Università di Pisa

Dipartimento di Matematica
Corso di Laurea Magistrale in Matematica

Tesi di Laurea Magistrale

Rigidity of Mapping Class Groups mod powers of Dehn Twists

Candidato:
Giorgio Mangioni
Matricola 587287

Relatore:
Prof. Roberto Frigerio

Correlatore:
Prof. Alessandro Sisto
Heriot-Watt University, Edimburgh

Anno Accademico 2022-23
9 Giugno 2023

Introduction

Dehn twist quotients

Let S be a compact, oriented surface, and let $MCG(S)$ be the group of its orientation-preserving self-homeomorphisms, up to isotopy. Among these are the Dehn twists, which are particular homeomorphisms that coincide with the identity map outside an annulus. Let DT_K be the normal subgroup generated by all K-th powers of Dehn Twists.
This thesis aims to study algebraic and geometric properties of the quotient $MCG(S)/DT_K$ that arise when K is suitably large. One reason of interest for these quotients is that they can be regarded as "Dehn filling quotients" of mapping class groups, as pointed out and explored in [DHS21]. In fact, Thurston's Dehn filling theorem [Thu97] states that, if M is a three-dimensional hyperbolic manifold with torus boundary, and one chooses a simple closed curve γ on the boundary, then the complete manifold obtained by filling the boundary with a solid torus, whose meridian is γ, is again hyperbolic, except for finitely many curves γ. In other words, if we take the fundamental group $\pi_1(M)$, which is relatively hyperbolic in the sense of Gromov [Gro87], then the quotient by the subgroup normally generated by γ is hyperbolic, except in finitely many cases. Now, mapping class groups are not hyperbolic, except if the surface is very simple, but it is still natural to consider the quotient by the subgroups generated by powers of Dehn twists, which correspond to the curves on the surface that one wants to "fill", and ask whether these quotients share the same "hyperbolicity features" as the mapping class group if one avoids a certain number of "bad" powers. This was done in both [DHS21] and [BHMS20], where the authors also use properties of these groups to relate questions on finite quotients of mapping class groups to residual finiteness of certain hyperbolic groups.
These quotients of mapping class groups also appear naturally in a somewhat unrelated are of mathematics, namely the study of topological quantum field theories, see e.g. [Fun99]. Yet another reason of interest is simply that given a group G and a collection of elements $g_1, \ldots, g_n$, 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 "stabilizing" 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, in order to avoid some complications that arise in the case of surfaces with genus.
The first result we state, which answers [BHMS20, Question 3] in the case of punctured spheres, is quasi-isometric rigidity of our groups of interest: if a group G is quasi-isometric to $MCG(S)/DT_K$, which roughly speaking means that G "has the same large scale geometry as $MCG(S)/DT_K$" (in a sense that shall be clarified in Section 5.1), then these two groups are isomorphic, up to taking finite index subgroups or quotienting by finite subgroups. To be more precise, 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 [BKMM12], see also [Ham07], but our argument is closer to the proof given in [BHS21, Section 5]. We cannot use the key theorem from [BHS21, Section 5] as stated, but we will instead give another version of it, Theorem 5.2.15, which is of independent interest. Roughly speaking, this theorem allows one to extract an automorphism of a certain graph from a quasi-isometry of a space with certain "hyperbolicity features".
In our case of interest, this graph 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 several arguments due to Bowditch, especially from [Bow20], where the author proved quasi-isometric rigidity of another metric graph associated to a surface, the pants graph. Indeed, in some sense quotients of mapping class groups by powers of Dehn twists more closely resemble the pants graph than the mapping class group, as will be clear in Chapter 6. This is why, in order to illustrate the strategy of the proof of Theorem 1, we also provide an alternative argument for a sizable part of the proof of quasi-isometric rigidity of pants graphs of spheres:

Theorem 2 (QI rigidity for pants graphs of spheres; [Bow20, Theorem 1.4])**.** *Let $S = S_{0,b}$ be a punctured sphere, with $b \geqslant 7$ punctures, and $\mathbb{P}(S)$ be its pants graph. Every self-quasi-isometry of $\mathbb{P}(S)$ is at bounded distance from the isometry induced by some element of the mapping class group.*

From quasi-isometric rigidity descend some other rigidity results. The first one can be regarded as a form of algebraic rigidity, and roughly speaking states that an automorphism of our quotients can only be a conjugation. Here, $MCG^{\pm}$ denotes the extended mapping class group, where orientation-reversing mapping classes are allowed.

Theorem 3 (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. *$Aut(MCG(S)/DT_K) = MCG^{\pm}(S)/DT_K$, and $Out(MCG(S)/DT_K) \cong \mathbb{Z}/2\mathbb{Z}$;*
2. *The abstract commensurator of $MCG^{\pm}(S)/DT_K$ is trivial, that is, 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 [Iva97]. There, he studied simplicial automorphisms of the *curve graph*, first introduced by Harvey [Har81], whose vertices are the simple closed curves on the surface, up to isotopy, and whose edges correspond to disjointness. Ivanov showed that every automorphism of the curve graph comes from a mapping class of the surface, which acts on the set of curves and preserves disjointness. Similarly, our quasi-isometric rigidity and in turn algebraic rigidity results rely on an analogue of Ivanov's theorem for quotients of curve graphs:

Theorem 4 (Combinatorial 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 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 4.4.1 and the discussion in the outline section. Moreover, an analogue of the theorem also holds for $b = 4$, but in that case the map has finite kernel, see Theorem 4.1.2.

Outline of proofs

Combinatorial rigidity

We will first prove combinatorial rigidity, Theorem 4. 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 [DHS21] to show that $\mathcal{C}/DT_K$ is hyperbolic by lifting geodesic triangles, and was developed further in [BHMS20]. We push these techniques even further to show that the graph constructed by Aramayona and Leiniger in [AL13] can be lifted. This graph is a finite rigid set for $\mathcal{C}$, that is, any isometric embedding of this finite subgraph extends to a unique mapping class. Consider now an automorphism ϕ of $\mathcal{C}/DT_K$ and a copy X of one such graph in $\mathcal{C}/DT_K$. One can consider the following diagram, where the hats denote lifts and π is the quotient projection:

$$\begin{tikzcd} \widehat{X} \arrow[r, "g"] \arrow[d, "\pi"] & \widehat{\phi(X)} \arrow[d, "\pi"] \\ X \arrow[r, "\phi"] & \phi(X) \end{tikzcd}$$

From the diagram we obtain a candidate element of the extended mapping class group that induces ϕ, namely the element $g \in MCG^{\pm}$ mapping $\widehat{X}$ to $\widehat{\phi(X)}$. Showing that this candidate is actually the desired element requires more care, and further lifts, but this is the basic idea.
In order to prove quasi-isometric rigidity, we will also need a version of the combinatorial rigidity theorem for some subgraphs of $\mathcal{C}/DT_K$, spanned by classes of curves that cut out certain subsurfaces. This is shown in Chapter 7, where we combine arguments due to Bowditch with further lifting techniques.

Quasi-isometric rigidity

To prove quasi-isometric rigidity, we will use a result about self quasi-isometries of *hierarchically hyperbolic* spaces, defined by Behrstock, Hagen and Sisto in [BHS17] and which, roughly speaking, are those spaces with the same "hyperbolicity features" as the mapping class group. For this class of spaces, [BHS21, Theorem 5.7] states that if X is hierarchically hyperbolic and satisfies some additional assumptions then every self-quasi-isometry f of X induce an automorphism ϕ of a certain graph called the hinge graph. This is a graph that encodes the "standard flats" of X, that is, it roughly describes how isometric copies of the Euclidean space $\mathbb{R}^n$ inside X are arranged. Actually, the additional assumptions of that theorem do not apply to our case, but we show that a similar statement still holds under different hypotheses that do apply; this is Theorem 5.2.15. The new requirements are also met by the pants graph, and indeed in Chapter 6 we show how to recover a result of Bowditch [Bow20] about quasi-isometric rigidity of pants graphs of spheres from Theorem 5.2.15. Our argument is not completely new, since it relies on lemmas from [Bow16] and Sections 6 and 7 of [Bow20].
Both in the case of pants graphs and in the case of mapping class groups mod powers of Dehn twists, starting with an automorphism ϕ of the hinge graph, we use combinatorial arguments to

recover automorphisms of the graphs from Chapter 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.

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 9.0.1, using results of Antolin, Minasyan and Sisto from [AMS16]. 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 9.0.7.

Outline of Chapters

In Chapter 1 we recall some basic definitions and facts about mapping class groups and curve graphs. Chapter 2 introduces the basic tools of lifting and projecting that we will use throughout the thesis, recalling and sometimes extending results from [DHS21, BHMS20]. In Chapter 3 we study the combinatorial properties of the finite rigid sets from [AL13]. The key results here are that our chosen finite rigid sets map injectively to the quotient, see Theorem 3.3.3, and that they can be lifted from quotients of curve graphs to curve graphs, as we prove in Theorem 3.4.1. We then use this in Chapter 4 to show the combinatorial rigidity Theorem 4, see Theorem 4.4.1. In Chapter 5 we first review some background on coarse geometry and geometric group theory, and then gather the relevant properties of hierarchically hyperbolic spaces. Finally, we show that quasi-isometries of suitable hierarchically hyperbolic spaces induce automorphisms of an associated graph, Theorem 5.2.15. In Chapter 6 we show how to use Theorem 5.2.15 to recover quasi-isometric rigidity of the pants graph (still using some of the original arguments), as this will serve as an outline for the proof of the quasi-isometry rigidity Theorem 1. In Chapter 7 we prove further combinatorial rigidity results that will be needed in Chapter 8 to prove Theorem 1, which is a combination of Theorem 8.0.16, which says that quasi-isometries are all at controlled distance from left-multiplications, and the general Lemma 8.0.18. Finally, in Chapter 9 we show algebraic rigidity, Theorem 3, see Theorems 9.0.1 and 9.0.8.

Acknowledgements

This thesis is an extended version of the paper [MS22], which is a joint work with Alessandro Sisto and was written during the six months I spent at the Heriot-Watt University of Edinburgh as a visiting student. I would like to thank him for always pushing me to do more and for his daily efforts in the making of that article. I would also like to thank Roberto Frigerio for his good early career advice, indirectly leading to this thesis. Finally, I thank my girlfriend for the constant support, even while I was abroad.

Contents

Chapter 1

Setting and notation

In this Chapter we gather all the notions and facts we need in order to understand and prove the Combinatorial Rigidity Theorem 4, postponing the prerequisites for the second half of the thesis to Chapter 5. We will always assume that the reader is familiar with the basic definitions and examples from the following areas:

- group theory (groups and subgroups, indices, centralizers, normal subgroups...);
- topology (topological spaces, homeomorphisms, homotopies and isotopies, coverings...);
- metric geometry (metric spaces, neighborhoods, balls, curves, length metric...).

