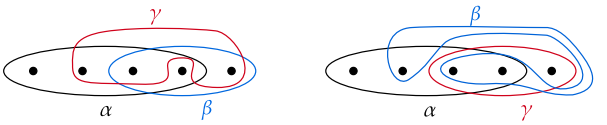


Figure 20. The shape of a possible curve γ depends on whether one of the disks contains the punctures of the other.

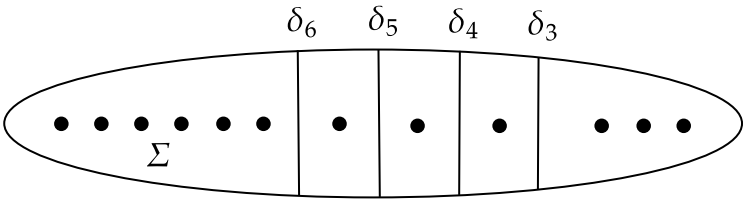


Figure 21. Starting from δ_3, Δ is constructed by consecutively cutting out once-punctured annuli.

Now we can exploit the previous results to intrinsically determine if two curves belong to a *peripheral* S_7, that is, an S_7 cut out by a single curve. Here is where we need the number of punctures to be at least 9, since in S_8 there is no strongly separating curve that cuts out some S_7.

LEMMA 7.18. (Bizarre simplices) *Let $b \geqslant 9$. If $\Delta = (\delta_3, \dots \delta_{b-6}) \subseteq \mathcal{C}^{ss}$ is an ordered simplex of dimension $b-9$ such that:*

(1) *$\delta_3 \in \mathcal{C}^1$;*
(2) *every δ_i separates δ_{i-1} from δ_{i+1};*
(3) *there is no strongly separating curve that separates two consecutive curves δ_i and δ_{i+1};*

then $Lk(\Delta)$ fills a peripheral S_7 cut out by δ_{b-6}.

Proof. The second property tells us that Δ cuts the surface into two disks and some punctured annuli between consecutive curves, as in Figure 21. Moreover each of these annuli should contain just one puncture, because otherwise we could find some curve that contradicts the third property. Then, since the $S_{0,4}$ cut out by δ_3 cannot contain any strongly separating curve, $\mathrm{Lk}(\Delta)$ must fill the other peripheral disk, call it Σ. Now it is enough to count the punctures in Σ, that must be six (δ_3 cuts out three punctures, and each of the $b-9$ annuli contains a single puncture). □

Notice that these bizarre properties are stated only in terms of some vertices separating some others, and are therefore preserved when passing to the quotient and taking lifts by Lemma 7.17. Hence we get the following:

COROLLARY 7.19. *Let $b \geqslant 9$. If $\overline{\Delta} = (\overline{\delta}_3, \dots \overline{\delta}_{b-6}) \subseteq \mathcal{C}^{ss}/DT_K$ is an ordered simplex of dimension $b-9$ such that:*

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

(1) $\overline{\delta}_3 \in \mathcal{C}^1/DT_K$;
(2) *every* $\overline{\delta}_i$ *separates* $\overline{\delta}_{i-1}$ *from* $\overline{\delta}_{i+1}$;
(3) *there is no* $\overline{\gamma} \in \mathcal{C}^{ss}/DT_K$ *that separates two consecutive* $\overline{\delta}_i$ *and* $\overline{\delta}_{i+1}$;

then every vertex $\overline{\alpha} \in Lk(\overline{\Delta})$ *lifts inside a peripheral* S_7 *cut out by some lift* δ_{b-6} *of* $\overline{\delta}_{b-6}$.

Now, given some peripheral S_7 and a bizarre simplex Δ for it, we are able to recognise $\mathcal{C}^{ss}(S_7)$, which is the subgraph of $\mathcal{C}^{ss}$ spanned by those curves $\gamma \in \mathrm{Lk}(\Delta)$ which do not cut out a pair of pants with the boundary of S_7, or equivalently for which there exists some other curve $\gamma' \in \mathrm{Lk}(\Delta)$ that separates γ from δ_{b-6}. Again, this property just involves some vertices separating some others, thus it holds in the quotient if and only if it holds in $\mathcal{C}^{ss}$. Therefore, if $\pi : \mathcal{C}^{ss} \to \mathcal{C}^{ss}/DT_K$ is the quotient projection, we can recognise $\pi(\mathcal{C}^{ss}(S_7))$ inside $\mathcal{C}^{ss}/DT_K$ in the same way. Notice that, as discussed in Corollary 3.2, $\pi(\mathcal{C}^{ss}(S_7))$ is isomorphic to $\mathcal{C}^{ss}(S_7)/DT_K(S_7)$, therefore we are exactly in the framework of the previous subsection. Then we can define surrounding pairs and triples as in Figures 15 and 16. The following is a summary of [12, Lemmas 6.1 and 6.3]:

LEMMA 7.20. *Let* $b \geqslant 9$. *Two curves* $\alpha, \beta \in \mathcal{C}^{ss}$ *form a surrounding pair if and only if:*

- *They belong to* $\mathcal{C}^1$;
- *They belong to some peripheral* S_7;
- *They are a surrounding pair intrinsically inside* S_7.

Three curves $\alpha, \beta, \gamma \in \mathcal{C}^{ss}$ *form a surrounding triple if they are pairwise surrounding pairs and they all belong to some peripheral* S_6 *(which is defined exactly as a peripheral* S_7*, but this time the bizarre simplex has one more vertex).*

Now we can almost repeat what we did for the case $b = 7$, being careful to remember in which link we are. All proofs will have the same flavour as the corresponding ones, though they will involve more cut set arguments.

- First notice that, if $\overline{T}$ is a heptagon which lies in $\mathrm{Lk}_{\mathcal{C}^{ss}/DT_K}(\overline{\Delta})$ for some simplex $\overline{\Delta}$, then $\overline{T} \star \overline{\Delta}$ is a generalised heptagon, thus we can find lifts Δ and $T \subset \mathrm{Lk}_{\mathcal{C}^{ss}}(\Delta)$ by Theorem 2.7.
- To prove the equivalent of Lemma 7.7 we must show that the vertex $\overline{\omega}$ surrounded by a pair is also independent of the chosen simplex $\overline{\Delta}$. Thus we take simplices $\overline{\Delta}, \overline{\Delta}'$ and heptagons $\overline{T}, \overline{T}'$ inside the respective links, and lift them to $T \in \mathrm{Lk}(\Delta)$ and $T' \in \mathrm{Lk}(\Delta')$. Then the proof follows the same steps, being careful to add Δ to the cut set containing ω (and similarly for Δ').
- It is clear that surrounding pairs (triples) project to surrounding pairs (triples) since the bizarre properties are preserved in the quotient. Then Lemma 7.9 follows.
- Lemma 7.13 deals with disjointness, and the same statement works in our case.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

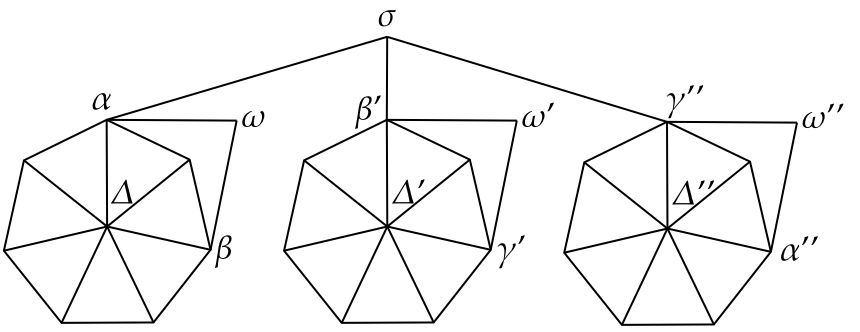


Figure 22. The various lifts from the proof of Lemma 7.21.

We are left to prove an analogous to Lemma 7.11, which needs a bit more care.

LEMMA 7.21. *Let $b \geqslant 9$. If $\overline{\alpha}, \overline{\beta}, \overline{\gamma} \in \mathcal{C}^{ss}/DT_K$ form a surrounding triple then they pairwise surround the same $\overline{\omega}$.*

Proof. As in the proof of Lemma 7.11 we actually want to show that a surrounding triple in $\mathcal{C}^{ss}/DT_K$ lifts to a surrounding triple in $\mathcal{C}^{ss}$. Since $\overline{\alpha}, \overline{\beta}$ are a surrounding pair they belong to some heptagon $\overline{T}$ inside the link of some simplex $\overline{\Delta}$, which corresponds to some peripheral S_7 (meaning that it lifts to a pants decomposition for the complement of a peripheral S_7). We can similarly find $\overline{T}', \overline{\Delta}'$ for $\overline{\beta}, \overline{\gamma}$ and $\overline{T}'', \overline{\Delta}''$ for $\overline{\alpha}, \overline{\gamma}$. Moreover, let $\overline{\sigma}$ be a simplex which corresponds to some peripheral S_6 and whose link contains $\overline{\alpha}, \overline{\beta}, \overline{\gamma}$. Now, choose lifts $\Delta, \Delta', \Delta'', \sigma$ for the various simplices. Lift $\overline{T}$ to a heptagon T inside $\mathrm{Lk}(\Delta)$, and let $\alpha \in T$ be the lift of $\overline{\alpha}$. If $\alpha \notin \mathrm{Lk}(\sigma)$ let $g \in DT_K$ be such that $g(\alpha) \in \mathrm{Lk}(\sigma)$, and replace T and Δ with $g(T)$ and $g(\Delta)$. Apply the same procedure to T', Δ' and T'', Δ''. Finally, let ω be the curve surrounded by α, β, and define similarly ω' and ω''. The situation is depicted in Figure 22.

Now let $\beta \in T$ and $\beta' \in T'$ be the corresponding lifts of $\overline{\beta}$, which we want to glue to each other. Let $g \in DT_K$ be an element mapping β to β'. If g is not the identity let (s, γ_s) be as in Proposition 2.3. Arguing as in the proof of Lemma 6.15 we see that γ_s must fix pointwise one of the following:

- β;
- ω, two curves δ, ε in T and the whole Δ;
- α;
- σ;

Then, we can apply γ_s to everything beyond the cut set, while preserving that all ωs are still the minimal curves surrounded by the corresponding pairs as in Lemma 7.7, and proceed by induction.

Thus we can glue β to β' and with the exact same argument we can glue γ' to γ''. Therefore, our data up to this point span a graph which schematically looks like in Figure 23.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

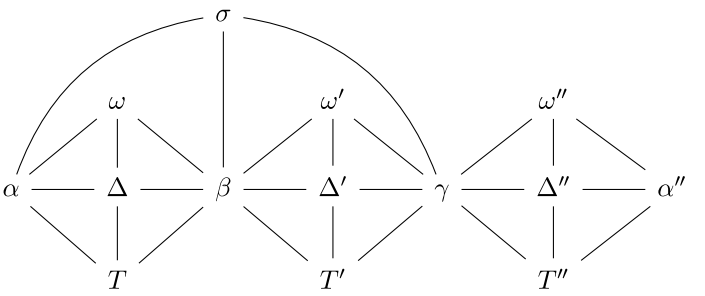


Figure 23. This graph represents all possible "paths" from α to α'', and elements on the same column correspond to the same cut set. Notice that every heptagon actually represents two paths.

We are left to glue α to α'', since then α, β, γ will be a surrounding triple and therefore surround the same minimal curve, which in turn will mean that $\omega = \omega' = \omega''$. Let $g \in DT_K$ be an element mapping α to α''. If g is not the identity let (s, γ_s) be as in Proposition 2.3. Then γ_s must fix pointwise one of the following cut sets:

- α;
- σ, ω, two curves δ, ε in T and the whole Δ;
- σ and β;
- σ, ω', two curves δ', ε' in T' and the whole Δ';
- γ;
- ω'', two curves δ'', ε'' in T'' and the whole Δ''.

Thus, as argued in Lemma 7.7, we can apply γ_s beyond the cut set, while preserving that every minimal curve is the one surrounded by the corresponding pair, and conclude by induction. □

To sum up we get the following, which is the case $b \geqslant 9$ of Theorem 7.1:

Corollary 7.22. *For every $b \geqslant 9$ and all large multiples* K*, every automorphism of $\mathcal{C}^{ss}/DT_K$ extends to an automorphism of $\mathcal{C}/DT_K$.*

7.3. 8 punctures

We are left to deal with the case $b = 8$. The idea will be to replace heptagons with a suitable subgraph which we will use to recognise surrounding pairs. Define $\mathfrak{O}$ as the graph obtained by adding the four longest diagonal to an octagon, as in Figure 24. A possible realisation of this graph inside $\mathcal{C}^1(S_8)$ is given by the eight curves in the same Figure, and [12, Lemma 7.2] shows that there is only one copy of $\mathfrak{O}$ inside $\mathcal{C}^1(S_8)$ up to the action of the mapping class group (that is, every $\mathfrak{O}$ corresponds to some curves arranged as in the Figure).