1.1 The mapping class group of a surface

For this and the following sections we mainly follow the book of Farb and Margalit [FM12]. By *surface of finite type*, or simply surface, we mean a compact, connected, oriented 2-manifold that is possibly *punctured*, that is, with a finite set of points removed (of course, after we puncture a compact surface, it ceases to be compact). The classification theorem of surfaces (see e.g. [Tho92] for a proof) states that any closed, connected, orientable surface is homeomorphic to the connected sum of a 2-dimensional sphere with $g \geqslant 0$ tori, and any finite type surface is obtained from a closed surface by removing $b \geqslant 0$ punctures. Then the set of homeomorphism types of compact surfaces is in bijective correspondence with the set $\{(g, b) : g, b \geqslant 0\}$.
Notice that, from a topological viewpoint, a surface with one point removed is homeomorphic to the same surface with a disk removed, thus the classification does not see the difference between punctures and boundary components. However, sometimes it is better to specify a surface by the triple (g, b, n), where n is the number of boundary components and b is the number of punctures.

Definition 1.1.1. Let $\mathrm{Homeo}(S, \partial S)$ denote the group of homeomorphisms of S that restrict to the identity on the boundary. The *extended mapping class group* of S, denoted $MCG^{\pm}(S)$, is the group of isotopy classes of elements of $\mathrm{Homeo}(S, \partial S)$, where isotopies are required to fix the boundary pointwise. The *mapping class group* of S, denoted $MCG(S)$, is the index two subgroup of $MCG^{\pm}(S)$ generated by classes of orientation-preserving homeomorphisms.

This is where the difference between boundary components and punctures is relevant: an element of $MCG(S)$, which we also call a *mapping class*, has to fix the boundary components pointwise, but may permute the punctures.
The *support* of a mapping class f is the subsurface (defined up to isotopy) given by the closure of the set of points $x \in S$ such that $f(x) \neq x$.

1.1.1 Curves and (half) Dehn twists

Definition 1.1.2. By a *curve* on a surface S we mean a simple closed curve, that is, the image of a continuous, injective map from the circle $\mathbb{S}^1$ to S. A curve is *essential* if it is not homotopic to a point, a puncture, or a boundary component.

A curve α is *separating* if $S\backslash\alpha$ has two connected components, and in this case we say that α *bounds* these two subsurfaces. Notice that if the surface S is a *sphere* (i.e. it has no genus, as will almost always be the case in this thesis) then every curve is separating, e.g. by the Jordan curve Theorem [Jor91].
The following definition is due to Dehn (see [Deh38, Section 2b]). Consider the Euclidean plane 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 $T(\theta, r) = (\theta + 2\pi r, r)$, which is an orientation-preserving homeomorphism that fixes ∂A pointwise. We call T the *left twist* of our annulus. Note that instead of using $\theta + 2\pi r$ we could have used $\theta - 2\pi r$, thus obtaining what is called a *right* twist.

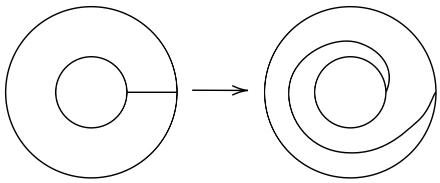

Figure 1.1: The left twist T of an annulus.

Now let S be an oriented surface and let α be a curve in S. Let N be a regular neighborhood 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, denoted T_α, is the homeomorphism

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

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. By construction, the support of T_α is the annulus A.

Lemma 1.1.3 (Properties of Dehn Twists)**.** *Let α be a curve on a surface S. The following facts hold for the Dehn twist T_α:*

- *If α is essential then T_α has infinite order;*
- *$T_\alpha(\beta) = \beta$ iff $i(\alpha, \beta) = 0$;*
- *$T_\alpha = T_\beta$ iff α and β are isotopic;*
- *If $f \in MCG^\pm(S)$ then $fT_\alpha f^{-1} = T^{\pm}_{f(\alpha)}$, where the sign is positive if and only if f is orientation-preserving.*

Proof. These facts are well-known, see e.g. [FM12, Sections 3.2 and 3.3] for proofs. ☐

If the surface has punctures, it is sometimes possible to define a *half* Dehn twist as follows. Let α be a curve that bounds a twice-punctured disk D. Let D_0 be the disk of radius 3 and centered at the origin inside the Euclidean plane, with the points $(\pm 1, 0)$ removed, and let $\phi : D_0 \to D$ be an orientation-preserving homomorphism mapping ∂D_0 to α and punctures to punctures. Then define the *half Dehn twist* H_α as (the isotopy class of) the composition $H_\alpha = \phi \circ H \circ \phi^{-1}$, where in polar coordinates (θ, r) the map H is defined as follows:

$$H(\theta, r) = \begin{cases} (\theta + \pi, r) \text{ if } r \leqslant 2, \\ (\theta + (3 - r)\pi, r) \text{ if } 2 \leqslant r \leqslant 3. \end{cases}$$

In other words, H does a half-rotation around the origin and swaps the punctures.

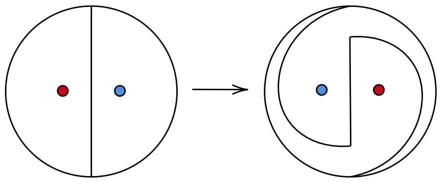

Figure 1.2: The half twist H of a twice-punctured disk. Notice that the punctures are swapped.

With similar arguments as for the proof of Lemma 1.1.3 one can show the following:

Lemma 1.1.4 (Properties of half Dehn Twists)**.** *Let α be a curve on a punctured surface S that bounds a twice-punctured disk. The following facts hold for the half Dehn twist H_α:*

- *$H_\alpha^2 = T_\alpha$, thus H_α has infinite order if α is essential;*
- *if $f \in MCG^{\pm}(S)$ then $fH_\alpha f^{-1} = H_{f(\alpha)}^{\pm}$, where the sign is positive if and only if f is orientation-preserving.*

We denote by $PMCG(S)$ the *pure* mapping class group, that is, the finite index subgroup of $MCG(S)$ generated by those mapping classes which do not permute the punctures. The quotient $MCG(S)/PMCG(S)$ is isomorphic to the symmetric group on the number of punctures and is generated by a finite number of half Dehn twists which act as transpositions. Moreover, the following celebrated theorem was first proven by Dehn [Deh87] and Lickorish [Lic64] independently:

Theorem 1.1.5. *The pure mapping class group of a surface is generated by a finite number of Dehn twists around non-separating curves. Therefore the mapping class group is finitely generated.*

We will be interested in the special case of the previous theorem when $S = S_b$ is a sphere with punctures, which actually follows from work by Artin [Art47] as explained in [FM12, Chapter 9]:

Theorem 1.1.6. *Let S_b be a sphere with $b \geqslant 4$ punctures. Order the punctures $\{1, \ldots, b\}$ and let β_i be a curve surrounding the punctures $i, i+1$, so that*

$$i(\beta_i, \beta_j) = \begin{cases} 2 \textit{ if } |i - j| = 1\,(mod\ b), \\ 0 \textit{ otherwise.} \end{cases}$$

Then the half Dehn twists around $\beta_1, \ldots, \beta_b$ generate $MCG(S_b)$.

An example of such curves are the minimal curves in the finite rigid set, defined in Chapter 3.

1.2 The curve graph

1.2.1 Simplicial graphs

This subsection gathers the basic notions of graph theory. The reader which is familiar with what a simplicial graph is could easily skip this section.

Definition 1.2.1. An *abstract simplicial graph* is a pair $X = (V, E)$, where V is a set of *vertices* and E is a set of unordered pairs of distinct vertices, that we call *edges*.

In other words, we do not allow multiple edges between the same two vertices, nor edges joining a vertex to itself. Moreover, we are considering *undirected* graphs, that is, edges have no preferred direction. We will often refer to V as the *0-skeleton* $X^{(0)}$ of X.

Definition 1.2.2. A *combinatorial path* is a sequence $x_0, x_1, \ldots, x_n \in V$ such that for every $i = 0, \ldots, n-1$ the vertices x_i and x_{i+1} are *adjacent*, i.e., they are connected by an edge. We say that the path *connects* the vertices x_0 and x_n, which are the *endpoints* of the path. A graph is *connected* if every two vertices are connected by some path.

A simplicial graph can be endowed with a metric by saying that every edge has length 1. In other words, a path $x_0, x_1, \ldots, x_n$ has length n, and for every two vertices $v, w \in V$ the distance $d_X(v, w)$ is given by the minimal length of a path connecting them. Therefore with a slight abuse of notation we will often view a graph X as the metric space defined as follows:

- for every $e \in E$ take an isometric copy of the unit interval $[0, 1]$ and identify the endpoints of e with the endpoints of the interval;
- glue two edges along an endpoint if they share a common vertex;
- endow the resulting space with the *length metric*.

We will call this metric space the *geometric realization* of the abstract simplicial graph. Notice that a graph is connected, in the sense of Definition 1.2.2, if and only if its geometric realization is path connected in the usual sense.

Definition 1.2.3. We say that a path is a *geodesic segment* if it realizes the distance between its endpoints, that is, if it has minimal length. One can similarly define geodesic *rays* and *lines*, which are infinite (resp. bi-infinite) paths which realize the distance between any two of their vertices. We will often refer to these objects simply as geodesics, when it will be clear from the context whether we are talking about lines, rays, or segments.

Definition 1.2.4. A *simplex* of a graph X is a (possibly empty) collection of pairwise adjacent vertices. The *dimension* of a simplex is the number of its vertices minus 1.

Vertices and edges are themselves simplices, respectively, of dimension 0 and 1. Moreover, it is sometimes convenient to see the empty set $\varnothing$ as a simplex without vertices.

Definition 1.2.5. A *subgraph* of $X = (V, E)$ is a graph $X' = (V', E')$ such that $V' \subseteq V$ and $E' \subseteq E$. A subgraph $X' \subseteq X$ is the *induced* subgraph spanned by V' if two vertices $v, w \in V'$ are adjacent in X' if and only if they are adjacent in X.

Definition 1.2.6. If K, L are disjoint induced subgraphs of X such that every vertex of K is adjacent to every vertex of L, then their *join* $K \star L$ is the induced subgraph with vertex set $K^{(0)} \cup V^{(0)}$.

Definition 1.2.7. For a simplex $\Delta \subseteq X$, the *link* $\mathrm{Lk}(\Delta)$ is the union of all simplices Σ of X such that $\Sigma \cap \Delta = \varnothing$ and $\Sigma \star \Delta$ is a simplex of X. The *star* of Δ is given by $\bigstar\Delta = \Delta \star \mathrm{Lk}(\Delta)$.