In [12] it was proved that $\alpha, \beta \in \mathcal{C}^{ss}$ form a surrounding pair if and only if:

- They belong to $\mathcal{C}^1$;
- They belong to some isometrically embedded copy of $\mathfrak{O}$ inside $\mathcal{C}^1$;

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

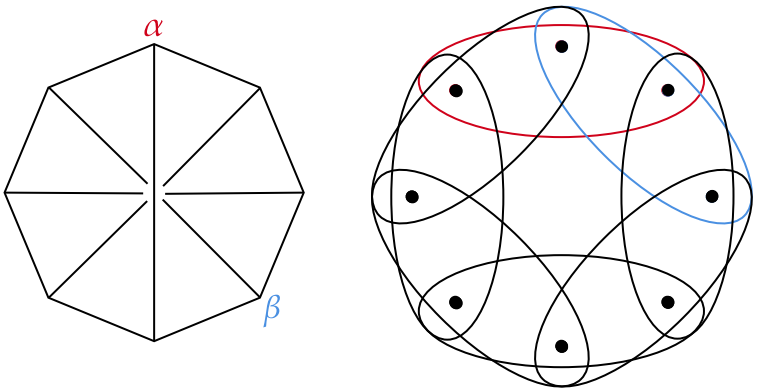


Figure 24. The graph $\mathfrak{O}$ and its realisation with 1-separating curves, each of which surrounds three consecutive punctures. Notice that two vertices correspond to a surrounding pair if and only if they are connected by exactly two paths of length 2 inside $\mathfrak{O}$.

- They are connected by exactly two geodesic paths of length 2 inside $\mathfrak{O}$.

Moreover, three curves $\alpha, \beta, \gamma \in \mathcal{C}^{ss}$ form a surrounding triple if and only if they are pairwise surrounding pairs and there is no $\delta \in \mathcal{C}^{ss} \setminus \mathcal{C}^1$ which is disjoint from all of them.

Now, in order to define surrounding pairs and triples inside $\mathcal{C}^{ss}/DT_K$ we need the following lemmas:

LEMMA 7.23. *For all large multiples* K *every isometrically embedded copy of* $\mathfrak{O}$ *inside* $\mathcal{C}$ *projects isometrically into* $\mathcal{C}^1/DT_K$. *In particular there exists an isometrically embedded copy of* $\mathfrak{O}$ *inside* $\mathcal{C}^1/DT_K$.

Proof. Every two copies of $\mathfrak{O}$ differ by a mapping class, thus it suffices to argue as in Lemma 3.7 with X_b replaced by $\mathfrak{O}$. □

LEMMA 7.24. *Every isometrically embedded copy* $\overline{\mathfrak{O}} \subset \mathcal{C}^1/DT_K$ *admits an isometrically embedded lift* $\mathfrak{O} \subset \mathcal{C}^1$. *Therefore, surrounding pairs lift to surrounding pairs.*

Proof. It suffices to find a lift $\mathfrak{O}$, which will automatically be isometrically embedded since the projection map is 1-Lipschitz (as argued in Lemma 2.16). Now, it is useful to see $\mathfrak{O}$ as the 1-skeleton of a Möbius band made of four squares, as in Figure 25.

Clearly the graph G in the Figure admits a lift, since we can lift each square and glue them together along common sides (recall that every 1-simplex admits a unique DT_K-orbit of lifts). Now we are left to glue 1 to $1'$ and 5 to $5'$. Let $g \in DT_K$ be an element that maps the edge $1, 5$ to the edge $1', 5'$. If g is not the identity let (s, γ_s) be as in Lemma 2.3, applied to $x = 1$. If $d_{\mathcal{C}}(s, 1) \leqslant 1$ we can apply γ_s to the whole data and proceed by induction. Otherwise every path from 1 to $1'$ must intersect the star of s. This means that there must be a square Q where it is not possible to move from the left side to the right side without crossing the star of s.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

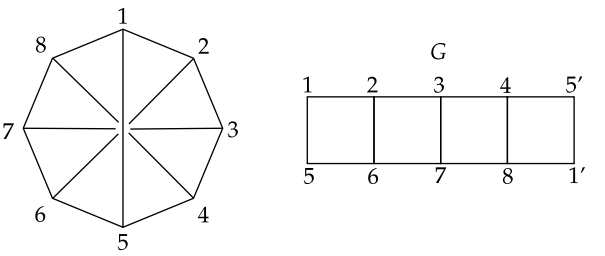


Figure 25. $\mathfrak{O}$ is obtained from the strip G by gluing 1 to $1'$ and 5 to $5'$.

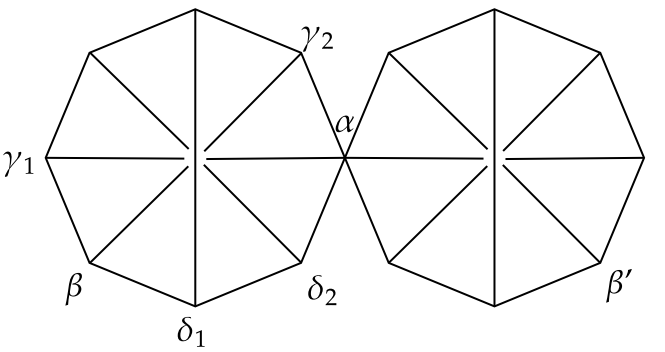


Figure 26. The two octagons from the proof of Lemma 7.25.

Thus there must be at least two vertices $p, q \in Q$ which lie in the star of s and that cut G in two connected components (more precisely, p, q must be the vertices of either a diagonal or a vertical side). Thus, we can apply γ_s beyond p and q. Notice that either $5'$ is one of p and q or $5'$ and $1'$ are on the same connected component cut out by p and q; either way γ_s is applied to the whole edge $1', 5'$, and we can proceed by induction. □

Now we can proceed exactly as for the case $b = 7$, with a few adjustments to cut sets arguments. Given a surrounding pair $\overline{\alpha}, \overline{\beta}$, define the vertex $\overline{\omega}$ surrounded by the pair by lifting the pair and projecting the curve surrounded by the lift, as before.

LEMMA 7.25. *The vertex $\overline{\omega}$ surrounded by a pair $\overline{\alpha}, \overline{\beta}$ is well-defined.*

Proof. Let $\overline{\mathfrak{O}}$ and $\overline{\mathfrak{O}}'$ be two copies of $\mathfrak{O}$ inside $\mathcal{C}^{ss}/DT_K$ that contain $\overline{\alpha}, \overline{\beta}$, and lift them to $\mathfrak{O}$ and $\mathfrak{O}'$. We can assume that the lifts of $\overline{\alpha}$ coincide, and let β, β' be the lift of $\overline{\beta}$. Let $\gamma_1, \gamma_2, \delta_1, \delta_2$ be the curves in Figure 26.

Let ω, ω' be the curves surrounded by the two pairs, and let $g \in DT_K$ be an element that maps β to β'. If g is not the identity let (s, γ_s) be as in Proposition 2.3. If $d_{\mathcal{C}}(\beta, s) \leqslant 1$ we can apply γ_s to the whole data and proceed by induction on the complexity. If $d_{\mathcal{C}}(\alpha, s) \leqslant 1$ we can apply γ_s just to the second octagon and again proceed by induction. If none of the previous hold we can assume without loss of generality that $d_s(\alpha, \beta)$ is large, and therefore every path from β to α must pass

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

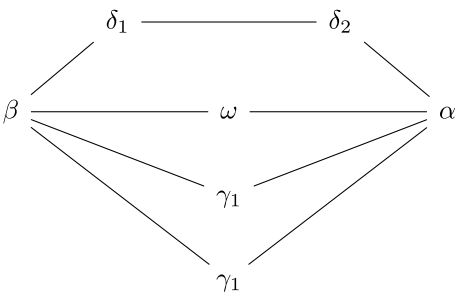


Figure 27. A schematic representation of $\{\omega\} \cup \mathrm{Lk}_{\mathfrak{O}}(\alpha) \cup \mathrm{Lk}_{\mathfrak{O}}(\beta)$ from Lemma 7.25.

through the star of s. In order to understand the following cut set argument we represent ω and part of the octagon as in Figure 27.

We already see that γ_s must fix $\omega, \gamma_1, \gamma_2$ and one between δ_1 and δ_2. In the first case γ_s fixes $\mathrm{Lk}_{\mathfrak{O}}(\beta) = \{\gamma_1, \gamma_2, \delta_1\}$, which is a cut set for $\mathfrak{O}$ because it is the link of a vertex. Thus, we can apply γ_s to all curves but β. In the second case γ_s fixes $\mathrm{Lk}_{\mathfrak{O}}(\alpha) = \{\gamma_1, \gamma_2, \delta_2\}$, and we can apply γ_s to $\mathfrak{O}' \cup \mathrm{Lk}(\alpha)$. In both cases ω remains the curve surrounded by β and $\gamma_s(\alpha)$, as argued in Lemma 7.7; therefore we can proceed by induction. □

The rest of the argument is as for the case $b = 7$.

- Arguing exactly as in Lemma 7.9 we see that surrounding triples project to surrounding triples.
- Adapting the proof of Lemma 7.11 with the cut set arguments from Lemma 7.25 we get that the vertex $\overline{\omega}$ surrounded by a surrounding triple is well-defined.
- Again, the conclusion of Lemma 7.13 holds, so we can recognise disjointness.

Thus we get the final piece of the proof of Theorem 7.1:

COROLLARY 7.26. *If* b = 8*, for all large multiples* K *every automorphism of* $\mathcal{C}^{ss}/DT_K$ *extends to an automorphism of* $\mathcal{C}/DT_K$.

8. Extracting combinatorial data from a quasi-isometry

The main goal of this section is, roughly speaking, to show that quasi-isometries of $MCG(S_b)/DT_K$ induce automorphisms of a certain graph, which we will later on relate to $\mathcal{C}/DT_K$. This strategy is inspired by [7, Section 5], where an approach to the quasi-isometric rigidity of mapping class groups is presented as a template to prove quasi-isometric rigidity of other *hierarchically hyperbolic spaces* (HHS), $MCG(S_b)/DT_K$ being a HHS (for suitable K) [8]. More precisely, [7, Theorem 5.7] states that quasi-isometries of a HHS satisfying three additional assumptions induce automorphisms of a suitable graph. The assumptions are not satisfied by $MCG(S_b)/DT_K$ but nevertheless we will show that a very similar statement applies to that case, see Theorem 8.7, which is the main result of this section. This theorem

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

in fact applies to any HHS satisfying certain assumptions (different from those of [7]), other examples of such HHSs being pants graphs.

We will begin by recalling all facts about HHSs that we will need.

8.1. HHS background

A hierarchically hyperbolic space is a metric space X that comes with certain additional data, most importantly a family of uniformly hyperbolic spaces $\{\mathcal{C}(Y)\}_{Y\in\mathfrak{S}}$, called *coordinate spaces*, and uniformly coarsely Lipschitz maps $\pi_Y : X \to \mathcal{C}(Y)$. For mapping class groups, these are curve graphs of subsurfaces and maps coming from subsurface projections. Moreover, the index set $\mathfrak{S}$ of the set of hyperbolic spaces has two relations, *orthogonality* and *nesting*, denoted $\perp$ and $\sqsubseteq$ respectively. For mapping class groups, these are containment and disjointness of subsurfaces (up to isotopy). When two elements U, V of $\mathfrak{S}$ are not $\sqsubseteq$- nor $\perp$-compatible, one says that they are *transverse*. It is customary to denote a HHS simply by $(X, \mathfrak{S})$.

A crucial fact about HHSs is that they satisfy a *distance formula*. This means that for all sufficiently large threshold s there exists D such that for all $x, y \in X$ we have

$$d_X(x,y) \asymp_{D,D} \sum \{\{d_Y(x,y)\}\}_s,$$

where

- $\asymp_{D,D}$ denotes equality up to multiplicative and additive error at most D,
- $\{\{A\}\}_s = A$ if $A \geqslant s$ and $\{\{A\}\}_s = 0$ otherwise,
- $d_Y(x,y)$ is shorthand for $d_{\mathcal{C}(Y)}(\pi_Y(x), \pi_Y(y))$.

The idea of orthogonality is that it corresponds to products, in the following sense. Given any $U \in \mathfrak{S}$, there is a corresponding space F_U associated to it, which is quasi-isometrically embedded in X (F_U is a HHS itself with index set $\mathfrak{S}_U = \{Y \in \mathfrak{S} : Y \sqsubseteq U\}$, and in mapping class groups these essentially correspond to mapping class groups of subsurfaces). Given a maximal set $\{U_i\}$ of pairwise orthogonal elements of $\mathfrak{S}$, there is a corresponding *standard product region* $P_{\{U_i\}}$ which is quasi-isometric to the product of the F_{U_i} (think of a Dehn twist flat as, coarsely, a product of annular curve graphs). The product region $P_{\{U_i\}}$ has the property that for all U with $U_{\{U_i\}}$ or $U \pitchfork U_i$ for some i, we have that $\pi_Y(P_{\{U_i\}})$ is a uniformly bounded set, that we denote $\rho_V^{\{U_i\}}$ (usually, this notation is used for a different bounded set at finite Hausdorff distance from the one defined here; this will not matter for us). Product regions can also be associated to non-maximal sets of pairwise orthogonal elements of $\mathfrak{S}$, but in this case there is an additional factor than the F_{U_i}; this will not be important for us.

The following is [7, Corollary 1.28] (which is stated for product regions corresponding to one element of the index set, but the proof does not use this):

LEMMA 8.1. *For all sufficiently large* s *there exists* D *such that, if* $\{U_i\}$ *is a maximal collection of pairwise orthogonal indices and* $x \in X$*, then*

$$d_X(x, P_{\{U_i\}}) \asymp_{D,D} \sum_W \left\{\left\{d_W\left(x, \rho_W^{\{U_i\}}\right)\right\}\right\}_s$$

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

where the sum is made on all W *such that* $U_{\{W_i\}}$ *or* $U_i \pitchfork W$ *for some* $\mathrm{U_i}$.

COROLLARY 8.2. Intersection of different product regions *If* $\{U_i\}$ *and* $\{V_j\}$ *are maximal collections of pairwise orthogonal indices, such that* $\{U_i\} \cap \{V_j\} = \emptyset$ *and every* $\mathrm{V_j}$ *is* $\sqsubseteq$*-minimal, then* $P_{\{U_i\}} \tilde{\cap} P_{\{V_j\}}$ *is coarsely a point.*

Proof. It is enough to apply Lemma 8.1 with $x \in P_{\{V_j\}}$. In fact, every $W \notin \{V_j\}$ contributes to the sum with a term that does not depend on x, but only on $\{U_i, V_j\}$ and W. Moreover, there is a finite number of terms that contribute to the sum, so if we choose a big enough threshold s we get that

$$d_X(x, P_{\{U_i\}}) \asymp_{D,D} \sum_j \left\{ \left\{ d_{V_j} \left(x, \rho_{V_j}^{\{U_i\}} \right) \right\} \right\}_s .$$

Therefore, all points $x \in P_{\{V_j\}}$ that lie within a given constant of $P_{\{U_i\}}$ have nearby projections to all $\mathcal{C}V_j$, and hence they are close to each other, as required. □

Finally, we will be interested in quasiflats in HHSs and, related to quasiflats, orthants. In X and the F_Us there are special (uniform) quasi-geodesics, called *hierarchy paths*, which are those that project monotonically (with uniform constants) to all $\mathcal{C}(Y)$. Similarly, there are hierarchy rays and hierarchy lines (which are quasi-geodesic rays and lines, respectively). Given a product region $P_{\{U_i\}}$ and hierarchy lines (resp., rays) in some of the F_{U_i}, the product region contains a product of those (where for the indices i for which a line/ray has not been assigned one chooses a point in F_{U_i} instead). This is what we will refer to as *standard k-flats (resp., orthants)*, where k is the number of lines/rays. The *support* of a standard k-flat (resp., orthant) is the set of U_i for which a line/ray has been assigned. Related to this, a *complete support set* is a subset $\{U_i\}_{i=1}^{\nu} \subseteq \mathfrak{S}$ of pairwise orthogonal indices with all $\mathcal{C}(U_i)$ unbounded, and with maximal possible cardinality ν among sets with these properties. This definition is relevant here as ν is then the maximal dimension of standard orthants (if the $\mathcal{C}(U_i)$ are not only unbounded, but also have non-empty Gromov boundary, as will be the case for us). Moreover, in this case points in the Gromov boundary $\partial C(U_i)$ each determine a hierarchy ray, and similarly pairs of points determine hierarchy lines. In particular, for every complete support set $\{U_i\}_i = 1^{\nu}$ and for every choice of distinct points $p_i \in \partial \mathcal{C}U_i$ there exists an associated standard flat $\mathfrak{F}_{\{(U_i, p_i^{\pm})\}}$, which is the product of the hierarchy lines in F_{U_i} whose projections to $\mathcal{C}U_i$ have endpoints $p_i^{\pm}$.

8.2. The hinge graph

In this section, we discuss the graph that quasi-isometries will induce automorphisms of. We will call it the hinge graph, and it is best thought of, for the purposes of this section, as encoding standard orthants, as well as their coarse intersections. Here, recall that given subsets A, B of a metric space X, the coarse intersection $A\tilde{\cap}B$, if well-defined, is a subspace of X within bounded Hausdorff distance of all $N_R(A) \cap N_R(B)$ for R sufficiently large.

REMARK 8.3. The coarse intersection of any two standard flats $\mathfrak{F}_1, \mathfrak{F}_2$ is well-defined by [7, Lemma 4.12]. In fact, the lemma gives a description of the coarse

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

intersection as a subspace F of the HHS whose projection to any hyperbolic space U is the coarse intersection of $\pi_U(\mathfrak{F}_1)$ and $\pi_U(\mathfrak{F}_2)$. For example, the coarse intersection is a standard 1-flat provided that all coarse intersections in the various hyperbolic spaces are bounded, except for one hyperbolic space where $\pi_U(\mathfrak{F}_1)$ and $\pi_U(\mathfrak{F}_2)$ are both quasilines with the same endpoints in the Gromov boundary.

DEFINITION 8.4. (Hinge graph) *As in [7, Definition 5.2], let* **Hinge**$(\mathfrak{S})$ *be the set of* hinges*, that is, pairs* (U, p) *where* U *is contained in a complete support set and* $p \in \partial CU$*. We say that two hinges* (U, p) *and* (V, q) *are compatible if they are orthogonal and there exists some complete support set that contains both. Then we give* **Hinge**$(\mathfrak{S})$ *a metric graph structure by declaring two hinges to be adjacent if and only if they are compatible.*

As in [7, Definition 5.3], one can associate to a hinge $\sigma = (U, p)$ a standard 1-orthant, denoted h_σ. This is a hierarchy ray whose projection to $\mathcal{C}(U)$ is a quasi-geodesic ray asymptotic to p, and whose projections to all other coordinate spaces are uniformly bounded. Another property of h_σ, stated in [7, Remark 5.4], is that if $\sigma \neq \sigma'$ are two different hinges then $d_{Haus}(h_\sigma, h_{\sigma'}) = \infty$.

LEMMA 8.5. *Let* $(X, \mathfrak{S})$ *be a HHS. If for every* $U \in \mathfrak{S}$ *either the space* $\mathcal{C}U$ *is bounded or* $|\partial \mathcal{C}U| \geqslant 4$ *then:*

(1) *For every hinge* $\sigma = (U, p)$ *there exist two standard flats* $\mathfrak{F}_1$, $\mathfrak{F}_2$ *whose coarse intersection is supported in* U*, is coarsely a standard 1-flat, and contains* h_σ*.*
(2) *For any two compatible hinges* $\sigma = (U, p)$ *and* $\sigma' = (V, q)$ *there exist two standard flats* $\mathfrak{F}_1$, $\mathfrak{F}_2$ *whose coarse intersection is a standard 2-flat supported in* $\{U, V\}$ *and containing* h_σ *and* $h_{\sigma'}$*.*

This lemma will replace the three assumptions from [7, Section 5].

Proof.

(1) Choose another ideal point $q \in \partial \mathcal{C}U \setminus \{p\}$, and let γ be a hierarchy path in F_U whose projection to $\mathcal{C}U$ has limit points p and q. Then pick a complete support set $\{U_i\}$ completing $U = U_1$, and for each $i \geqslant 2$ choose four distinct ideal points $p_i^\pm, q_i^\pm \in \mathcal{C}U_i$. Let γ_i, δ_i be two bi-infinite hierarchy paths in each F_{U_i} whose projection to $\mathcal{C}U_i$ have limit points, respectively, $p_i^\pm$ and $q_i^\pm$. Then $\mathfrak{F}_1 = \gamma \times \prod_i \gamma_i$ and $\mathfrak{F}_1 = \gamma \times \prod_i \delta_i$ have the required coarse intersection in view of the discussion in Remark 8.3.
(2) The proof is very similar, in this case fixing two hierarchy lines rather than one.

□

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

REMARK 8.6. (Hidden hypothesis: asymphoricity) In the following theorem the HHSs will be required to be *asymphoric*, as in [7, Definition 1.14]. For our purposes, it suffices to know that asymphoric HHS include all hierarchically hyperbolic groups, such as our quotients of interest MCG/DT_K (see Remark 9.2).

THEOREM 8.7. *Let $(X, \mathfrak{S})$, $(Y, \mathfrak{T})$ be asymphoric HHSs such that for every $U \in \mathfrak{S}$ either the space $\mathcal{C}U$ is bounded or $|\partial \mathcal{C}U| \geqslant 4$, and similarly for every $V \in \mathfrak{T}$. Then every quasi-isometry $f : X \to Y$ induces an isomorphism $\mathrm{f}_{\mathrm{hin}}$ between the hinge graphs, such that for all $\sigma \in \mathbf{Hinge}(\mathfrak{S})$ we have $d_{Haus}(h_{f_{hin}(\sigma)}, f(h_\sigma)) < \infty$.*

Proof. This is a very similar statement to [7, Theorem 5.7], which applies to HHSs satisfying three additional assumptions (not satisfied in our case). An inspection of the proof shows that each use of the three assumptions can be replaced by our Lemma 8.5. □

The following is the analogue of [7, Lemma 5.9], which is a corollary of [7, Theorem 5.7] and whose proof does not use the additional assumptions on the HHSs further. Given a quasi-isometry between HHSs, the lemma gives a condition for the image of a standard flat to lie within bounded Hausdorff distance of a standard flat in terms of the induced map f_{hin} on hinge graphs. Crucially for applications, the bound is only in terms of the HHSs and the quasi-isometry constants.

LEMMA 8.8. (Flats go to flats) *Let $(X, \mathfrak{S})$, $(Y, \mathfrak{T})$ be asymphoric HHSs, such that for every $U \in \mathfrak{S}$ either the space $\mathcal{C}U$ is bounded or $|\partial \mathcal{C}U| \geqslant 4$, and similarly for every $V \in \mathfrak{T}$. Let $f : X \to Y$ be a quasi-isometry. There exists a constant C_0, depending only on the quasi-isometry constants of f, with the following property. Let $\{U_i\}_{i=1}^{\nu} \subseteq \mathfrak{S}$ be a complete support set, and let $p_i^{\pm}$ be distinct points in $\partial \mathcal{C}U_i$. Suppose that there exist a complete support set $\{V_i\}_{i=1}^{\nu} \subseteq \mathfrak{T}$ and distinct points $q_i^{\pm} \in \partial \mathcal{C}V_i$ such that for every $i = 1, \dots, \nu$ we have $f_{hin}(U_i, p_i^{\pm}) = (V_i, q_i^{\pm})$. Then, $d_{Haus}(f(\mathfrak{F}_{\{(U_i, p_i^{\pm})\}}), \mathfrak{F}_{\{(V_i, q_i^{\pm})\}}) \leqslant C_0$.*

9. Quasi-isometric rigidity

This section is devoted to the proof of Theorem 1. For the rest of the section, let $S = S_b$ be a sphere with $b \geqslant 7$. Let us start from the HHS structure of $MCG^{\pm}/DT_K$.