Definition 1.2.8. Let X, Y be two graphs. A *graph morphism* $\phi : X \to Y$ consists of a map $X^{(0)} \to Y^{(0)}$ mapping adjacent vertices to adjacent vertices. A graph morphism is *locally injective* if it is injective when restricted to the star of any vertex; it is an *isometric embedding* if for every two vertices $v, w \in X^{(0)}$ we have that

$$d_X(v, w) = d_Y(\phi(v), \phi(w)).$$

An isometric embedding which is surjective (hence also bijective) is called an *isomorphism.* An isomorphism from a graph X to itself is called an *automorphism.* The group of automorphisms of a graph will be denoted by $\mathrm{Aut}(X)$.

1.2.2 Definition of the curve graph

Recall that by curve we mean a simple closed curve on a surface S. We will often conflate a curve α with its *free homotopy class*, that is, the class of all curves on S which admit an unbased homotopy to α.
The *geometric intersection number* between free homotopy classes a and b of curves in a surface S is defined as the minimal number of transverse intersection points between a representative curve in the class a and a representative curve in the class b:

$$i(a, b) = \min\{\#(\alpha \cap \beta)\,|\alpha \in a, \beta \in b\}.$$

We often abuse notation slightly by writing $i(\alpha, \beta)$ for the intersection number between the homotopy classes of simple closed curves α and β. We will say that two curves intersect essentially, or simply that they intersect, if their intersection number is positive. Notice that, since the surface is oriented, the regular neighborhood of a curve is an oriented annulus, and therefore every curve can be homotoped away from itself. This shows that the intersection number of a curve with itself is always zero.
The following definition is due to Harvey [Har81]. First define the *complexity* of a surface with genus g, b punctures, and n boundary components as $3g + b + n - 3$.

Definition 1.2.9. Let S be a surface of complexity at least 2. The *curve graph* of a surface S is the simplicial graph whose vertices are the free homotopy classes of simple, closed, essential curves, and where two curves are joined by an edge if they are *disjoint*, that is, if $i(\alpha, \beta) = 0$. We denote by d_S the distance in the curve graph.

Notice that the definition does not see if we replace the punctures with boundaries, because we only consider essential curves (which cannot be homotopic to punctures or boundary components). Therefore, unless otherwise stated, we will always assume that our surfaces only have punctures.
By construction, simplices of the curve graphs are tuples of pairwise disjoint essential curves, and we can cut S along these surfaces. In particular, a maximal simplex must cut the surface into the finest subdivision possible. It can be shown (see e.g. [FM12, Section 8.3.1]) that such a subdivision decomposes S into *pairs of pants*, that is, subsurfaces homeomorphic to a sphere with three disks removed; therefore we say that a maximal simplex in the curve graph is a *pants decomposition.* One can then easily see that the complexity of a surface corresponds to the number of curves in a pants decomposition. This is why we asked the complexity to be at least 2 in Definition 1.2.9, so that the curve graph has some edges, and it is actually connected by a

result of Lickorish [Lic64]. However, for surfaces of complexity 1 the graph defined in Definition 1.2.9 is just a collection of countably many points with no edges. Therefore, in order to define the curve graph of a complexity 1 surface, one needs to replace the adjacency relation in order to detect the minimum intersection number possible:

- If S is a torus with one puncture, two curves α, β are adjacent if and only if $i(\alpha, \beta) = 1$;
- If S is a sphere with four punctures, two curves α, β are adjacent if and only if $i(\alpha, \beta) = 2$.

In both cases, the curve graph is isomorphic to the 1-skeleton of the *Farey complex*, which is the triangulation of the hyperbolic plane depicted in Figure 1.3 (see e.g. [Min96] for a proof).

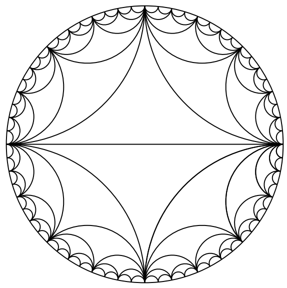

Figure 1.3: The Farey complex.

The following lemma summarizes some well-known properties of the Farey complex (see e.g. Bowditch-Epstein [BE88], Hatcher-Thurston [HT80], and Series [Ser85]):

Lemma 1.2.10. *The following facts hold for the Farey complex:*

- *Every edge belongs to exactly two triangles (that is, 2-dimensional simplices).*
- *For every two triangles T, T' there exists a sequence of triangles $T = T_1, \dots, T_k = T'$ such that T_i and T_{i+1} share an edge.*
- *Any automorphism of the Farey complex is uniquely determined by its restriction to a triangle.*

Remark 1.2.11. If $S = \bigsqcup_{i=1}^k S_i$ is a disconnected subsurface, and each connected component S_i has complexity at least 1, then we can define the curve graph of S as in Definition 1.2.9, which is a nontrivial join since every curve lying on a component is disjoint from every curve lying on another one. Hence curve graphs of disconnected subsurfaces have diameter at most 2, and therefore are rarely interesting from a large-scale viewpoint.

1.2.3 Annular curve graph

Recall that an annulus A on a surface S is the regular neighborhood of a curve $\gamma \subset S$, and we say that the annulus is essential in S if so is its core. We cannot define its curve graph as before, since it would be empty: the only simple closed curve on A is the core, which is homotopic to both boundary components and therefore not essential in A. Instead, following Masur and Minsky [MM00], we fix a point $x_0 \in \gamma$ and an orientation of γ, so that we can see γ as an element of the fundamental group $\pi_1(X, x_0)$. Let $p_\gamma : S_\gamma \to S$ be the *annular covering*, that is, the covering associated to the subgroup $\langle\gamma\rangle \leqslant \pi_1(X, p)$, to which A lifts homeomorphically. There is a natural compactification of S_γ to a closed annulus $\widehat{S_\gamma}$.

Definition 1.2.12. The *annular curve graph* associated to the curve γ is the simplicial graph 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).

One sees that the definition does not depend on the representative of the isotopy class of γ, nor on the choices of a base point x_0 and of an orientation for γ. We will denote the annular curve graph of the annulus with core γ by $\mathcal{C}(\gamma)$, and by d_γ the distance in this graph.

1.3 Subsurface projections

Let S be a surface and Y be a *subsurface*, that is, Y is a (possibly disconnected) surface and there is an embedding $Y \to S$. A subsurface is *essential* if it is not a disk, a pair of pants or an annulus with core a boundary curve of S. We will always conflate a subsurface with the image of the embedding, and two subsurfaces will be considered equivalent up to isotopy.
The set-theoretic boundary of Y inside S is a collection of disjoint curves on S, which we call the *relative boundary* of Y inside S and denote by ∂Y. Unless otherwise stated, we will adopt the convention that the curves of the relative boundary do not lie on Y, since they can be homotoped outside Y. Thus we will always see Y as a surface with punctures but without boundary, and every puncture of Y will correspond either to a puncture of S or to a relative boundary curve via the inclusion $Y \to S$.
Two subsurfaces are said to be:

- *disjoint* if (up to isotopy) they have empty intersection;
- *nested* if (up to isotopy) one is contained in the other;
- *transverse* otherwise.

In particular, two subsurfaces Y and Z are transverse if and only if one of the components of ∂Y intersects one of the components of ∂Z.
Given a subsurface $Y \subset S$, if Y has complexity at least 2 we can see the curve graph of Y as the induced subgraph of $\mathcal{C}(S)$ spanned by those curves which (up to isotopy) lie on Y. Equivalently, $\mathcal{C}Y$ can be seen as the link of a simplex $\Delta \subset \mathcal{C}S$, which corresponds to a pants decomposition of $S \backslash Y$ (including the relative boundary). This is not true if Y has complexity 1, since its curve graph has a different adjacency relation, nor if Y is an annulus, whose annular curve graph is different in nature.
Now let Y be a subsurface of complexity at least 1. Again following Masur and Minsky [MM00] we will define a set valued projection $\pi_Y : \mathcal{C}(S) \to 2^{\mathcal{C}(Y)}$, where $2^{\mathcal{C}(Y)}$ denotes the set of subsets of $\mathcal{C}(Y)$. Let $\alpha \in \mathcal{C}(S)$ be a curve. There are three possible cases:

- If $\alpha \subset Y$ then $\alpha \in \mathcal{C}(Y)$, and we set $\pi_Y(\alpha) = \alpha$.
- If α is disjoint from Y we set $\pi_Y(\alpha) = \emptyset$.
- Otherwise α intersects the boundary of Y, and $\alpha \cap Y$ is a collection of finitely many disjoint arcs $a_1, \dots, a_k$. For every arc a_i let N_i be a regular neighborhood inside Y of the union of this arc and the component(s) of ∂Y on which its endpoints lie, and let ∂N_i be the boundary components of N_i inside Y which are essential curves of Y. Finally, we let $\pi_Y(\alpha) = \bigcup_i \partial N_i$.

With a similar procedure one can define the *annular projection* $\pi_\gamma : \mathcal{C}(S) \to 2^{\mathcal{C}(\gamma)}$, where γ is an essential curve:

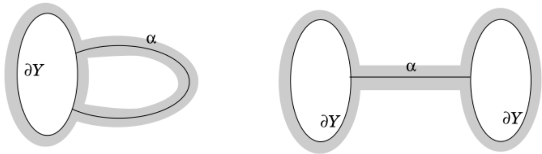

Figure 1.4: The regular neighborhood N of one of the arcs coming from α and the component(s) of ∂Y that it intersects. Notice that N can have one or two boundary curves, depending on whether the endpoints of the arc lie on the same component or not. This is [MM00, Figure 6].

- If α is a simple closed curve in S crossing γ then a lift of α to S_γ has at least one component that connects the two boundaries of $\widehat{S_\gamma}$, and together these components make up a (finite) set of diameter 1 in $\mathcal{C}(\gamma)$. Let $\pi_\gamma(\alpha)$ be this set.
- If α does not intersect γ essentially (including the case when $\gamma = \alpha$ is the core of the annulus) then $\pi_\gamma(\alpha) = \varnothing$.