DEFINITION 9.1. *For Δ, Δ' simplices of some graph G, we write $\Delta \sim \Delta'$ to mean $Lk(\Delta) = Lk(\Delta')$. We denote by $[\Delta]$ the $\sim$-equivalence class of Δ, and by $\mathfrak{S}$ the set of equivalence classes of non-maximal simplices. Finally, we define the saturation of Δ as the set of vertices $v \in G$ for which there exists a simplex Δ' such that $v \in \Delta'$ and $\Delta' \sim \Delta$, i.e.*

$$Sat(\Delta) = \left(\bigcup_{\Delta' \in [\Delta]} \Delta' \right)^{(0)}.$$

REMARK 9.2. By [8, Theorem 7.1], MCG/DT_K is a hierarchically hyperbolic *group*, that is, it has a "nice" proper and cocompact action on some hierarchically hyperbolic space from which it inherits a HHS structure (see [6, Definition

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

1.21]). Moreover, combining [8, Theorem 7.1 and Proposition 8.13], we see that the HHS structure has the following properties:

(i) The index set is the set $\mathfrak{S}$ of $\sim$-equivalence classes of simplices inside $\mathcal{C}/DT_K$;
(ii) There is a bijection $j : \mathfrak{S} \to \mathfrak{S}^{\geqslant 1}/DT_K$, where $\mathfrak{S}^{\geqslant 1}$ is the set of essential, non-annular, possibly disconnected subsurfaces $U \subseteq S$.
(iii) Two elements $[\overline{\Delta}], [\overline{\Sigma}] \in \mathfrak{S}$ are orthogonal (resp. nested) if $j\left([\overline{\Delta}]\right)$ and $j\left([\overline{\Delta}]\right)$ admit representatives, which we also call *lifts*, that are disjoint (resp. nested);
(iv) If $\overline{\Delta}$ is not a facet, then the coordinate space $\mathcal{C}(\overline{\Delta}) := \mathcal{C}([\overline{\Delta}])$ associated to $[\overline{\Delta}]$ is quasi-isometric to $\mathrm{Lk}(\overline{\Delta})$ (see [8, Claim 6.12]). Otherwise, the pointwise stabiliser of the saturation $P(\overline{\Delta}) := \mathrm{Pstab}(\mathrm{Sat}(\overline{\Delta}))$ is a hyperbolic group acting properly and cocompactly on $\mathcal{C}(\overline{\Delta})$.

We will use *convex-cocompact subgroups* to show that the $\mathcal{C}(\overline{\Delta})$ have at least 4 points at infinity (or they are bounded). This notion was introduced in [22], and it admits several characterisations. The properties that we will need are that there exist free convex-compact subgroups inside every finite index subgroup (see e.g. [22, Theorem 1.4]), and the following fact from [30]. If a subgroup $Q < MCG(S)$ is convex-cocompact then there exists a constant D such that, for every element $h \in Q$ and for every two curves $x, s \subset S$ we have $d_s(x, h(x)) \leqslant D$ whenever the quantity is defined.

In our context, convex-cocompact subgroups survive in deep enough quotients:

LEMMA 9.3. *Let* S *be a connected surface of finite type and whose complexity is at least* 2. *Given a convex-cocompact subgroup* $Q < MCG(S)$, *for all large multiples* K *the projection* $\pi|_Q$ *is injective and the orbit maps of* Q *to* $\mathcal{C}(S)/DT_K$ *are quasi-isometric embeddings.*

Proof. This is just [8, Theorem 7.1.iii]. □

Now, in order for our machinery from § 8 to work, we must ensure that we are in the assumptions of Theorem 8.7:

LEMMA 9.4. *For all large multiples* K *the following holds. For any* $[\overline{\Delta}] \in \mathfrak{S}$, *either the space* $\mathcal{C}(\overline{\Delta})$ *is bounded or it has at least four points at infinity.*

Proof. First, assume that $\overline{\Delta}$ has codimension 1. Then $\mathcal{C}(\overline{\Delta})$ is quasi-isometric to the hyperbolic group $P(\overline{\Delta})$, and we claim that there exists a copy of the free group on two generators F_2 inside $P(\overline{\Delta})$. If this is the case then $P(\overline{\Delta})$ is an infinite, non-virtually cyclic hyperbolic group, hence it has at least four boundary points.

Let $\Delta \subset \mathcal{C}$ be a lift of $\overline{\Delta}$. By [8, Proposition 8.13.vii] we have that $\pi(\mathrm{Sat}(\Delta)) = \mathrm{Sat}(\overline{\Delta})$, therefore $P(\overline{\Delta})$ contains the quotient projection of $P(\Delta) := \mathrm{Pstab}(\mathrm{Sat}(\Delta))$. Thus it is enough to show that there is a copy of F_2 inside $P(\Delta)$ whose projection to $MCG^{\pm}/DT_K$ is injective. Now, Δ is a facet in $\mathcal{C}$, thus it cuts out a four-holed sphere U. Notice that a curve $\gamma \in \mathcal{C}$ is in the link of Δ if and only if γ lies in U and is not one of its boundary curves. Hence, a simplex Δ' which has the same link of

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

Δ must be a pants decomposition for $S \setminus U$, including its boundary curves. This in turn means that, if we see U just as a four-punctured sphere S_4 (that is, we forget about the difference between punctures and boundary curves), then $P(\Delta)$ contains the pure mapping class group $PMCG(S_4)$.

Inside $PMCG(S_4)$, which is finite-index inside $MCG(S_4)$, we may find a convex-cocompact copy of F_2, call it H. Suppose by contradiction that $\pi|_H$ is not injective, which means that there is some $h \in H \cap DT_K \setminus \{1\}$. Now h can not fix every curve, since the only elements of $MCG(S_4)$ with this property are the hyperelliptic involutions, which permute the punctures. Thus let $x \subseteq U$ be a curve such that $h(x) \neq x$ and let (s, γ_s) be as in Proposition 2.3. Notice that $h(x)$ still lies on U, therefore both x and $h(x)$ complete Δ to maximal simplices.

Now, suppose by contradiction that $d_{\mathcal{C}}(x, s) > 1$, and therefore $d_s(x, h(x)) > \Theta$. If we argue as in Lemma 2.8 we get that γ_s must fix Δ pointwise. Hence s is disjoint from all curves in Δ, but $s \notin \Delta$ since otherwise γ_s, which is a power of T_s, would fix x. Therefore, s lies in U, and by convex-cocompactness of H there exists a constant D (which we can choose independently of U) such that $d_s(x, h(x)) \leqslant D$. This is a contradiction if $\Theta > D$, which happens for all large multiples K.

Then we must have that $d_{\mathcal{C}}(x, s) \leqslant 1$, i.e. x must be fixed by γ_s, and we can apply γ_s to the whole data and proceed by induction on the complexity of h. In the end we must have that $h = \prod_{i=1}^{r} \gamma_{s_i}$ and every γ_{s_i} fixes x. But then $h(x) = x$, which contradicts our hypothesis. Thus we proved the Lemma for $\mathcal{C}(\overline{\Delta})$ whenever $\overline{\Delta}$ has codimension 1.

Now suppose that $\overline{\Delta}$ is not a facet. Then $\mathcal{C}(\overline{\Delta})$ is quasi-isometric to $\mathrm{Lk}(\overline{\Delta})$, which is the projection of $\mathrm{Lk}(\Delta)$ for some lift Δ. Notice that $\mathrm{Lk}(\Delta) = \mathcal{C}(U)$ is the curve graph of the subsurface U cut out by Δ. If U has at least two connected components which are not pairs of pants, then $\mathcal{C}(U)$ is bounded, and so is its projection. Otherwise U is some connected subsurface of complexity at least 2, and $\mathcal{C}(\overline{\Delta})$ is the image of the curve graph $\mathcal{C}(U)$ under the quotient map, which is isomorphic to $\mathcal{C}(U)/DT_K(U)$ by Corollary 3.2.

Now our goal has become to show that, for every surface U of complexity at least 2, the quotient of the curve graph $\mathcal{C}(U)/DT_K(U)$ has at least four boundary points. But now we can just apply Lemma 9.3: if we choose a convex-cocompact copy of F_2 inside $MCG(U)$ then its projection is still a copy of F_2 whose orbit maps to $\mathcal{C}(U)/DT_K(U)$ are quasi-isometric embeddings, and we are done. □

Thus we get:

Corollary 9.5. of Theorem 8.7 *Every quasi-isometry* $f : MCG/DT_K \to MCG/DT_K$ *induces an automorphism* $\mathrm{f}_{\mathrm{hin}}$ *of the hinge graph, such that for all* $\sigma \in \mathbf{Hinge}(\mathfrak{S})$ *we have* $d_{Haus}(h_{f_{hin}(\sigma)}, f(h_\sigma)) < \infty$.

9.1. Complete support sets for the quotient

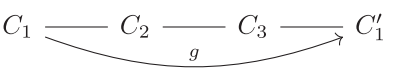
Moving forward, our next goal is to understand the structure of complete support sets. For the following theorems we think of $\mathfrak{S}$ as the set of classes of subsurfaces (thus omitting the bijection j whenever possible).

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

LEMMA 9.6. *The following holds for all large multiples* K. *Let* $\{\overline{U}_i\}_{i=1}^r$ *be a collection of pairwise orthogonal indices. Then there exist pairwise disjoint representatives* $\{U_i\}_{i=1}^r$.

Notice that this lemma is not at all obvious: we just know that every two indices in $\{\overline{U}_i\}_{i=1}^r$ have disjoint representatives, but this does not mean a priori that this conditions can all be satisfied simultaneously. Notice moreover that the $\overline{U}_i$s in the statement can be any equivalence classes of subsurfaces, not necessarily elements of a complete support set.

Proof of Lemma 9.6. We proceed by induction on r, the base case $r \leqslant 2$ being true by the description of the HHS structure. Then let $r \geqslant 3$ and choose three indices $\overline{U}_1, \overline{U}_2, \overline{U}_3$. Take lifts U_1, U_2, U_3, U_1' such that $U_1 \perp U_2$, $U_2 \perp U_3$ and $U_3 \perp U_1'$. Let $g \in DT_K$ be an element mapping U_1 to U_1'. For each of these surfaces we choose a family C_i of filling curves, in such a way that $C_1' = g(C_1)$. Since U_1 and U_2 are disjoint, we have a join $C_1 \star C_2$, and similarly for the other disjoint pairs. Then morally we have a "path" C_1, C_2, C_3, C_1', and we want to show that we can glue C_1 to C_1', as if we were looking for a "closed lift" of this path. The situation in the curve graph is as follows:



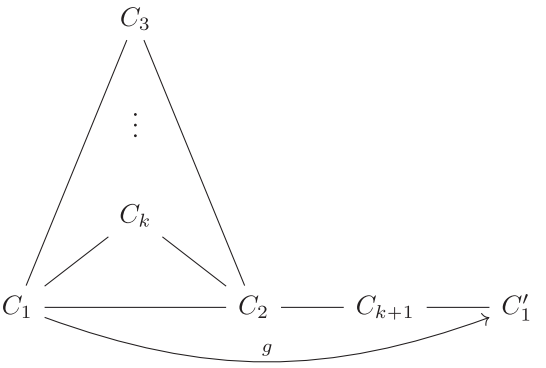
We proceed by induction on the complexity of g. If g is the identity we are done; otherwise fix a curve $x \in C_1$, let $x' = g(x)$, and let (s, γ_s) be as in Proposition 2.3, applied to x and g. If $d_{\mathcal{C}}(x, s) \leqslant 1$ we can apply γ_s to everything and proceed by induction. Otherwise, one between C_2 and C_3 must be fixed by γ_s pointwise, since if not we can find a path from x to x' that does not intersect the star of s, thus violating the Bounded geodesic image Theorem 2.4 (here we use the fact that consecutive C_is form a join). Either way we can apply γ_s to part of our chain and proceed by induction.

At the end of this process we get that U_1, U_2, U_3 are pairwise disjoint. Now suppose that, for $3 \leqslant k < r$ we can find representatives $U_1, \dots, U_k$ such that U_1, U_2, U_i are pairwise disjoint for all $3 \leqslant i \leqslant k$, and we want to show that the same holds for $k+1$. As before, let U_{k+1} be a representative for $\overline{U}_{k+1}$ which is disjoint from U_2 and let U_1' be a representative for $\overline{U}_1$ disjoint from U_{k+1}. Let $C_1, \dots, C_{k+1}, C_1'$ be sets of filling curves for the corresponding subsurfaces. In the curve graph there is a "triangle" of the form C_1, C_2, C_i for every $i = 3, \dots, k$, as shown in this schematic picture:

The same argument as before shows that we can glue C_1 to C_1', this time without breaking the "triangles": in every inductive step, γ_s fixes either C_1 (and we apply γ_s to the whole data) or one between C_2 and C_{k+1} (and we can apply γ_s beyond these curves, without moving the triangles).

If we do this for $k = 3, \dots, r-1$ we can find representatives $\{U_i\}_{i=1}^r$ such that $U_1 \perp U_2$ and both U_1 and U_2 are disjoint from every other U_i. But now we can consider U_1 and U_2 as a single (possibly disconnected) subsurface and conclude by induction on r. □

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press



In the proof of Lemma 9.6 we actually showed that every "simplex of indices" admits a lift. Now we want to prove the uniqueness of these lifts, up to elements of DT_K:

LEMMA 9.7. *For all large multiples* K *the following holds. Let* $\{\overline{U_i}\}_{i=1}^r$ *be a collection of pairwise orthogonal indices. Any two collections of pairwise orthogonal representatives* $\{U_i\}_{i=1}^r$, $\{U_i'\}_{i=1}^r$ *are obtained one from the other via some element* $g \in DT_K$.

Proof. We proceed by induction on r, the base case $r=1$ being clear. If the conclusion holds for $r-1$ then, up to some element $g \in DT_K$, we can assume that $U_i = U_i'$ for $i = 1, \dots, r-1$. Now let $h \in DT_K$ be an element that maps U_r to U_r', and let $C_1, \dots, C_r, C_r'$ be sets of filling curves such that $h(C_r) = C_r'$. If h is not the identity, fix a point $x \in C_r$ and let (s, γ_s) be as in Proposition 2.3. If $d_{\mathcal{C}}(x,s) \leqslant 1$ we can apply γ_s to everything and proceed by induction on the complexity of h; otherwise γ_s must fix $C_1, \dots, C_{r-1}$ pointwise, so we can apply γ_s to the "simplex" $\{C_1, \dots, C_{r-1}, C_r'\}$ and proceed by induction. □

The main reason we established Lemmas 9.6 and 9.7 is that every complete support set for MCG/DT_K lifts to a collection of pairwise orthogonal S_4 and S_5 of the maximal cardinality, which we call a *complete support set* for the surface S; moreover, such a collection is unique up to the action of DT_K.

Recall that a surface is *odd* if its complexity is odd, otherwise it is *even* (in our case, S_b is odd iff b is even). Bowditch [13, Section 6] pointed out that if a subsurface U belongs to a complete support set for S, then U must have one of the following shapes, which ensure that every complete support set containing U can cut out at most one pair of pants.

- If S is odd, then U is an S_4, and each connected component of the complement $S - U$ is odd and meets U in exactly one curve.
- If S is even then U is either an S_4 with all but one of the complementary components odd, or an S_5 with all complementary components odd.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

REMARK 9.8. (Cutting out a pair of pants) Notice that the only case where U does not admit a complete support set whose union is the whole S is when S is even and the even connected component of $S - U$ is a pair of pants. Indeed:

- If S is odd, every connected component of $S - U$ is an odd sphere, and can be cut into four-holed spheres. The same holds if S is even and U is an S_5.
- If S is even and U is an S_4, we can fill the odd complementary components with four-holed spheres as above. Then, if the even component is not a pair of pants then it must be a sphere with at least five holes, and one can fill it with an S_5 and some S_4.

9.2. Unambiguous domains

Let $\sigma = (\overline{U}, p)$ be a hinge for MCG/DT_K, and let $\mathrm{Compl}(\sigma)$ be the set of all tuples of pairwise orthogonal hinges that σ completes to a complete support set. In other words, $\mathrm{Compl}(\sigma)$ is the set of all facets in $\mathbf{Hinge}(\mathfrak{S})$ that σ completes to a maximal simplex.

DEFINITION 9.9. *Two hinges σ, σ' are said to be* equally completable*, or to have the same completions, if $Compl(\sigma) = Compl(\sigma')$.*

This definition clearly induces an equivalence relation which is preserved by any automorphism of the hinge graph. One would like to think that if two hinges $(\overline{U}, p)$ and $(\overline{U}', p')$ are equally completable then $\overline{U} = \overline{U}'$. This is unfortunately not the case. Indeed, if S is even and $\overline{U}$ is the class of a subsurface U which is an S_4 whose even complementary component is a pair of pants P, then $\overline{U}$ has the same completions as the class of the S_5 given by the union of U and P. See Figure 28 to understand the situation. Luckily, we will see that this is the only problem that could arise.

DEFINITION 9.10. *A hinge σ is* unambiguous *if any other equally completable hinge σ' has the same support. A support $\overline{U}$ is unambiguous is every hinge supported on* U *is unambiguous.*

We need another definition that characterises (some) unambiguous domains, and can be recognised from the hinge graph.

DEFINITION 9.11. *A hinge σ is* minimal *if $Compl(\sigma)$ is maximal by inclusion, among completions.*

LEMMA 9.12. *For all large multiples* K*, a hinge $(\overline{U}, p)$ for MCG/DT_K is minimal if and only if it is unambiguous and its support is the class of a four-holed sphere.*

Proof. Let U be a lift of $\overline{U}$. We examine all possible shapes of U.

- Suppose U is an S_5, and choose some subsurface $V \sqsubseteq U$ such that V is an S_4 and $S - V$ does not contain a pair of pants. Such a V exists, as one can choose a pair of pants Q whose boundary intersects the boundary of $S - U$, and set $V = U - Q$. Every completion for $\overline{U}$ admits representatives

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

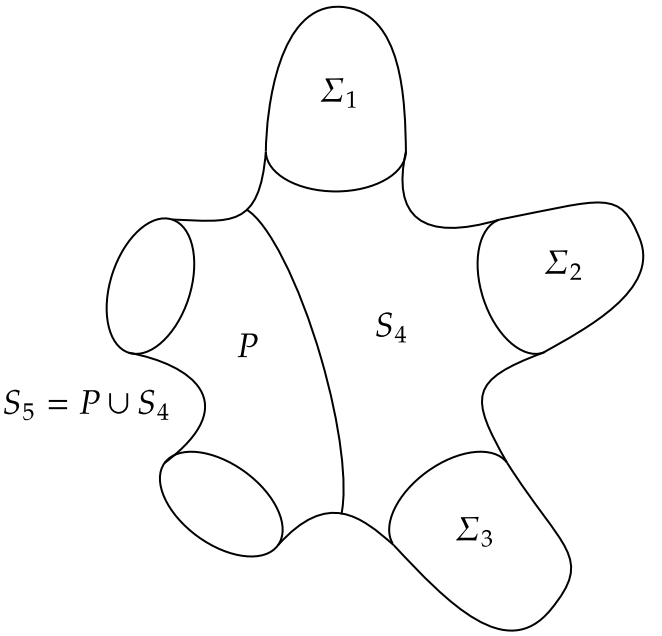


Figure 28. The classes of the S_4 and the S_5 in the Figure cannot be distinguished just by their completions, all of which lift inside the union of the odd complementary regions Σ_1, Σ_2, and Σ_3.

that are disjoint from U by Lemma 9.6, and therefore also from V. Thus every completion for $\overline{U}$ is also a completion for $\overline{V}$. Moreover we claim that there exists a completion $\{\overline{V}_i\}$ for $\overline{V}$ but not for $\overline{U}$. To this purpose choose a completion $\{V_i\}$ for V, such that some boundary curve δ crosses one of the boundary components η of U, and let $\{\overline{V}_i\}$ be their classes. Now recall that, as a consequence of Lemma 3.3, if we fix some finite subset $F \subset \mathcal{C}$ then for all large multiples K the projection map is an isometry on $g(F)$, for every mapping class $g \in MCG$. In our case, there is a finite number of S_5 inside S, up to the action of the mapping class group, and for each of these we can find δ and η as above. Then we set F as the union of these curves. Now, any lift $\{V_i\}$ of $\{\overline{V}_i\}$ must intersect U. More precisely, η must cross the boundary curve δ' corresponding to δ, since

$$d_{\mathcal{C}}(\delta', \eta) \geqslant d_{\mathcal{C}/DT_K}(\overline{\delta'}, \overline{\eta}) = d_{\mathcal{C}/DT_K}(\overline{\delta}, \overline{\eta}) = d_{\mathcal{C}}(\delta, \eta) \geqslant 2,$$

where we used that the projection is 1-Lipschitz and that its restriction to F is isometric. This shows that $\overline{U}$ is not minimal.

- Suppose U is an S_4 which cuts out a pair of pants, as in Figure 28. Let V be the S_5 given by the union of U and the pair of pants cut out by U. Then, slightly abusing notation, we have that $\mathrm{Compl}(\overline{U}) = \mathrm{Compl}(\overline{V})$. Indeed, every support $\overline{W}$ which completes $\overline{V}$ also completes $\overline{W}$; conversely, if $\overline{W}$ completes $\overline{U}$ then it admits a lift W inside one of the complementary

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

components of U, which is therefore disjoint from V. Thus $\overline{U}$ is ambiguous. Moreover, by the above argument $\overline{V}$ is not minimal as it is an S_5, and therefore neither is $\overline{U}$ which has the same completions as $\overline{V}$.

- The only case left is when U is an S_4 which does not cut out a pair of pants. Choose a completion $\{U_i\}$ such that $U = S \setminus \bigcup U_i$, as in Remark 9.8, and let $\{\overline{U}_i\}$ be its projection. If $\overline{V}$ completes $\{\overline{U}_i\}$ then we may lift it to some $V \sqsubseteq U$ (here we used that, up to elements of DT_K, the lift of $\{\overline{U}_i\}$ is unique). But then $V = U$ since U has minimal complexity. This proves that $\overline{U}$ is minimal and unambiguous.

□

9.3. Minimal product regions are preserved

REMARK 9.13. Any automorphism of the hinge graph must map minimal hinges with the same completions to minimal hinges with the same completions, therefore it acts on the set of minimal supports. Thus, if $f\colon MCG/DT_K \to MCG/DT_K$ is a quasi-isometry, and $\overline{U}$ is a minimal support, we can define $f_{supp}(\overline{U})$ as the support of $f_{hin}(\overline{U}, p)$ for any hinge $(\overline{U}, p)$ supported on $\overline{U}$, where f_{hin} is the map from Corollary 9.5.

COROLLARY 9.14. (Products go to products) *Let $f\colon MCG/DT_K \to MCG/DT_K$ be a quasi-isometry. There exists* C*, depending only on the quasi-isometry constants of* f*, with the following property. Let $\{\overline{U}_i\} \subseteq \mathfrak{S}$ be a complete support set made of minimal supports, and let $f_{supp}(\overline{U}_i) = \overline{V}_i$ for all* i*. Let $P_{\{\overline{U}_i\}}$ and $P_{\{\overline{V}_i\}}$ be the standard product regions defined by $\{\overline{U}_i\}$ and $\{\overline{V}_i\}$, respectively. Then $d_{Haus}\left(f\left(P_{\{\overline{U}_i\}}\right), P_{\{\overline{V}_i\}}\right) \leqslant C$.*

Proof. Another way of stating Remark 9.13 is that, if $(\overline{U}, p^{\pm})$ are two hinges with the same minimal support and $f_{hin}(\overline{U}, p^+) = (\overline{V}, q^+)$, then there exists $q^- \in \partial\mathcal{C}V$ such that $f_{hin}(\overline{U}, p^-) = (\overline{V}, q^-)$. Thus the Flats to Flats Lemma 8.8 says that, if $\{(\overline{U}_i, p_i^{\pm})\}$ is a complete support set made of minimal supports, with a choice of two points in every $\partial\mathcal{C}\overline{U}_i$, and if we set $(\overline{V}_i, q_i^{\pm}) := f_{hin}(\overline{U}_i, p_i^{\pm})$, then $d_{Haus}(f(\mathfrak{F}_{\{(\overline{U}_i, p_i^{\pm})\}}), \mathfrak{F}_{\{(\overline{V}_i, q_i^{\pm})\}}) \leqslant C_0$ for some constant C_0 depending only on the quasi-isometry constants of f. Hence

$$d_{Haus}\left(f\left(\bigcup \mathfrak{F}_{\{(\overline{U}_i, *)\}}\right), \bigcup \mathfrak{F}_{\{(f_{supp}(\overline{U}_i), *)\}}\right) \leqslant C_0$$

where $\bigcup \mathfrak{F}_{\{(\overline{U}_i, *)\}}$ is the union of all standard flats supported in $\{\overline{U}_i\}$.

Now, the conclusion follows if we prove the existence of some constant C_1 such that, for every complete support set $\{\overline{U}_i\}$ made of minimal supports,

$$d_{Haus}\left(\bigcup \mathfrak{F}_{\{(\overline{U}_i, *)\}}, P_{\{\overline{U}_i\}}\right) \leqslant C_1$$

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

In fact, if this is the case then we also get that,

$$d_{Haus}\left(f\left(\bigcup \mathfrak{F}_{\{(\overline{U}_i,*)\}}\right), f\left(P_{\{\overline{U}_i\}}\right)\right) \leqslant C_2$$

where C_2 is some constant depending only on C_1 and the quasi-isometry constants of f. In turn, since $P_{\{\overline{U}_i\}} = \prod F_{\overline{U}_i}$, to prove the existence of such C_1 it suffices to select a close enough hierarchy line in each coordinate, as we shall do in Lemma 9.15. □

LEMMA 9.15. *There exists a constant* C_3 *such that whenever* $\overline{U}$ *is a minimal domain and* $x \in F_{\overline{U}}$*, there exists a hierarchy line* $\gamma \subset F_{\overline{U}}$ *such that* $d_{F_{\overline{U}}}(x, \gamma) \leqslant C_3$*.*

Proof. Since $P_{\{\overline{U}_i\}} = \prod F_{\overline{U}_i}$ it suffices to work in each coordinate, that is, we will show that there exists a constant C_3 such that whenever $\overline{U} \in \mathfrak{S}$ is a minimal domain and $x \in F_{\overline{U}}$, there exists a hierarchy line $\gamma \subset F_{\overline{U}}$ such that $d_{F_{\overline{U}}}(x, \gamma) \leqslant C_3$. Notice that, by the distance formula and the fact that $\overline{U}$ is a $\sqsubseteq$-minimal, we have that $F_{\overline{U}}$ is uniformly quasi-isometric to $\mathcal{C}\overline{U}$, and hierarchy lines simply correspond to bi-infinite quasigeodesics (with certain constants). Hence we just need to show that every point $x \in \mathcal{C}\overline{U}$ is within uniformly bounded distance from a bi-infinite quasigeodesic. If $\overline{U} = j([\overline{\Delta}])$, as in Remark 9.2.(iv), then $\mathcal{C}\overline{U}$ is uniformly quasi-isometric to the hyperbolic group $P(\overline{\Delta})$, which has at least 2 points at infinity and therefore admits bi-infinite quasigeodesics. Since a group acts transitively on itself the conclusion follows. □

9.4. Terminal supports and 1-separating curves

We want to translate the automorphism f_{supp} into an automorphism of the graph $\mathcal{C}^1/DT_K$ from § 6. In other words, we need a way to identify the minimal domains which are cut out by 1-separating curves. To this purpose, we introduce the following definition.