We will say that a *domain* is a subsurface Y which is either an essential annulus or a subsurface of complexity at least 1 (which are the cases for which there is a well-defined subsurface projection). Given two curves $\alpha, \beta \in \mathcal{C}(S)$ and a domain Y, if both $\pi_Y(\alpha)$ and $\pi_Y(\beta)$ are non-empty we define $d_Y(\alpha, \beta) := d_{\mathcal{C}(Y)}(\pi_Y(\alpha), \pi_Y(\beta))$, where the latter distance is the minimum distance between pairs of points. The following is [MM00, Lemma 2.3]:

Lemma 1.3.1. *Let Y be a domain. For any simplex $\Delta \subset \mathcal{C}(S)$, if $\pi_Y(\Delta) \neq \varnothing$ then $diam_Y(\Delta) \leqslant 2$.*

Corollary 1.3.2. *The projection π_Y sends curves to uniformly bounded finite subsets. Moreover, it is 2-Lipschitz (where the distance between two finite subsets is the minimum distance between pairs of points).*

Proof. The first statement follows from Lemma 1.3.1 with Δ equal to a point. The second follows with Δ equal to an edge, because then the distance between the images of two curves is at most twice the length of a geodesic path between them. □

Subsurface projection also makes sense between two domains Y and Z. For example, if Y is nested in Z we can restrict the projection $\pi_Y : \mathcal{C}(S) \to 2^{\mathcal{C}(Y)}$ to $\mathcal{C}(Z)$ and get a map $\rho^Z_Y : \mathcal{C}(Z) \to 2^{\mathcal{C}(Y)}$. Conversely, we can define $\rho^Y_Z \subset \mathcal{C}(Z)$ as the relative boundary of Y inside Z, which is a simplex in $\mathcal{C}(Z)$. Notice that we can see ρ^Y_Z as a constant map $\mathcal{C}(Y) \to 2^{\mathcal{C}(Z)}$, which is why the notation is similar. Finally, if Y and Z are transverse then there is at least a curve in ∂Y which intersects a curve in ∂Z, and therefore the projections $\rho^Y_Z := \pi_Z(\partial Y)$ and $\rho^Z_Y := \pi_Y(\partial Z)$ are well-defined.

1.4 The action of the mapping class group

Given any homeomorphism f of our surface and a curve α, $f(\alpha)$ is again a curve, and up to isotopy the result does not depend on the isotopy classes of f and α. Moreover, f preserves disjointness of curves, thus we get an action of the mapping class group on the curve graph by simplicial automorphisms. An important property of this action is the following:

Lemma 1.4.1. *There are finitely many MCG-orbits of simplices inside the curve graph.*

The previous result relies on the following, which is a special case of the so-called *change of coordinates principle* (see [FM12, Section 1.3.1]):

Lemma 1.4.2. *Up to the action of the mapping class group, there are only finitely many representatives of simple closed essential curves, and only one of them is non-separating. More precisely, given two curves α and β, there is a homeomorphism of S mapping α to β if and only if the two surfaces obtained by cutting S along α and β, respectively, are homeomorphic, and up to homeomorphism there are finitely many surfaces of the form $S\backslash\{\gamma\}$ with $\gamma \in \mathcal{C}(S)$.*

We will say that the action $MCG \circlearrowleft \mathcal{C}(S)$ is *cofinite*, since the orbit of a finite set of representatives covers the whole curve graph.

Proof of Lemma 1.4.1. Let $\Delta = \{\gamma_1, \ldots, \gamma_k\}$ be a simplex. Let $R(S)$ be a finite set of representatives of the MCG-orbits of curves in $\mathcal{C}(S)$, which exist by Lemma 1.4.2. Let $\alpha_1 \in R(S)$ be such that there exists $g \in MCG$ such that $g(\gamma_1) = \alpha_1$. Thus, up to replacing Δ with $g\Delta$, we can assume that γ_1 and α_1 coincide. Now let $S_1 = S\backslash\{\alpha_1\}$ be the surface with boundary obtained by cutting along α_1, which is possibly disconnected. Again, by Lemma 1.4.2, there are finitely many topological types of S_1. Now we can repeat the same argument with S_1: there is a finite family $R(S_1)$ of representatives of $MCG(S_1)$-orbits of curves in $\mathcal{C}(S_1)$, and up to the action of $MCG(S_1)$ (which fixes α_1 pointwise because the latter is a boundary curve of S_1) we can assume that $\gamma_2 \in R(S_1)$. Proceeding this way we see that Δ must fall into one of finitely many MCG-orbits of simplices, obtained by successively choosing a curve in $R(S)$, a curve in $R(S_1)$ and so on. □

Corollary 1.4.3. *There are finitely many MCG-orbits of subsurfaces $Y \subset S$.*

Proof. Let Y be a subsurface, and let ∂Y its relative boundary. Complete ∂Y to a pants decomposition Δ for the complementary surface $S\backslash Y$. Now, if $f \in MCG$, then $f(\Delta)$ is a pants decomposition for $f(S\backslash Y) = S\backslash f(Y)$. Therefore simplices in the same MCG-orbit correspond to homeomorphic subsurfaces, and since there are finitely many orbit of simplices the thesis follows. □

Another important fact about the action $MCG \circlearrowleft \mathcal{C}$ is that it induces an action on the set of domains, preserving disjointness and nesting. Moreover, this action preserves subsurface projections:

Lemma 1.4.4. *If $\alpha, \beta \in \mathcal{C}$, Y is a domain and $f \in MCG$, then*

$$d_{f(Y)}(f(\alpha), f(\beta)) = d_Y(\alpha, \beta) \tag{1.1}$$

Proof. Equation (1.1) clearly holds if Y has complexity at least 1, by how subsurface projection is defined. If Y is an annulus with core γ, we first notice that $\alpha \in \mathcal{C}(S)$ crosses γ if and only if $f(\alpha)$ crosses $f(\gamma)$. Moreover, every homeomorphism $f : S \to S$ lifts to a covering map between the annular covers $\tilde{f} : S_\gamma \to S_{f(\gamma)}$, meaning that the following diagram commutes:

$$\begin{tikzcd} S_\gamma \arrow[r, "\tilde{f}"] \arrow[d, "p_\gamma"] & S_{f(\gamma)} \arrow[d, "p_{f(\gamma)}"] \\ S \arrow[r, "f"] & S \end{tikzcd}$$

This is because the image of $\pi_1(S_\gamma)$ via the map induced by $f \circ p_\gamma$ is the subgroup generated by $f(\gamma) \in \pi_1(S)$, which in turn is the image of $\pi_1(S_{f(\gamma)})$ via the map induced by $p_{f(\gamma)}$. Thus such a covering map $\tilde{f}$ exists by basic facts of covering theory. One can show that $\tilde{f}$ is a homeomorphism and extends to a map $\hat{f} : \widehat{S_\gamma} \to \widehat{S_{f(\gamma)}}$ between the compactifications, which again commutes with the covering maps. This means that, if $\alpha \in S$ is a curve crossing γ, then a lift of $f(\alpha)$ to $\widehat{S_{f(\gamma)}}$ is the image under $\hat{f}$ of a lift of α to $\widehat{S_\gamma}$. Now the thesis clearly follows. $\square$

1.4.1 Pseudo-Anosov mapping classes

It is possible to classify mapping classes according to how they act on the curve graph. For example, the Dehn twist along a curve γ fixes the base curve and every curve which does not cross γ, thus it acts as the identity on the star of γ inside $\mathcal{C}$. More generally, Dehn twists are examples of what are called *reducible* mapping classes, which are isotopic to the identity on an essential subsurface (in the case of T_γ, outside an annulus with core γ). On the other hand, there are mapping classes which do not fix any finite collection of curves, which are called *pseudo-Anosov* elements. We will not need the actual definition of what a pseudo-Anosov is; it will suffice to know that these mapping classes exist and have the following properties:

Lemma 1.4.5. *Let S be a surface of complexity at least 1. The following facts hold for a pseudo-Anosov mapping class $f \in MCG(S)$:*

1. *f has infinite order;*
2. *Every power f^k, with $k \in \mathbb{Z}\backslash\{0\}$, is again pseudo-Anosov. Therefore every finite-index subgroup $H \leqslant MCG$ contains pseudo-Anosov elements.*
3. *f is* loxodromic*, meaning that for every curve $\alpha \in \mathcal{C}$ there exists $c \geqslant 0$ (depending on α and f) such that $d_{\mathcal{C}}(f^k(\alpha), \alpha) \geqslant c|k|$.*

Proof. The first two items follow from the Nielsen-Thurston classification Theorem ([Nie44, HT85]), which says that a mapping class is Pseudo-Anosov if and only if it is neither of finite order nor reducible. For a proof of the third fact see e.g. [MM99, Proposition 3.6]. $\square$

We should think of a pseudo-Anosov element as a mapping class that moves curves around the surface in a certain "wild" manner. More precisely, if f is pseudo-Anosov and α, β are any two curves then $\lim_{k\to+\infty} d_{\mathcal{C}}(f^k(\alpha), \beta) = +\infty$ by triangle inequality and Lemma 1.4.5.3, which in particular means that $f^k(\alpha)$ intersects β if k is large enough. We can state a more powerful result, for which we need the following definition.

Definition 1.4.6. A family of curves $\{\gamma_1, \ldots, \gamma_k\} \subset \mathcal{C}(S)$ *fills* the surface if for any curve η there exists $1 \leqslant i \leqslant k$ such that η intersects γ_i.

The following lemma shows that we can always find two filling curves.

Lemma 1.4.7. *Let $f \in MCG$ be a pseudo-Anosov element. For every two curves α, β there exists k_0 such that if $k \geqslant k_0$, $f^k(\alpha)$ and β fill the surface.*

Proof. It suffices to notice that two curves fill the surface if and only if their distance in the curve graph is at least 3, because then any other curve γ must be at distance at least 2 from (hence intersect) one of the two. Then the thesis follows from Lemma 1.4.5.3. $\square$

From Item 3 of Lemma 1.4.5 we also get the following:

Corollary 1.4.8. *The curve graph of a connected surface of complexity at least 1 (hence including the Farey complex) is unbounded.*

1.4.2 Ivanov's Theorem

We conclude this Chapter with a celebrated theorem of Ivanov [Iva97], later extended by Korkmaz [Kor99] to surfaces of low genera. Since in this thesis we will only deal with punctured spheres we state the theorem only in this case.

Theorem 1.4.9 (Ivanov-Korkmaz)**.** *Let $S = S_b$ be a sphere with b punctures, and let $\mathcal{C}(S)$ its curve graph. Let $i : MCG^{\pm}(S) \to \mathrm{Aut}(\mathcal{C}(S))$ be the action of the extended mapping class group.*

- *If $b \geqslant 5$ then i is an isomorphism.*
- *If $b = 4$ then i is surjective. The kernel $\mathcal{K}$ is a Klein four-group generated by the two involutions in Figure 1.5.*

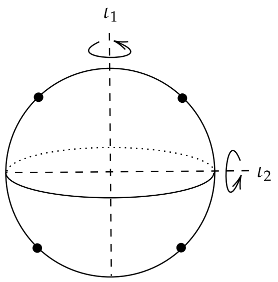

Figure 1.5: The involution ι_1 rotates the sphere by 180 degrees around the vertical axis, while ι_2 is a half rotation around the horizontal axis. These two elements generate a Klein four-group, whose elements are called *hyperrelliptic involutions*.