DEFINITION 9.16. *A support* $\overline{U}$ *is* terminal *if it is the class of a subsurface* U *of complexity* 1 *which is cut out by a single curve. Notice that the boundary curve of such a* U *is* 1*-separating.*

LEMMA 9.17. *For all large multiples* K*, a hinge* $\sigma = (\overline{U}, p)$ *has terminal support if and only if it is minimal and there exists a hinge* $(\overline{V}, q)$*, compatible with* σ*, such that any complete support set containing* $(\overline{V}, q)$ *must contain some* σ' *supported in* $\overline{U}$*.*

Proof. Suppose $\overline{U}$ is terminal, and choose a lift U of $\overline{U}$. Let V be a support such that U is one of the connected components of $S - V$; such a V exists as $S - V$ has complexity at least 2, since we assumed that $S = S_b$ with $b \geqslant 7$. Then every complete support set containing $\overline{V}$ lifts to a complete support set containing V, and therefore also U. Moreover, by the discussion in the proof of Lemma 9.12, a terminal support is also minimal.

Conversely, suppose that $\overline{U}$ is not terminal. If $\overline{U}$ is not minimal we have nothing to prove. Otherwise $\overline{U}$ is the class of a non-terminal S_4, and we want to show

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

that every $\overline{V} \perp \overline{U}$ which is compatible with $\overline{U}$ admits a completion that does not contain $\overline{U}$. First notice that, up to the action of the mapping class group, there is a finite number of possible pairs $U \perp V$ of orthogonal subsurfaces. For each of these possibilities, let $\{V_i\}_{i=2}^{\nu}$ a completion for $V = V_1$, that we can choose in such a way that some boundary curve δ of U crosses some boundary curve η of some V_i, say, V_2. This is always possible because U, which is a non-terminal S_4, cannot coincide with the connected component Σ of $S \setminus V$ it belongs to, and therefore there must be a relative boundary curve of U inside Σ that we can choose as δ. Let F be the finite union of all δs and ηs that arise from these possibilities, and for all large multiples K the projection is an isometry on F. Thus $\{\overline{V}_i\}$ is a completion for $\overline{V}$ that cannot contain $\overline{U}$, since any lift U' of $\overline{U}$ contains a boundary curve δ' which must cross η (we can argue precisely as in Lemma 9.12). Hence $\overline{U} \pitchfork \overline{V}_2$ since $U' \pitchfork V_2$ for every lift U'. This implies that $\overline{U}$ cannot be an element of $\{\overline{V}_i\}$, and we are done. □

COROLLARY 9.18. *For all large multiples* K*, any self-quasi-isometry* f *of* MCG/DT_K *induces an automorphism* ϕ *of* $\mathcal{C}^1/DT_K$.

Proof. For every terminal support $\overline{U}$, let $\partial\overline{U}$ be the element $\overline{\gamma} \in \mathcal{C}^1/DT_K$ such that there exist lifts U and $\gamma = \partial U$. This is a one-to-one correspondence between terminal supports and vertices of $\mathcal{C}^1/DT_K$ (here we are using that every 1-separating curve cuts out exactly one S_4, as $S = S_b$ with $b \geqslant 7$). Now, since being terminal has a completely combinatorial characterisation, in view of Lemma 9.17, f_{supp} must map terminal supports to terminal supports, and we can set $\phi(\partial\overline{U}) = \partial f_{supp}(\overline{U})$. Moreover, f_{supp} preserves compatibility, which, for terminal subsurface, is equivalent to the fact that their boundaries admit disjoint representatives. Thus ϕ is an automorphism of $\mathcal{C}^1/DT_K$. □

By Theorem 6.4, ϕ comes from some $\overline{g} \in MCG^{\pm}/DT_K$. Finally, we need to show that, if f_{supp} and $\overline{g}$ agree on terminal subsurfaces then they agree on every minimal surface. More precisely we claim the following:

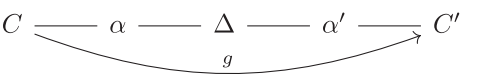
LEMMA 9.19. *For all large multiples* K*, every minimal support* $\overline{U}$ *for* MCG/DT_K *is uniquely determined by the terminal supports it is compatible with.*

Proof. Let U be a lift of $\overline{U}$ and let $S \setminus U = \bigsqcup_{i=1}^{4} \Sigma_i$. Moreover, let $\overline{V}$ be another support such that every terminal support $\overline{T}$ compatible with $\overline{U}$ is also compatible with $\overline{V}$. We claim that, for $i = 1, \dots, 4$, there exists a representative V_i for $\overline{V}$ which is disjoint from Σ_i. If this is the case then there exists a representative V which is disjoint from all Σ_is (more precisely, we can use Lemma 9.6 to lift the "simplex" $\{\overline{V}\} \cup \{\overline{\Sigma}_i\}_{i=1}^4$, and Lemma 9.7 shows that we can choose $\{\Sigma_i\}_{i=1}^4$ as lifts of $\{\overline{\Sigma}_i\}_{i=1}^4$). Therefore, $V \sqsubseteq U$, and equality holds since U has already minimal complexity.

First notice that there is a finite number of possibilities for U, up to the action of the mapping class group. For each of these possibilities look at its complementary components. Whenever one of these, call it Σ, is not terminal we choose two 1-separating curves α, α' that fill Σ and a pants decomposition Δ for $S \setminus \Sigma$, including its boundary. Let F be the finite set of curves given by the union of all these α, α' and Δ. Notice that Corollary 3.9 tells us that for all large multiples K every lift of $\overline{\alpha}, \overline{\alpha}'$ inside $\mathrm{Lk}(\Delta)$ is still a pair of filling curves for Σ.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

Now we go back to our proof. If Σ is already a terminal support then we can find a representative V which is disjoint from Σ, and we are done. Otherwise let α, α', Δ be the image under some mapping class of the corresponding elements of F, which satisfy the property that every two lifts of $\overline{\alpha}, \overline{\alpha}'$ inside $\mathrm{Lk}(\Delta)$ fill Σ. Let W, W' be the terminal subsurfaces cut out by α and α', and let V, V' be some representatives of $\overline{V}$ such that $V \perp W$ and $V' \perp W'$. Let $g \in DT_K$ that maps V to V'. Let C be a collection of filling curves for V, and let $C' = g(V)$. Then in the curve graph the situation is as follows:



Now if g is not the identity pick some $x \in C$ and let (s, γ_s) be as usual. If $d_{\mathcal{C}}(x, s) \leqslant 1$ we can apply γ_s to everything and proceed by induction on the complexity of g. Otherwise we must be in one of these three cases:

- If $d_{\mathcal{C}}(\alpha, s) \leqslant 1$ we can apply γ_s to everything after α. Now $\gamma_s(\alpha)$ and $\gamma_s(\alpha')$ still fill $\gamma_s(\Sigma)$, and we can proceed.
- If $d_{\mathcal{C}}(\alpha', s) \leqslant 1$ we can apply γ_s just to C', without touching α and α'.
- If Δ is fixed pointwise by γ_s then we can apply γ_s to everything after Δ. Notice that $\gamma_s(\alpha')$ is still a lift of $\overline{\alpha}'$ in the link of Δ, hence it still fills Σ together with α.

At the end of the induction we have that $C = C'$ is disjoint from both α and α', which may differ from the original curves but remain a pair of filling curves for a representative of $\overline{\Sigma}$. This gives us the required representative for $\overline{V}$. □

9.5. Quasi-isometric rigidity

We are finally ready to prove quasi-isometric rigidity of $MCG^{\pm}/DT_K$, which is Theorem 1. We subdivide the proof in two steps.

THEOREM 9.20. *Let $S = S_{0,b}$ be a punctured sphere, with $b \geqslant 7$ punctures. For all large multiples* K *the following holds. For every* $\mathrm{T} > 0$ *there exists* $\mathrm{D} > 0$ *such that every* (T, T)*-self-quasi-isometry* f *of $MCG(S)/DT_K$ lies within distance* D *of the left multiplication by some element $\overline{g} \in MCG^{\pm}/DT_K$, which depends only on the restriction of* $\mathrm{f}_{\mathrm{supp}}$ *to terminal supports.*

Proof. Let f be a self-quasi-isometry of MCG/DT_K, and let f_{supp} be the induced map on minimal supports. With a slight abuse of notation, the restriction of f_{supp} to terminal supports is an automorphism of $\mathcal{C}^1/DT_K$ which comes from some element $\overline{g} \in MCG^{\pm}/DT_K$ by Theorem 7.1. Then f_{supp} and g agree on terminal supports, and therefore also on every minimal support by Lemma 9.19. Now, the Products to products Corollary 9.14 tells us that there exists some constant C, depending only on the quasi-isometry constants of f, such that, if $\{\overline{U}_i\}$ is a complete support set made of minimal supports, f maps the corresponding product region $P_{\{U_i\}}$ within Hausdorff distance at most C from $\overline{g}(P_{\{U_i\}})$.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

We are left to prove that every point $x \in MCG/DT_K$ is the (uniform) coarse intersection of two standard product regions $P\tilde{\cap}P'$, coming from minimal complete support sets. If this is the case then $f(x)$ will be the coarse intersection of $f(P)$ and $f(P')$, which lie at uniformly bounded distance from the coarse intersection of $\overline{g}(P)$ and $\overline{g}(P')$, which is coarsely $\overline{g}(x)$. Notice that it is enough to show that there is *some* point $x_0 \in MCG/DT_K$ which is the coarse intersection of two standard product regions with minimal supports, since MCG/DT_K acts transitively on itself and maps product regions to product regions. In turn, by Corollary 8.2 we are left to prove that there exist two complete support sets $\{\overline{U}_i\}$ and $\{\overline{V}_i\}$ with minimal, pairwise distinct supports (notice that we can apply Corollary 8.2 since minimal supports are S_4, and therefore they are also minimal with respect to nesting).

Firstly, fix a complete support set $\{\overline{U}_i\}$ as follows:

- If S is odd, choose any complete support set $\{\overline{U}_i\}$.
- If S is even, let $\tilde{U}$ be the class of an S_5 cut out by a single curve, and complete $\tilde{U}$ to a complete support set.

In both cases, by inspection of the proof of Lemma 9.12, one gets that $\{\overline{U}_i\}$ is made of minimal supports. Let $\{U_i\}$ be a lift of $\{\overline{U}_i\}$, and choose a pseudo-Anosov mapping class h such that every boundary curve of $\{U_i\}$ crosses every boundary curve of $\{V_i\} = h\{U_i\}$. Now for all large multiples K the projection is an isometry on the finite set F of boundary curves of $\{U_i\}$ and $\{V_i\}$. Therefore, we must have that $\overline{U}_i \neq \overline{V}_j$ for every choice of i and j, because any two lifts U_i' and V_j' must have crossing boundaries. This proves the theorem. □

COROLLARY 9.21. *For every* $\mathrm{T} > 0$ *there exists* D *such that, if a* (T, T)*-self-quasi-isometry* f *of* MCG/DT_K *lies within finite distance of the identity, then it lies within distance* D *of the identity.*

Proof. By the previous theorem we know that f lies within distance D of the left multiplication by some $\overline{g}$, which depends only on the induced map f_{supp} on terminal supports. In turn f_{supp} is induced by f_{hin}, thus if we show that this map is the identity then $\overline{g}$ can be chosen to be the identity, and the corollary follows. Now, Theorem 8.7 tells us that $d_{Haus}(h_{f_{hin}(\sigma)}, f(h_\sigma)) < \infty$. But then, since $d_{Haus}(f(h_\sigma), h_\sigma) < \infty$ we must also have that $d_{Haus}(h_{f_{hin}(\sigma)}, h_\sigma) < \infty$. This in turn means that $f_{hin}(\sigma) = \sigma$, since h_σ has the property that if $\sigma \neq \sigma'$ then $d_{Haus}(h_\sigma, h_{\sigma'}) = \infty$. □

Thus Theorem 1 is implied by the general statement below, which follows from standard arguments (see e.g. [31, Section 10.4]):

LEMMA 9.22. *Let* H *be a finitely generated group. Suppose that for every* $\mathrm{T} > 0$ *there exists* $\mathrm{D} > 0$ *such that:*

(1) *Every* (T, T)*-self-quasi-isometry of* H *lies within distance* D *of the left multiplication by some element of* H*;*
(2) *If a* (T, T)*-self-quasi-isometry of* H *lies within finite distance of the identity then it lies within distance* D *of the identity.*

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

Then H is a finite extension of the group $QI(H)$ of its self-quasi-isometries up to bounded distance. Moreover, if a finitely generated group G is quasi-isometric to H then G and H are weakly commensurable, meaning that there exist two finite normal subgroups $L \trianglelefteq H$ and $M \trianglelefteq G$ such that the quotients H/L and G/M have two finite index subgroups that are isomorphic.