In other words, every simplicial automorphism of the curve graph comes from a homeomorphism of the surface. The intuition that this proposition conveys is that almost every algebraic property of the mapping class group should be already witnessed by its action on the curve graph. For example, in the same papers the two authors showed the following "algebraic rigidity" result:

Theorem 1.4.10. *Let $S = S_b$ be a sphere with at least five punctures, and let H_1 and H_2 be two subgroups of $MCG^{\pm}(S)$ of finite index. Then any isomorphism $\phi : H_1 \to H_2$ is induced by some inner automorphism of $MCG^{\pm}(S)$. In particular, the outer automorphism group $Out(MCG^{\pm}(S))$ is trivial, while $Out(MCG(S)) \cong \mathbb{Z}/2\mathbb{Z}$.*

The strategy of the proof is to show that an automorphism ϕ between finite index subgroups of the mapping class group induces an automorphism of the curve graph, which by Theorem 1.4.9 comes from an element g of the mapping class group. Then with a little more effort one can show that ϕ is actually the restriction of the conjugation by g. One of the goals of this thesis is to prove equivalents Theorems 1.4.9 and 1.4.10 for a particular quotient of the mapping class group, following a similar strategy (though with significant differences).
Later we will need an extension of Ivanov's theorem to a certain subgraph of the curve graph. We will say that a curve $\gamma \in \mathcal{C}$ is 1-*separating* if it cuts out a subsurface of complexity one, that is, either a once-punctured torus or a sphere with four punctures. We will denote by $\mathcal{C}^1(S)$ the induced subgraph of the curve graph spanned by 1-separating curves. The following result was proven by Bowditch [Bow20, Bow16], building on results by Korkmaz [Kor99] and Luo [Luo00]:

Theorem 1.4.11. *Let $S = S_b$ be a sphere with at least seven punctures. The action of $MCG^{\pm}(S)$ on $\mathcal{C}^1(S)$ gives an isomorphism $MCG^{\pm}(S) \to Aut(\mathcal{C}^1(S))$.*

Chapter 2

Modding by suitable high powers of Dehn twists

For the rest of the thesis, let $S_b = S_{0,b}$ be a sphere with b punctures, on which we fix some orientation, and let $\mathcal{C}_b = \mathcal{C}(S_b)$ be its curve graph, which we will also denote by $\mathcal{C}$ whenever it will not cause ambiguity. Recall that, given two curves α, β (meaning two isotopy classes of simple closed essential curves), we denote their geometric intersection number by $i(\alpha, \beta)$, and if both intersect a third curve s let $d_s(\alpha, \beta)$ be the distance of their annular projections over s. Let $MCG^{\pm}(S_b)$ be the extended mapping class group, where we allow orientation-reversing homeomorphisms.
Here we introduce the main object of this study.

Definition 2.0.1. For any $K \in \mathbb{N}$, let $DT_K = \langle T_\alpha^k | \alpha \in \mathcal{C} \rangle$ be the subgroup generated by all K-th powers of Dehn twists.

Notice that DT_K is a normal subgroup. Indeed, by the properties of a Dehn twist under conjugation (Lemma 1.1.3), if $f \in MCG$ and $\alpha \in C$ then $fT_\alpha^K f^{-1} = T_{f(\alpha)}^K \in DT_K$. Hence we can consider the quotient group MCG/DT_K. Moreover, if we restrict the action $MCG \circlearrowleft \mathcal{C}$ to DT_K we can consider the quotient graph $\mathcal{C}/DT_K$, whose vertices and edges are DT_K-equivalence classes of vertices and edges in $\mathcal{C}$. Notice that, at the moment, we don't know if $\mathcal{C}/DT_K$ is a simplicial graph, since it may have multiple edges. This will be proven in Corollary 2.3.6.
In this Chapter we gather various results about the nature of these two quotients.

2.1 Basic lifting properties

First, notice that the projection map $\pi : \mathcal{C} \to \mathcal{C}/DT_K$ is 1-Lipschitz, since every combinatorial path c in $\mathcal{C}$ is mapped to a combinatorial path $\bar{c} = \pi(c)$ in $\mathcal{C}/DT_K$. On the other hand we have the following:

Lemma 2.1.1. *For every combinatorial path $\bar{c} = \{\bar{\gamma}_1, \ldots, \bar{\gamma}_k\} \subset \mathcal{C}/DT_K$ there exists a path $c = \{\gamma_1, \ldots, \gamma_k\} \subset \mathcal{C}$ such that $\pi(\gamma_i) = \bar{\gamma}_i$ for all i.*

We will say that any such c is a *lift* of $\bar{c}$.

Proof. We proceed by induction on k. If $k = 1$ we have nothing to prove. Now suppose the thesis is proven for $k - 1$. Let $\bar{c}$ be as above. Lift the sub-path $\{\bar{\gamma}_1, \ldots, \bar{\gamma}_{k-1}\}$ to some path

$\{\gamma_1, \ldots, \gamma_{k-1}\}$. Moreover, since an edge in $\mathcal{C}/DT_K$ is an equivalence class of edges in $\mathcal{C}$, there exists an edge $\{\gamma'_{k-1}, \gamma'_k\}$ whose projection is the edge $\{\overline{\gamma}_{k-1}, \overline{\gamma}_k\}$. Since γ_{k-1} and γ'_{k-1} lie in the same DT_K-orbit there exists an element $g \in DT_K$ such that $g(\gamma'_{k-1}) = \gamma_{k-1}$. Thus, since g acts as a simplicial automorphism, we see that γ_{k-1} and $\gamma_k := g(\gamma'_k)$ are adjacent. This way we patched the last edge to the partial lift, and we got a lift of $\bar{c}$. □

Corollary 2.1.2. *Geodesic segments in $\mathcal{C}/DT_K$ lift to geodesic segments in $\mathcal{C}$.*

Proof. Let $\bar{c}$ be a geodesic segment with endpoints $\overline{\alpha}$ and $\overline{\beta}$, and let c one of its lifts. Let α and β be the endpoints of c. Then $d_{\mathcal{C}}(\alpha, \beta)$ is at most the length of c, which is the length of $\bar{c}$. Moreover, since $\bar{c}$ is a geodesic its length is the distance between the endpoints. Thus $d_{\mathcal{C}}(\alpha, \beta) \leqslant d_{\mathcal{C}/DT_K}(\overline{\alpha}, \overline{\beta})$, and since the projection is 1-Lipschitz we must have equality. Hence c already realizes the distance between its endpoints, and therefore is a geodesic. □

2.2 More lifting properties: complexity reduction

The quotient MCG/DT_K was already extensively studied by Dahmani, Hagen and Sisto in [DHS21], and then by Behrstock, Hagen, Martin, and Sisto in [BHMS20]. This section collects a number of results from these papers, which we sometimes enhance. In all our arguments, we will only need one technical property of DT_K (for suitable K), shown in [DHS21] using results of Dahmani [Dah18], which we now state.

Proposition 2.2.1. *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 $\alpha : DT_K \to \mathcal{O}$ such that the following holds. For all $x \in \mathcal{C}$ and $g \in DT_K - \{1\}$ there exist $s \in \mathcal{C}$ and some power $\gamma_s \in DT_K$ 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 [DHS21, Corollary 3.6] (which applies to DT_K in view of [DHS21, Proposition 5.1]). □

The second case will always be used in conjunction with the following, which is a special case of the Bounded geodesic image Theorem, first proven by Masur and Minsky [MM00, Theorem 3.1]:

Theorem 2.2.2 (BGI)**.** *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.*

Roughly speaking, this theorem says that, when moving from x to y along a geodesic, in order to change the projection on the annular curve graph $\mathcal{C}(s)$ one must pass close to s.

Definition 2.2.3. For short, we will say that a normal subgroup $\mathcal{N} \trianglelefteq MCG(S_b)$ is *deep enough* if it satisfies the conclusion of Proposition 2.2.1 for some Θ depending on the data that has been fixed up to that point.

Remark 2.2.4. Notice that a deep enough subgroup $\mathcal{N}$ is generated by powers of Dehn twist. 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. This $g = \gamma_s^{-1} g'$ is a product of Dehn twists.

The key use of the conclusion of Proposition 2.2.1 will be to lift a variety of subgraphs from a quotient of the curve graph to the curve graph. This is also done in [DHS21] as well as generalized in [BHMS20]. First, a definition from the latter paper.

Definition 2.2.5. A *generalized* m-gon in a graph is a sequence $\tau_0, \ldots, \tau_{m-1}$ 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.

Roughly speaking, a generalized m-gon is a "necklace" whose "beads" are simplices or geodesic segments, glued together along sub-simplices or endpoints. Notice that the necklace could self-intersect, that is, we do not require that two non-adjacent beads do not intersect, though this will almost always be the case in this thesis.
The following is a small variation on a result from [BHMS20].

Theorem 2.2.6. *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 generalized m-gon in $\mathcal{C}/\mathcal{N}$ can be lifted in $\mathcal{C}$.*

Proof. This is essentially [BHMS20, Proposition 8.29], whose proof only uses [DHS21, Corollary 3.6] and therefore works for any deep enough subgroup. Said proposition is stated for a precise threshold m_0, but the proof works for any fixed m_0. Finally, [BHMS20, Proposition 8.29] 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 2.2.7 and 2.2.8, whose proofs are prototypical of many arguments throughout the thesis. $\square$

Lemma 2.2.7. *Whenever $\mathcal{N}$ is deep enough, every edge $\overline{e} \subset \mathcal{C}_4/\mathcal{N}$ admits a unique $\mathcal{N}$-orbit of lifts.*

Proof. Let $\overline{x}, \overline{y}$ be the vertices of $\overline{e}$, and let $e = \{x, y\}$ and $e' = \{x', y'\}$ be two lifts of $\overline{e}$. 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 2.2.1. If $d(x, s) \leqslant 1$, 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 the projection $\pi_s(y)$ 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 2.2.2, 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 $x = x'$, and since $\mathcal{C}$ is a simplicial graph we must also have that $e = e'$. Therefore $\overline{e}$ has a unique orbit of lifts. $\square$

Lemma 2.2.8. *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 Lemma 2.1.1, 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 2.2.1. 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 these 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 2.2.2, 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 2.2.7 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 2.2.1. 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 2.2.7 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 2.2.6, which shows that, when lifting a quadrilateral (that is, a geodesic quadrangle), two opposite sides can be chosen somewhat "rigidly":

Lemma 2.2.9. *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 [DHS21, Proposition 4.3], whose proof only uses [DHS21, Corollary 3.6]. □

2.3 Projection lemmas

While the previous sections were about lifting subgraphs, the present aims to study how subgraphs of $\mathcal{C}$ get projected to the quotient.

2.3.1 Puncture separations

Firstly, we make some observations on how two curves which separate the punctures on the sphere "in the same way" must have different projections. Choose an enumeration of the punctures of S_b. A curve α on S_b is always separating, therefore it splits the punctures into two sets $\mathcal{B}^+(\alpha)$ and $\mathcal{B}^-(\alpha)$.