Proof. Set $\mu : H \to \mathrm{QI}(H)$ by mapping h to the left multiplication by h. This map is surjective by Item (1); moreover it has finite kernel, since if $\mu(h)$ is within finite distance of the identity then it is within distance $D = D(1)$ of the identity, which means that $d_H(h, 1) = d_H(\mu(h)(1), 1) \leqslant D$ and we conclude since balls in Cayley graphs are finite.

Regarding the second statement let L be such that $H/L \cong \mathrm{QI}(H)$, and let $\phi : G \to H$ be a quasi-isometry with quasi-inverse ϕ^{-1}. We define a group homomorphism $\psi : G \to \mathrm{QI}(H)$ by setting

$$\psi(g)(h) = \phi(g\phi^{-1}(h))$$

Notice that $\psi(g)$ is a self-quasi-isometry of H, whose constants (T, T) depend only on the quasi-isometry constants of ϕ and ϕ^{-1}. Let $D = D(T)$ as in the hypothesis. The same argument as before shows that ψ has finite kernel: if $g \in \ker \psi$ then $\phi(g\phi^{-1}(1))$ is D-close to 1, hence $g\phi^{-1}(1)$ is D'-close to $\phi^{-1}(1)$ for some other constant $D'(\phi, \phi^{-1}, D)$. Thus the Lemma follows if we prove that ψ is coarsely surjective.

For every $g \in G$ let $\theta(g) \in H$ be the element whose multiplication is D-close to $\psi(g)$. Then by construction the following diagram commutes:

$$\begin{array}{ccc} & & H \\ & \overset{\theta}{\nearrow} & \downarrow \mu \\ G & \xrightarrow{\psi} & \mathrm{QI}(H) \end{array}$$

Since μ is a quotient map it is 1-Lipschitz and surjective. Hence it is enough to show that θ is coarsely surjective, which in turn will follow if we prove that θ lies within bounded distance from the quasi-isometry ϕ. In fact $\theta(g)$ is D-close to $\phi(g\phi^{-1}(1))$, which in turn is uniformly close to $\phi(g)$ since $d_G(g, g\phi^{-1}(1)) = \|\phi^{-1}(1)\|_G$ is constant. □

10. Algebraic rigidity

This section is devoted to the proof of Theorem 2, which is covered by Theorems 10.1 and 10.9 below.

THEOREM 10.1. *For every $b \geqslant 7$ and for all large multiples K the following holds. Let $\phi : H \to H'$ be an isomorphism between finite index subgroups of $MCG^{\pm}(S_b)/DT_K$. Then ϕ is the restriction of an inner automorphism. In particular $Out(MCG^{\pm}(S_b)/DT_K)$ is trivial.*

We recall some definitions from [2] that we will need to state the exact theorem.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

DEFINITION 10.2. *Two elements* h *and* g *of a group* G *are* commensurable*, and we write* $h \overset{G}{\approx} g$*, if there exist* $m, n \in \mathbb{Z} \setminus \{0\}$*,* $k \in G$ *such that* $kg^m k^{-1} = h^n$ *(that is, if they have non-trivial conjugate powers).*

DEFINITION 10.3. *If a group* G *acts by isometries on a hyperbolic space* $\mathcal{S}$*, an element* $g \in G$ *is* loxodromic *if for some* $x \in \mathcal{S}$ *the map* $\mathbb{Z} \to \mathcal{S}$*,* $n \mapsto g^n(x)$ *is a quasi-isometric embedding. In the same setting, an element* $g \in G$ *is* weakly properly discontinuous*, or* WPD*, if for every* $\varepsilon > 0$ *and any* $x \in \mathcal{S}$ *there exists* $N_0 = N_0(\varepsilon, x)$ *such that whenever* $N \geqslant N_0$ *we have*

$$\left|\left\{h \in G \middle| \max\left\{d_S\left(x, h(x)\right), d_S\left(g^N(x), hg^N(x)\right)\right\} \leqslant \varepsilon\right\}\right| < \infty$$

We denote by $\mathcal{L}_{WPD}$ *the set of loxodromic WPD elements.*

The following result is a special case of [2, Theorem 7.1]. Roughly speaking, the theorem says that if H is a subgroup of G and both act "interestingly enough" on some hyperbolic space, then any homomorphism $\phi : H \to G$ is either (the restriction of) an inner automorphism or it maps some loxodromic WPD to an element which is not commensurable to it.

THEOREM 10.4. *Let* G *be a group acting coboundedly and by isometries on a hyperbolic space* $\mathcal{S}$*, with loxodromic WPD elements. Let* $H \leqslant G$ *be a non-virtually-cyclic subgroup such that* $H \cap \mathcal{L}_{WPD} \neq \emptyset$*, and let* $E_G(H)$ *be the unique maximal finite subgroup of* G *normalised by* H*, whose existence is proven in [2, Lemma* 5.6*]. Let* $\phi : H \to G$ *be a homomorphism such that whenever* $h \in H \cap \mathcal{L}_{WPD}$ *then* $\phi(h) \overset{G}{\approx} h$*. If* $E_G(H) = \{1\}$ *then* ϕ *is the restriction of an inner automorphism.*

Proof. Outline of Theorem 10.1. We just need to verify that, for all $b \geqslant 7$ and for all large multiples K, the hypotheses of Theorem 10.4 are satisfied for $G = MCG^{\pm}(S_b)/DT_K$ and any isomorphism $\phi : H \to H'$ between subgroups of finite index. In [19, Theorems 2.1 and 5.2] it was proven that G is not virtually cyclic, it acts coboundedly and by isometries on the hyperbolic space $S = \mathcal{C}/DT_K$, and the action admits loxodromic WPD elements. In particular, every subgroup of finite index is not virtually cyclic and it contains some power of every loxodromic WPD element, which remains loxodromic WPD. In Lemma 10.5 we show that ϕ has the required commensurating property, that is, for every $h \in H \cap \mathcal{L}_{WPD}$ we have that $\phi(h) \overset{G}{\approx} h$. Moreover, in view of the general Lemma 10.6, in order to prove that $E_G(H) = \{1\}$ it will suffice to prove that $E_G(G) = \{1\}$, i.e. that $MCG^{\pm}/DT_K$ has no non-trivial finite normal subgroups. This is done in Lemma 10.8. □

LEMMA 10.5. *Let* $G = MCG^{\pm}(S_b)/DT_K$ *for* $b \geqslant 7$*. For every isomorphism* $\phi : H \to H'$ *between finite index subgroups and every element* $h \in H$ *of infinite order,* h *and* $\phi(h)$ *are commensurable.*

Proof. Fix a finite generating set for G, and let d_G be the corresponding word metric. Since H has finite index there exists a quasi-isometry $f : G \to H$, which we can choose to be the identity on H (for example, every $g \in G$ can be sent to one of the closest elements of H). Let $\Phi = \phi \circ f : G \to H'$, which coincides with

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

ϕ on H and is a self-quasi-isometry of G because ϕ is an isomorphism between finite index subgroups. Then by Theorem 9.20 there exist a constant D and an element $g \in G$ such that Φ is D-close to the left multiplication by g. In particular $d_G(g,1) = d_G(g,\Phi(1)) \leqslant D$, and since the word metric is invariant under left multiplication we also have that $d_G(g^{-1},1) = d_G(1,g) \leqslant D$. Then Φ is also $2D$-close to the conjugation by g, since for every $k \in G$ we have that

$$d_G(gkg^{-1},\Phi(k)) \leqslant d_G(gkg^{-1},gk)+d_G(gk,\Phi(k)) = d_G(g^{-1},1)+d_G(gk,\Phi(k)) \leqslant 2D$$

where again we used the left-invariance of the word metric. Choosing $k = h^l$ for every $l \in \mathbb{Z}$ we get that the infinite subgroups $H_1 = \langle ghg^{-1}\rangle$ and $H_2 = \langle \phi(h)\rangle$ lie at Hausdorff distance at most $2D$ (here we used that $\Phi|_H \equiv \phi$). But now [14, Proposition 9.4] states that, whenever G has a left-invariant proper metric (in our case, the word metric) and H_1, H_2 are two subgroups, for every D there exists a constant D' such that, if we denote by $N_R(S)$ the R-neighbourhood of a set S,

$$N_{2D}(H_1) \cap N_{2D}(H_2) \subseteq N_{D'}(H_1 \cap H_2)$$

Thus $N_{D'}(H_1 \cap H_2)$ is infinite because it contains H_1. This in turn implies that $H_1 \cap H_2$ is infinite, since balls in the word metric are finite. Therefore, there exist common powers $gh^m g^{-1} = \phi(h)^n$, as required. □

LEMMA 10.6. *Let* G *be a group acting coboundedly, non-elementarily on a hyperbolic space* $\mathcal{S}$*, and suppose the action admits loxodromic WPD elements. Let* $H \leqslant G$ *be a finite index subgroup. If* $E_G(G) = \{1\}$ *then* $E_G(H) = \{1\}$*.*

Proof. For every element $g \in G$ set

$$E_G(g) = \left\{k \in G \,|\, \exists m,n \in \mathbb{Z}\setminus\{0\} \text{ s.t. } kg^m k^{-1} = g^n\right\}.$$

It follows from the proof of [18, Lemma 6.18] that there exists a loxodromic WPD element g_0 such that $E_G(g_0) = \langle g_0\rangle \ltimes E_G(G)$ (there the notation $K(G)$ is adopted for $E_G(G)$). Since $E_G(G) = \{1\}$ we have that $E_G(g_0) = \langle g_0\rangle$.

Now, in [2, Lemma 5.6] it is proved that, whenever H is a non-virtually-cyclic subgroup such that $H \cap \mathcal{L}_{WPD} \neq \emptyset$, then

$$E_G(H) = \bigcap_{h \in H \cap \mathcal{L}_{WPD}} E_G(h).$$

Choose $k \in \mathbb{N}_{>0}$ such that $g_0^k \in H$, and notice that $E_G(g_0^k) = E_G(g_0)$ by definition. Thus $E_G(H)$ is a finite subgroup of $E_G(g_0^k)$, which is infinite cyclic, and therefore $E_G(H) = \{1\}$. □

The following statement is well-known to experts, but we provide a proof since we could not find a suitable reference.

LEMMA 10.7. *Let* $S = S_{g,b}$ *be a surface of finite type with genus* g *and* b *punctures, with* $(g,b) \notin \{(0,2),(0,3),(0,4),(1,0),(1,1),(1,2),(2,0)\}$*. Then* $MCG^{\pm}(S)$ *has no non-trivial finite normal subgroups. In other words* $E_{MCG^{\pm}}(MCG^{\pm}) = \{1\}$*.*

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

Proof. Let $N \trianglelefteq MCG^{\pm}(S)$ be a finite normal subgroup and let $f \in N$. If f fixes every isotopy class of simple closed curves then it must be the identity, for example by Ivanov's Theorem [28, Theorem 1] and its extension to lower genera [25, Theorem 1]. Then suppose by contradiction that $f(\gamma) \neq \gamma$ for some curve γ, and let T_γ be the corresponding Dehn twist. Since N is normal there exists $g \in N$ such that $fT_\gamma^2 = T_\gamma^2 g$. Thus, using how Dehn twists behave under conjugation, we have $T_{f(\gamma)}^{\pm 2} = fT_\gamma^2 f^{-1} = T_\gamma^2 g f^{-1}$, where the sign depends on whether f is orientation preserving or reversing. This can be rewritten as

$$T_\gamma^{-2} T_{f(\gamma)}^{\pm 2} = gf^{-1}. \tag{10.1}$$

Now, referring to the table at the end of [21, Subsection 3.5.2] we see that T_γ^2 and $T_{f(\gamma)}^2$ are the generators of a subgroup isomorphic to either $\mathbb{Z}^2$ or F_2, hence the left-hand side of (Equation 10.1) has infinite order. This is impossible, since the right-hand side is an element of the finite subgroup N. □

LEMMA 10.8. *For every $b \geqslant 7$ and for all large multiples* K*, $MCG^{\pm}(S_b)/DT_K$ has no finite normal subgroups.*

Proof. For short, we denote $MCG^{\pm}(S_b)$ simply by $MCG^{\pm}$. The mapping class group is acylindrically hyperbolic, hence by [18, Lemma 6.18] there exists a pseudo-Anosov element $g \in MCG^{\pm}$ such that $E_{MCG^{\pm}}(g) = \langle g \rangle \ltimes E_{MCG^{\pm}}(MCG^{\pm}) = \langle g \rangle$. Moreover, if we fix a curve $x \in \mathcal{C}$, arguing as in the proof of Corollary 3.8 we can find $\Theta > 0$ such that $\sup_{s \in \mathcal{C}, n \in \mathbb{Z}} d_s(x, g^n(x)) < \Theta$. Then for all large multiples K the projection map $\mathcal{C} \to \mathcal{C}/DT_K$ is an isometry on the axis $\{g^n(x)\}_{n \in \mathbb{Z}}$, thanks to Lemma 3.3. In particular the element $\overline{g} \in MCG^{\pm}/DT_K$ induced by g has infinite order, since $\overline{g}^n$ maps $\overline{x}$ to the projection of $g^n(x)$, which is not $\overline{x}$. Now, we claim that every finite normal subgroup $N \leqslant MCG^{\pm}/DT_K$ must be trivial. First notice that N moves $\overline{x}$ within distance M for some $M \geqslant 0$, since it is finite. Then the whole axis $\{\overline{g}^n(\overline{x})\}_{n \in \mathbb{Z}}$ is moved within Hausdorff distance M, since

$$\sup_{n \in \mathbb{Z},\, \overline{\phi} \in N} d_{\mathcal{C}/DT_K}\left(\overline{g}^n(\overline{x}), \overline{\phi} \circ \overline{g}^n(\overline{x})\right) = \sup_{n \in \mathbb{Z},\, \overline{\psi} \in N} d_{\mathcal{C}/DT_K}\left(\overline{g}^n(\overline{x}), \overline{g}^n \circ \overline{\psi}(\overline{x})\right),$$

where we used that N is normal. But then by left-invariance of the word metric we get that

$$\sup_{n \in \mathbb{Z},\, \overline{\psi} \in N} d_{\mathcal{C}/DT_K}\left(\overline{g}^n(\overline{x}), \overline{g}^n \circ \overline{\psi}(\overline{x})\right) = \max_{\overline{\psi} \in N} d_{\mathcal{C}/DT_K}\left(\overline{x}, \overline{\psi}(\overline{x})\right) \leqslant M$$

Now, let $\overline{\phi} \in N$ and let $\phi \in MCG^{\pm}$ be one of its preimages. Fix $n \in \mathbb{Z}$, let $l = [x, g^n(x)]$ be a geodesic and let $s = \phi(l)$, which is a geodesic between $\phi(x)$ and $\phi \circ g^n(x)$. Notice that $\sup_{s \in \mathcal{C}} d_s(\phi(x), \phi(g^n(x))) = \sup_{s \in \mathcal{C}} d_{\phi^{-1}(s)}(x, g^n(x)) < \Theta$, therefore both l and s project isometrically to geodesics $\overline{l} = [\overline{x}, \overline{g}^n(\overline{x})]$ and $\overline{s} = \overline{\phi}(\overline{l})$, again by Lemma 3.3. We can complete these two segments to a quadrilateral $\overline{Q}$ of vertices $\overline{x}, \overline{\phi}(\overline{x}), \overline{\phi} \circ \overline{g}^n(\overline{x}), \overline{g}^n(\overline{x})$, by adding two geodesic segments of length at most M. By Lemma 2.7 there exists a lift Q of $\overline{Q}$, and by Lemma

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

2.10 the lifts l', s' of $\bar{l}, \bar{s}$ are DT_K-translates of l, s respectively. Up to elements of DT_K we can assume that $l = l'$. Moreover, let $k \in DT_K$ be such that $k(s) = s'$. Setting $\psi = k \circ \phi$, which still induces $\overline{\phi} \in MCG^{\pm}/DT_K$, we see that the vertices of Q are $x, \psi(x), \psi(g^n(x)), g^n(x)$. But since Q lifts $\overline{Q}$ we have that $d_{\mathcal{C}}(x, \psi(x)) = d_{\mathcal{C}/DT_K}(\overline{x}, \overline{\phi}(\overline{x})) \leqslant M$, and similarly $d_{\mathcal{C}}(g^n x, \psi(g^n x)) \leqslant M$. Now, since M is independent of n we can choose n big enough (that is, x and $g^n(x)$ far enough on the axis) that ψ must belong to $E_{MCG^{\pm}}(g) = \langle g \rangle$ because of [18, Lemma 6.7] which says, roughly, that coarsely stabilising a large segment of an axis is equivalent to stabilising the whole axis (as noted in [18], the lemma has the same proof as [11, Proposition 6], which has more restrictive hypotheses). Thus $\overline{\phi} = \overline{g}^m$ for some $m \in \mathbb{Z}$, and we must have that $m = 0$ since $\overline{\phi}$ has finite order. □

As a consequence of Theorem 10.1 we can also describe the automorphism group of MCG/DT_K:

THEOREM 10.9. *For every $b \geqslant 7$ and for all large multiples K the following hold:*

- *$Aut(MCG(S_b)/DT_K) \cong MCG^{\pm}(S_b)/DT_K$;*
- *$Out(MCG(S_b)/DT_K) \cong \mathbb{Z}/2\mathbb{Z}$.*

We will need this auxiliary lemma:

LEMMA 10.10. *For every $b \geqslant 7$ and for all large multiples* K, *$MCG(S_b)/DT_K$ has trivial centre.*

Proof. Let $G = MCG^{\pm}(S_b)/DT_K$ and $H = MCG(S_b)/DT_K$. The centre $Z(H)$ is contained in $E_G(h)$ for every $h \in H$, by definition of $E_G(h)$. Then

$$Z(H) \leqslant \bigcap_{h \in H \cap \mathcal{L}_{WPD}} E_G(h) = E_G(H)$$

and we know that $E_G(H) = \{1\}$ by Lemma 10.6. □

Proof of Theorem 10.9. Again, let $G = MCG^{\pm}/DT_K$ and $H = MCG/DT_K$. Theorem 10.1 gives a surjective map $\Phi : G \to \mathrm{Aut}(H)$ mapping an element $g \in G$ to the restriction of the conjugation by g. Notice that Φ is injective when restricted to H, since this group has trivial centre by Lemma 10.10. Thus $\ker \Phi \cap H = \{1\}$, which in turn means that $\ker \Phi$ injects in the quotient $G/H \cong \mathbb{Z}/2\mathbb{Z}$. Hence $\ker \Phi$ is a finite normal subgroup of G, and it must be trivial by Lemma 10.8. This proves that $G \cong \mathrm{Aut}(H)$, and since H has trivial centre we also get

$$\mathrm{Out}(H) = \mathrm{Aut}(H)/\mathrm{Inn}(H) \cong G/H \cong \mathbb{Z}/2\mathbb{Z},$$

as required. □

Acknowledgements
G.M. would like to thank Roberto Frigerio for his good early career advice, indirectly leading to this paper. The authors would also like to thank Beatrice Pelloni for her crucial help in making it possible for G.M. to visit A.S. Finally, we thank the referee for the numerous suggestions on how to improve the paper.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

References

1 I. Agol. The virtual Haken conjecture. In D. Groves, and J. Manning, Eds. *Doc. Math.* Vol. 18, pp. 1045–1087 (EMS press, Bielefeld, 2013). With an appendix by Agol.

2 Y. Antolín, A. Minasyan and A. Sisto. Commensurating endomorphisms of acylindrically hyperbolic groups and applications. *Groups Geom. Dyn.* **10** (2016), 1149–1210.

3 J. Aramayona and C. J. Leiniger. Finite rigid sets in curve complexes. *J. Topol. Anal.* **5** (2013), 183–203.

4 J. Aramayona, and J. Souto. Rigidity phenomena in the mapping class group. In A. Papadopoulos (Ed.), *Handbook of Teichmüller theory*, Vol. VI, pp. 131–165 (EMS press, Bielefeld, 2016).

5 J. Behrstock, M. F. Hagen and A. Sisto. Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups. *Proc. Lond. Math. Soc. (3).* **114** (2017), 890–926.

6 J. Behrstock, M. F. Hagen and A. Sisto. Hierarchically hyperbolic spaces II: Combination theorems and the distance formula. *Pacific J. Math.* **299** (2019), 257–338.

7 J. Behrstock, M. F. Hagen and A. Sisto. Quasiflats in hierarchically hyperbolic spaces. *Duke Math. J.* **170** (2021), 909–996.

8 J. Behrstock, M. Hagen, A. Martin and A. Sisto. A combinatorial take on hierarchical hyperbolicity and applications to quotients of mapping class groups. *J. Topol.* **17** (2024), 94 Id/No e12351.

9 J. Behrstock, B. Kleiner, Y. Minsky and L. Mosher. Geometry and rigidity of mapping class groups. *Geom. Topol.* **16** (2012), 781–888.

10 M. Bestvina, K. Bromberg and K. Fujiwara. Constructing group actions on quasi-trees and applications to mapping class groups. *Publ. Math. Inst. Hautes Études Sci.* **122** (2015), 1–64.

11 M. Bestvina and K. Fujiwara. Bounded cohomology of subgroups of mapping class groups. *Geom. Topol.* **6** (2002), 69–89.

12 B. H. Bowditch. Rigidity of the strongly separating curve graph. *Michigan Math. J.* **65** (2016), 813–832.

13 B. H. Bowditch. Large-scale rank and rigidity of the Weil-Petersson metric. *Groups Geom. Dyn.* **14** (2020), 607–652.

14 G. Christopher Hruska. Relative hyperbolicity and relative quasiconvexity for countable group. *Algebr. Geom. Topol.* **10** (2010), 1807–1856.

15 M. Clay and J. Mangahas. Hyperbolic quotients of projection complexes. *Groups Geom. Dyn.* **16** (2022), 225–246.

16 R. Coulon and Z. Sela. Equations in burnside groups (2021), *arXiv preprint arXiv:2112.07409*.

17 F. Dahmani. The normal closure of big Dehn twists and plate spinning with rotating families. *Geom. Topol.* **22** (2018), 4113–4144.

18 F. Dahmani, V. Guirardel and D. Osin. Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *Mem. Am. Math. Soc.* **245** (2017), v+152.

19 F. Dahmani, M. F. Hagen and A. Sisto. Dehn filling Dehn twists. *Proc. Roy. Soc. Edinburgh Sect. A.* **151** (2021), 28–51.

20 M. Dehn. *Papers on Group Theory and topology*. Translated from the German and with introductions and an appendix by John Stillwell, With an appendix by Otto Schreier (Springer-Verlag, New York, 1987).

21 B. Farb and D. Margalit. *A primer on mapping class groups*. Of *Princeton Mathematical Series.*, Vol. 49 (Princeton University Press, Princeton, NJ, 2012).

22 B. Farb and L. Mosher. Convex cocompact subgroups of mapping class groups. *Geom. Topol.* **6** (2002), 91–152.

23 L. Funar. On the TQFT representations of the mapping class groups. *Pacific J. Math.* **188** (1999), 251–274.

24 U. Hamenstaedt. Geometry of the mapping class groups III: quasi-isometric rigidity (2005).

25 M. Korkmaz. Automorphisms of complexes of curves on punctured spheres and on punctured tori. *Topol. Appl.* **95** (1999), 85–111.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press

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

27 J. Nielsen. Untersuchungen zur Topologie der geschlossenen zweiseitigen Flächen. *Acta Math.* **50** (1927), 189–358.

28 N. V. Ivanov. Automorphisms of complexes of curves and of Teichmüller spaces. *Travaux en Cours*, Vol.56, pp. 113–120 (Hermann, Paris, 1997).

29 R. C. H. Wade. Uniform bounds for bounded geodesic image theorems. *J. Reine Angew. Math.* **709** (2015), 219–228.

30 R. P. Kent IV and Should be C. J. Leiniger. Shadows of mapping class groups: capturing convex cocompactness. *Geom. Funct. Anal.* **18** (2008), 1270–1325.

31 R. E. Schwartz. The quasi-isometry classification of rank one lattices. *Inst. Hautes Études Sci. Publ. Math.* **82** (1996), 133–168.

https://doi.org/10.1017/prm.2025.23 Published online by Cambridge University Press