Definition 2.3.1. The *puncture separation* induced by α is the unordered pair $\{\mathcal{B}^+(\alpha), \mathcal{B}^-(\alpha)\}$.

Remark 2.3.2. 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, 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 2.2.4.

Definition 2.3.3. 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 2.3.2). The following lemma shows a peculiar behavior of the projection $\mathcal{C} \to \mathcal{C}/DT_K$ which holds for all $K \in \mathbb{N}$:

Lemma 2.3.4. *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 $b = 4$, 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.

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

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 g is a product of Dehn twists, which preserve puncture separations by Remark 2.3.2.

□

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

Corollary 2.3.5. *For every $b \geqslant 4$, the projection map is an isometry when restricted to any simplex $\Delta \subset \mathcal{C}$.*

Moreover, we can finally show that the quotient graph is simplicial:

Corollary 2.3.6. *If $\mathcal{N}$ is deep enough the graph $\mathcal{C}/\mathcal{N}$ is simplicial. In particular, two vertices $\overline{\alpha}, \overline{\beta} \in \mathcal{C}/\mathcal{N}$ are adjacent if and only if they have disjoint lifts $\alpha, \beta \in \mathcal{C}$.*

Proof. First we check that $\mathcal{C}/\mathcal{N}$ has no double edges. Suppose by contradiction that there exist two edges $\overline{e}, \overline{e}'$ between the same two vertices $\overline{\alpha}, \overline{\beta} \in \mathcal{C}/\mathcal{N}$. By Lemma 2.1.1 (whose proof works also for any deep enough subgroup $\mathcal{N}$) there exist two consecutive edges e, e' that lift $\overline{e}, \overline{e}'$. Let $e = \{\alpha, \beta\}$ and $e' = \{\beta, \alpha'\}$. Since α and α' both lift $\overline{\alpha}$, with the same proof of Lemma 2.2.7 we can find an element $g \in \mathcal{N}$ mapping α to α' and fixing β. Thus $g(e) = e'$, and therefore $\overline{e} = \overline{e}'$, against our hypothesis.
We are left to prove that no edge $\overline{e}$ of $\mathcal{C}/\mathcal{N}$ joins a vertex α to itself. If this was the case, we could lift $\overline{e}$ to an edge e of $\mathcal{C}$ joining two curves α, α' with the same projection, thus contradicting Item 1 or 2 of Lemma 2.3.4. ☐

2.3.2 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 [DHS21, 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:

Lemma 2.3.7. *Fix $\Theta > 0$. Let $\mathcal{N}$ be deep enough and let $\overline{\Delta} \subseteq \mathcal{C}/\mathcal{N}$ be a (possibly empty) simplex and $\Delta \subseteq X$ be one of its lift. Suppose that $x, y \in Lk(\Delta)$ project to $\overline{x}, \overline{y} \in Lk(\overline{\Delta})$ and have the property that $d_s(x, y) < \Theta$ whenever 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 [DHS21, 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 2.2.6. 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 2.2.1. 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 2.2.7 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 2.3.8. *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 2.3.7 applies. One way to see this is to complete x, z to complete clean markings μ, ν, in the sense of [MM00, Section 2.5]. By the Distance formula [MM00, 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 realize $d_s(\mu, \nu)$. 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 by Lemma 1.3.1 annular projections are 2-Lipschitz we have that

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

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

There are also two other consequences of Lemma 2.3.7 that will be quite useful later. The first one deals with the cardinality of the quotient:

Corollary 2.3.9. *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 [DHS21, Theorem 2.1] which just uses the Bounded geodesic image Theorem 2.2.2 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 2.3.8. Hence, if $\mathcal{N}$ is deep enough with respect to Θ, by Lemma 2.3.7 the projection is an isometry on the orbit $\{g^n(x)\}_{n \in \mathbb{Z}}$, which is infinite since g is pseudo-Anosov. $\square$

The second consequence deals with filling curves, as in Definition 1.4.6:

Corollary 2.3.10. *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 $d_{\mathrm{Lk}(\Delta)}(x, y) \geqslant 3$. Now if x and y satisfy the hypothesis of Lemma 2.3.7 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. $\square$

2.3.3 Quotient curve graphs of subsurfaces

We end this section by showing that the restriction of π to the curve graph of a subsurface is what one would expect.

Lemma 2.3.11. *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 2.2.1.
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 (Lemma 1.1.3).
Otherwise $d_s(x, x') > \Theta$. Arguing exactly as in the proof of Lemma 2.2.7 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 2.3.12. *Let U be a connected subsurface of S of complexity at least 2, and let $\mathcal{N}$ be a deep enough subgroup. If π is the quotient projection then $\pi(\mathcal{C}(U)) \cong \mathcal{C}(U)/\mathcal{N}(U)$.*

Proof. By Lemma 2.3.11 these graphs have the same vertices, and since U has complexity at least 2 the adjacency relation is having disjoint lifts in both cases. □

Chapter 3

The finite rigid set

By results of Aramayona and Leiniger [AL13] there exists a finite subgraph $X_b \subset \mathcal{C}_b$ such that every two copies of X_b inside the curve graph are obtained one from the other via a mapping class. In this Chapter we will show that, whenever $\mathcal{N}$ is a deep enough subgroup, the same subgraph can be found inside the quotient $\mathcal{C}/\mathcal{N}$, and that conversely any copy of X_b inside $\mathcal{C}/\mathcal{N}$ admits a unique orbit of lifts. This will allow us to extend the result of Aramayona and Leiniger to $\mathcal{C}/\mathcal{N}$ in Chapter 4, and this will be the key piece in the proof of the Combinatorial Rigidity Theorem 4.

3.1 Special intersections

We first make a digression and talk about when it is possible to recognize that two curves on our sphere have intersection number 2, just by looking at the curve graph. We will define what we call "special intersection" between curves, which should be compared with the notion of $\mathcal{X}$-detectable intersection from [AL13]. For the following definitions, let Γ be either $\mathcal{C}$ or $\mathcal{C}/\mathcal{N}$.

Definition 3.1.1. A *chain* is a collection of vertices $v_1, \ldots, v_k \in \Gamma$ 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.

Definition 3.1.2 (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;
- both γ and δ intersect exactly one vertex $\varepsilon \in P$.

We say that γ and δ are the auxiliary vertices that detect the intersection of α and β.

Lemma 3.1.3. *Let α and β be two curves in $\mathcal{C}$ with special intersection. Then $i(\alpha, \beta) = 2$.*

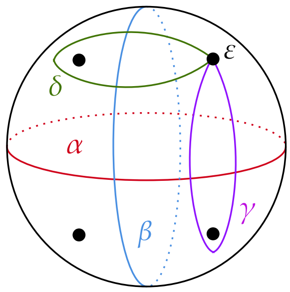

Figure 3.1: The various curves involved in the definition of a special intersection, on the four-holed sphere that the facet P cuts out.

Proof. Let P be the facet that detects the special intersection. Since P has dimension one less that the maximal, it cuts out a subsurface of complexity 1, which can only be a four-punctured sphere S_4 since S has no genus. Moreover, α and β must lie on this S_4, because the union $\alpha \cup \beta$ is connected and disjoint from every curve of P.
Now, we claim that the intersection number between α and β must be 2. This can be shown by successively drawing the curves on S_4, so that the final situation will be as in Figure 3.1. First draw α on the sphere, which must surround two punctures on each side since it is not homotopic to any of the punctures (because these correspond to some curves in P). Now, δ intersects S_4 because it crosses $\beta \subset S_4$; moreover the only curve in P that δ crosses is ε, thus ε must correspond to one of the punctures and the trace of δ on S_4 must be a collection of arcs with endpoints on the puncture corresponding to ε. However, δ is disjoint from α, hence every arc of $\delta \cap S_4$ is a non-homotopically-trivial arc lying on one of the two twice-punctured disks that α cuts out. Since, up to isotopy of a disk, there is only one such arc, δ must be as in Figure 3.1.
Now, γ crosses α but not β, then again the trace of γ on the pair of pants obtained by cutting S_4 along δ is an arc with endpoints on the puncture corresponding to $\varepsilon \cup \delta$, and therefore γ must be as in Figure 3.1. Finally, β does not cross γ and is not homotopic to any puncture, thus with the same arguments it must be as in Figure 3.1. Now we can finally see that α and β have intersection number 2. $\square$

Remark 3.1.4. Definition 3.1.2 actually describes an isometrically embedded generalized pentagon $\mathfrak{P}$, as in Figure 3.2. 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 which 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 3.1.5. We will call an isometrically embedded generalized pentagon $\mathfrak{P} \subset \Gamma$ which is isomorphic to the one in Figure 3.2 a *special* pentagon.

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

Lemma 3.1.6. *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.*

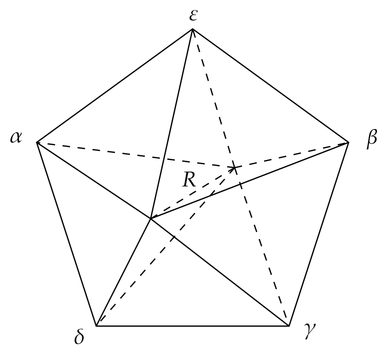

Figure 3.2: The generalized pentagon $\mathfrak{P}$ described in 3.1.2, which detects that any two non-adjacent vertices in the link of R have special intersection.

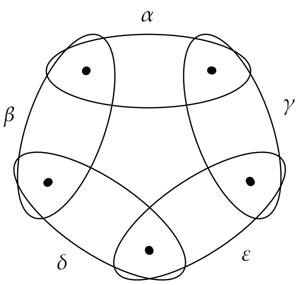

Figure 3.3: The five curves involved in Definition 3.1.2 for $\Gamma = \mathcal{C}$, forming a chain on the five-holed sphere that R cuts out. Every intersection is special, therefore the intersection number is always 2.

Proof. By Theorem 2.2.6, whenever $\mathcal{N}$ is deep enough every special pentagon $\overline{\mathfrak{P}}$ as in Definition 3.1.5 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 thesis, given a curve $\beta \in \mathcal{C}$ that surrounds 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 3.1.7. *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 of Lemma 3.1.7. P cuts out a sphere S_4 with four punctures, and by Lemma 3.1.3 having special intersection implies that the intersection number is 2. Thus we want to show that, if α, α' are adjacent to β in the Farey complex, then $\alpha' = H_\beta^k$ for some $k \in \mathbb{Z}$. By properties of the Farey complex (see Lemma 1.2.10) there exists a sequence of triangles $T_1, \ldots, T_k$ 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 vertices of a triangle then $\alpha' = H_\beta^{\pm 1}(\alpha)$. This is true, since the two curves $H_\beta^{\pm 1}(\alpha)$ are adjacent to both α and β, and α and β belong to exactly two triangles (again, by Lemma 1.2.10). □

Remark 3.1.8. 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 3.1.7 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.*

3.2 Definition of 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 [AL13, 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 3.4 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 by Aramayona and Leiniger [AL13, Theorem 3.1]. Recall that a simplicial map $\phi : G \to G'$ between simplicial graphs is *locally injective* if its restriction to the star of every vertex $v \in G$ is injective (in particular, isometric embeddings are locally injective).

Theorem 3.2.1. *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*

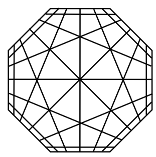

Figure 3.4: Doubling these arcs gives a copy of X_8.

of the pointwise stabilizer 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 [AL13, Lemma 3.2]). For example, the special intersection between α and β in Figure 3.5 is detected by the special pentagon spanned by α, β, the simplex $R \subset X_b$ of codimension two and the three curves $x, y, z \in X_b$.

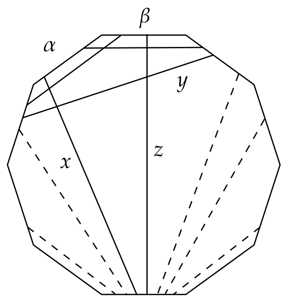

Figure 3.5: The dashed lines represent the codimension two simplex R. Any two intersecting curves of this "star" have special intersection, with respect to some facet that extends R.

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.

Definition 3.2.2. If $b \geqslant 6$, we say that a vertex β in a copy of X_b is *minimal* if $\mathrm{Lk}_{X_b}(\beta) \cong X_{b-1}$.

Lemma 3.2.3 (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* 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* 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 $S_b \backslash \beta$, 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, since its complexity cannot be too low. Thus the other connected component of $S_b \backslash \beta$ 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 neighborhood 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 neighborhood 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 intersection with β with respect to the same P). By Lemma 3.1.7 and the discussion in Remark 3.1.8 there is a suitable power H^k_β 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. $\square$

Lemma 3.2.4 (Uniqueness of the extension)**.** *Let $b \geqslant 6$. In any copy of 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 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 3.5. 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 3.6. 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. $\square$

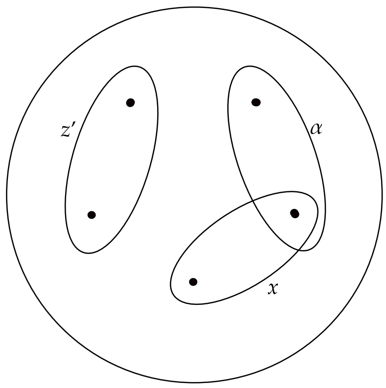

Figure 3.6: The five-punctured sphere that R cuts out.

Remark 3.2.5. When $b = 5$ we cannot define minimal curves as in Definition 3.2.2. Thus, with a slight abuse of notation, we will say that every curve in X_5 is minimal, because it bounds a twice-punctured disk. Moreover Lemma 3.2.3 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 3.2.4, with $R = \varnothing$, to get the following:

Lemma 3.2.6. *For $b = 5$, if two copies $X_5, X_5' \subset \mathcal{C}_5$ share four curves then they must coincide.*

3.3 Projecting the rigid set

Here we aim to show that, whenever $\mathcal{N}$ is deep enough, there is at least a copy of the finite rigid set inside $\mathcal{C}/\mathcal{N}$. This will follow from Corollary 2.3.8, which dealt with isometric projections.

Definition 3.3.1. Given a copy of $X_b \subset \mathcal{C}$ for $b \geqslant 5$, set

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

Lemma 3.3.2. *For every $b \geqslant 5$ there exists $M \in \mathbb{R}^+$ with the following property. Given a copy $X_b \subset \mathcal{C}$ set Y_b as in Definition 3.3.1. 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 2.3.8. 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 3.2.1. In particular f maps minimal curves to minimal curves. Therefore

$$Y_b' = X_b' \bigcup_{\beta' \in X_b'\, minimal} H_{\beta'}^{\pm 1}(X_b') = f(X_b) \bigcup_{\beta \in X_b\, minimal} H_{f(\beta)}^{\pm 1} \circ f(X_b) =$$

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

Now the thesis 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 3.3.3. *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 Y_b be as in Definition 3.3.1. Then the restriction $\pi|_{Y_b}$ is an isometry. In particular there exists an isometrically embedded copy of X_b inside $\mathcal{C}/\mathcal{N}$.*

Proof. Just combine Lemma 3.3.2 and Corollary 2.3.8. □

3.4 Lifting the rigid set

In the previous section we showed that the finite rigid set can be isometrically projected to the quotient. Conversely, our current goal is to prove the following lifting result:

Theorem 3.4.1. *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 $\mathcal{N}$.*

The key point of both the existence and uniqueness of lifts will be the following lemma:

Lemma 3.4.2 (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 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 3.1.7 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 2.3.2. Moreover x and y intersect twice, thus they separate the punctures of S into four 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 2.2.1. 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

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

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 3.3.2, which does not depend on Θ. This is a contradiction if $\mathcal{N}$ is deep enough. ☐

3.4.1 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 2.2.6. 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 3.4.3. *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 2.2.6, 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 $S \backslash \beta$, or in other words that, given any two curves $\gamma, \delta \in \mathrm{Lk}_{\mathcal{C}_b}(\beta)$, we can find a chain $\gamma = \eta_0, \eta_1, \ldots, \eta_k = \delta$ inside $\mathrm{Lk}_{\mathcal{C}_b}(\beta)$. Since every curve in $\mathrm{Lk}_{\mathcal{C}_b}(\beta)$ 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 2.3.12 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 $\mathrm{Lk}_{\mathcal{C}_b}(\beta)$.
In order to complete our lift of $\overline{X}_b$ we still need to lift the vertices $\overline{\alpha}_1, \ldots, \overline{\alpha}_{b-3} \in \overline{X}_b$ outside the link of $\overline{\beta}$. For every $i = 1, \ldots, b-3$ 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 3.5, 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 3.4.2 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 3.2.3 provides a copy X_b^i which extends $\beta \star X_{b-1}$ and contains α_1^i. By Lemma 3.2.4, 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 α_1^2 and α_1^j project to the same element $\overline{\alpha}_1$. Therefore, if we apply Lemma 3.4.2 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 α_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 3.2.4, 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. □

3.4.2 Uniqueness of the lift

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

Lemma 3.4.4. *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 3.4.2. More precisely, let $\{\overline{\gamma}_i\}_{i=1,\ldots,5}$, 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 2.2.6. Now, if we apply Lemma 3.4.2 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 3.2.6. □

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 3.4.5. *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. $\square$

Now we are finally able to prove the theorem. 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 3.4.2 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 3.2.4 we see that the two lifts now coincide. $\square$

Chapter 4

Combinatorial rigidity

The aim of this Chapter is to prove the Combinatorial Rigidity Theorem 4, which is implied by Theorem 4.4.1. First we show an analogue for the special case $b = 4$. The proof will highlight the core of the general case, though the latter will require some more machinery.

4.1 The four-punctured case

The subgraph playing the role of the rigid set in this case will be any triangle T inside the Farey complex, since the following well-known result holds (for example, it follows from the discussions in [AL13, Section 3] and in [FM12, Section 3.4]):

Theorem 4.1.1. *Given two triangles $T, T' \subset \mathcal{C}_4$ there exists an element $h \in MCG^{\pm}(S_{0,4})$ mapping T to T. Any two such h differ by an element of the Klein four-group $\mathcal{K}$ of hyperelliptic involutions (see Figure 1.5).*

Now we want to prove the following:

Theorem 4.1.2. *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 1.5.

The proof will actually work for the quotient by any deep enough subgroup $\mathcal{N}$. We split the proof in two propositions below.

Proposition 4.1.3. *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 4.1.4. *The following facts hold:*

1. *Any two adjacent curves α, $\beta \in \mathcal{C}_4$ 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 properties of the Farey complex (see Lemma 1.2.10). 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 $\{\alpha', \beta', \gamma'\}$, and by uniqueness of the orbit of lifts of an edge (which is a particular case of Theorem 2.2.6) 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 4.1.3. 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 2.3.4. Let $\widetilde{T}$ be a lift of $\overline{\phi}(\overline{T})$, which exists by Theorem 2.2.6. By Theorem 4.1.1 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 $T \subset \mathcal{C}_4$:

$$\begin{tikzcd}
\mathcal{C}_4 \arrow[r, "h"] \arrow[d, "\pi"] & \mathcal{C}_4 \arrow[d, "\pi"] \\
\mathcal{C}_4/\mathcal{N} \arrow[r, "\overline{\phi}"] & \mathcal{C}_4/\mathcal{N}
\end{tikzcd} \tag{4.1}$$

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 for any two triangles in the Farey complex there exists a sequence of triangles connecting them, such that two consecutive triangles share an edge (see Lemma 1.2.10). 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 (4.1) commutes also on T', as required. ☐

Proposition 4.1.5. *For $b = 4$ 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 2.2.8 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 4.1.1. ☐

While the injectivity part of the case $b \geqslant 5$ will be very similar to the proof of Proposition 4.1.5, for the surjectivity part we will need a more general version of Lemma 4.1.4. The following sections are devoted to establish such a result, the main tool being half twists around minimal vertices.

4.2 Characterization 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 characterization, which therefore is preserved by any automorphism of the curve complex. More precisely we want to prove the following:

Lemma 4.2.1. *For any $\alpha \in X_b$:*

- $H_\beta(\alpha) = \alpha$ *iff* $d_{\mathcal{C}}(\alpha, \beta) \leqslant 1$;
- *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 characterize $H_\beta(\alpha)$, since it is impossible to distinguish it from $H_\beta^{-1}(\alpha)$ just from the curve graph (since 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. □

4.3 Definition of half twists in the quotient

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

Definition 4.3.1. Let $\overline{X}_b$ be a copy of 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 $H_{\overline{\beta}} := \pi \circ H_{s(\overline{\beta})} \circ s$.

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 of $\overline{X}_b$ 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 3.3.3, 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 characterization of the half twist described in Theorem 4.2.1 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 4.2.1;
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.

4.3.1 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 4.3.2. $H_{\overline{\beta}}(\overline{\alpha}) = \overline{\alpha}$ *iff* $d_{\mathcal{C}/DT_K}(\overline{\alpha}, \overline{\beta}) \leqslant 1$.

Proof. By construction $H_{\overline{\beta}}$ is the identity on the link of $\overline{\beta}$. Conversely, if $\overline{\alpha}$ is at distance 2 from $\overline{\beta}$ then α and $H_\beta(\alpha)$ are different curves in Y_b, and by Theorem 3.3.3 they project to different points. □

Lemma 4.3.3. *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 3.5. 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 4.1 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

4.2. Since G it is a subgraph of the graph Y_b, defined in Definition 3.3.1, by Theorem 3.3.3 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 $\overline{\beta}$. □

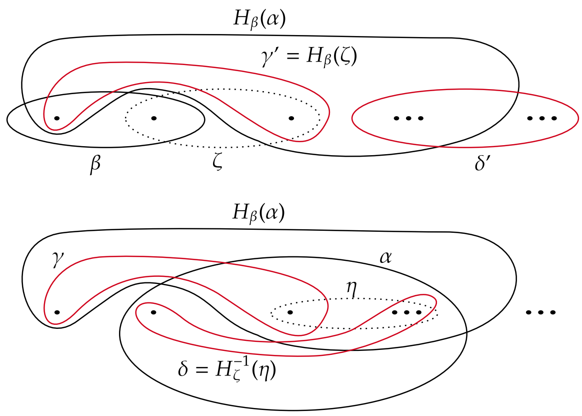

Figure 4.1: An example of how the auxiliary curves (here in red) may be chosen only using half twists around β and ζ, which are both in X_b. In our example $\gamma = \gamma'$, but this is not necessary. The case where α bounds a twice-punctured disk can be dealt with similarly.

4.3.2 If it looks like a twist it is the projection of a twist

Lemma 4.3.4. *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 4.2. 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 3.4.2 the union of the two simplices $\overline{\alpha} \star \overline{P} \cup \overline{\beta} \star \overline{P}$ admits a unique 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. □

4.4 The general case

We are finally ready to prove the following combinatorial rigidity statement, which implies Theorem 4 when $\mathcal{N} = DT_K$ for suitable K.

Theorem 4.4.1. *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.

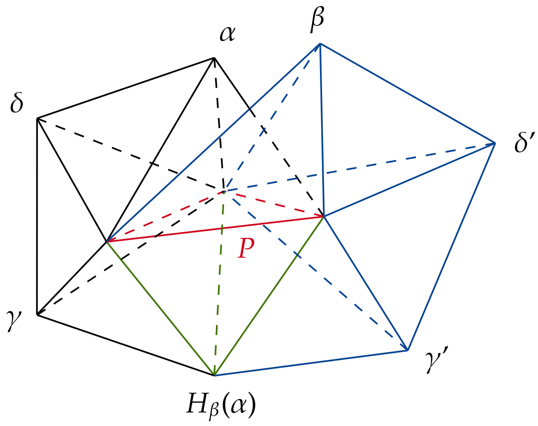

Figure 4.2: The configuration G that detects the special intersection of $H_\beta(\alpha)$ with both α and β, that consists in two special pentagons (here, in black and blue) which overlap over a simplex (here, in green and red). Possibly $\gamma = \gamma'$ and $\delta = \delta'$, but each of the pentagons is isometrically embedded.

Theorem 4.4.2. *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 3.3.3. Moreover let $\widetilde{X_b}$ be a lift of $\overline{\phi}(\overline{X}_b)$. By Theorem 3.2.1 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 $X_b \in \mathcal{C}$:

$$\begin{tikzcd} \mathcal{C} \arrow[r, "h"] \arrow[d, "\pi"] & \mathcal{C} \arrow[d, "\pi"] \\ \mathcal{C}/\mathcal{N} \arrow[r, "\overline{\phi}"] & \mathcal{C}/\mathcal{N} \end{tikzcd} \tag{4.2}$$

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 characterization of the half twist, Theorem 4.2.1. Section 4.3 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 4.2 commutes also on $H_\beta(\alpha_1)$, and hence on the whole $H_\beta(X_b)$, which by Lemma 3.2.4 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 $H^{-1}_\zeta(\beta) = H_\beta(\zeta)$, as shown in Figure 4.3, therefore we already know that the diagram commutes on $H^{-1}_\zeta(\beta)$ (and therefore on $H_\zeta(\beta)$ which is the only other curve that satisfies the same characterization as $H^{-1}_\zeta(\beta)$). 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 4.2 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:

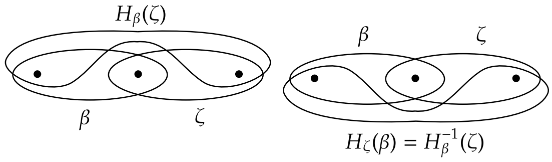

Figure 4.3: The two minimal curves from the proof of theorem 4.4.2 and their respective half twists.

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

Proof of Claim. To show this let h be a (orientation-preserving) mapping class that maps X_b to X_b'. Now, if $\beta_1, \ldots, \beta_b$ are the minimal curves in X_b we know that $H_{\beta_1}, \ldots, H_{\beta_b}$ generate the mapping class group, by Theorem 1.1.6. Therefore let $h = H_{\beta_{i_1}}^\pm \ldots H_{\beta_{i_k}}^\pm$, for some possibly repeated indices $i_1, \ldots, i_k \in \{1, \ldots, b\}$. Now we proceed by induction on k. If $k = 1$ we are done. Otherwise notice that, by the properties of half Dehn twists (Lemma 1.1.4),

$$h = H_{\beta_{i_1}}^\pm \ldots H_{\beta_{i_k}}^\pm = H_{H_{\beta_{i_1}}^\pm(\beta_{i_2})}^\pm \ldots H_{H_{\beta_{i_1}}^\pm(\beta_{i_k})}^\pm H_{\beta_{i_1}}^\pm.$$

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

This concludes the proof of Theorem 4.4.2. □

Now we turn to the injectivity part:

Theorem 4.4.3. *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 $g \in \mathcal{N}$.*

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 3.4.1 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 3.2.1 we must have that $h^{-1} \circ g \in \langle r \rangle$. Now if we show that r does not induce the identity on $\mathcal{C}/\mathcal{N}$ we must have that $g = h$, and we are done. 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, Lemma 1.1.4). 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 3.2.1 there is only one only 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 2.3.9). □

4.5 Finite rigidity of the quotient

As a byproduct of the proofs of Theorems 4.4.1 and 4.1.2 we get an analogue of Theorem 3.2.1 for the quotient.

Corollary 4.5.1 (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 stabilizer $PStab(\overline{X}_b)$, 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 $\mathrm{Pstab}(X_b)$. ☐

Chapter 5

Extracting combinatorial data from a quasi-isometry

The main goal of this Chapter 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 [BHS21, 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 by results of [BHMS20]. More precisely, [BHS21, Theorem 5.7] states that quasi-isometries of a HHS satisfying three additional assumptions induce automorphisms of a suitable graph. Those assumptions are not satisfied by $MCG(S_b)/DT_K$, but nevertheless we will show that a very similar statement applies, see Theorem 5.2.15 which is the main result of this Chapter. The Theorem will also apply to pants graphs (see Chapter 6).
In Section 5.1 we will recall the basic notions of coarse geometry, including what a quasi-isometry is, and of geometric group theory. Then in Section 5.2 we will give an overview on hierarchically hyperbolic spaces and gather all facts about HHSs that we will need.

5.1 A crash course in coarse geometry

For this section we will mainly follow the book of Bridson and Haefliger [BH99].

5.1.1 Coarse maps

Definition 5.1.1. Let $(X, d_X), (Y, d_Y)$ be two metric spaces. A map $f : X \to Y$ is said to be:

- *coarsely Lipschitz* if there exist two constants $k, c \geqslant 0$ such that for every $x, x' \in X$ we have that
$$d_Y(f(x), f(x')) \leqslant k d_X(x, x') + c;$$
- *coarsely surjective* if there exists a constant $c \geqslant 0$ such that, for every $y \in Y$ there exists $x \in X$ such that $d_Y(f(x), y) \leqslant c$;
- a *quasi-isometric embedding* if there exists two constants $k, c \geqslant 0$ such that for every $x, x' \in X$ we have that
$$\frac{1}{k} d_X(x, x') - c \leqslant d_Y(f(x), f(x')) \leqslant k d_X(x, x') + c;$$

- a *quasi-isometry* if it is a coarsely surjective quasi-isometric embedding.

We say that two maps $f, g : X \to Y$ *coarsely coincide*, or are at bounded distance, if there exists a constant c such that, for every $x \in X$, $d_Y(f(x), g(x)) \leqslant c$.

Notice that coarsely Lipschitz maps and quasi-isometric embeddings need not be continuous, since we get no information at all at scales below c. In other words, this kind of maps only capture the "large scale" behavior of distances in the two spaces: for example, a good way of thinking about quasi-isometric embeddings is that they are bi-Lipschitz maps at a large scale.
Two metric spaces are *quasi-isometric* if there exists a quasi-isometry between them (that is, if they look the same "from far enough"). Being quasi-isometric is an equivalence relation, since it is easy to show that a quasi-isometric embedding $f : X \to Y$ is a quasi-isometry if and only if admits a *quasi-inverse* $g : Y \to X$, that is, a quasi-isometric embedding such that $g \circ f$ coarsely coincide with the identity map id_X on X, and similarly $f \circ g$ and id_Y are at bounded distance.
Notice moreover that the composition of two quasi-isometries is a quasi-isometry (with possibly larger constants). Summing up these facts we see that, given a metric space X, the set $QI(X)$ of self-quasi-isometries of X up to bounded distance forms a group.

Definition 5.1.2. Let X be a metric space. A *(quasi)geodesic line* is a (quasi-)isometric embedding $\gamma : \mathbb{R} \to X$. One can similarly define (quasi)geodesic rays and segments.

We will often conflate a (quasi-)geodesic with its image. Notice that, if X is a metric simplicial graph, this new notion of geodesic agrees with the one we gave before, which is Definition 1.2.3.

Definition 5.1.3. A metric space X is said to be:

- *geodesic* if for every two points $x, x' \in X$ there exists a geodesic segment γ whose endpoints are x and x';
- *quasi-geodesic* if there exist two constants k, c such that for every two points $x, x' \in X$ there exists a (k, c)-quasigeodesic segment γ whose endpoints are x and x'.

We will denote a geodesic segment with endpoints x and x' by $[x, x']$, though there might be more than one geodesic segment between two points.
Though in what follows we could work with quasigeodesic metric spaces, we prefer to consider geodesic metric spaces, which are nicer since they are path connected. This is why it is sometimes better to replace an abstract graph (which is just a collection of points with the discrete topology) with its geometric realization. In the latter case the new notion of geodesic agrees with the one given in Definition 1.2.3.

5.1.2 The Cayley graph of a group

Here we recall how to endow a finitely generated group G with a metric graph structure. Let S be a finite set of generators, which we assume to be symmetric (meaning that if $s \in S$ then its inverse s^{-1} belongs to S as well).

Definition 5.1.4. The *Cayley graph* $\mathrm{Cay}(G, S)$ is the simplicial graph whose vertices are the elements of the group, and two elements $g, h \in G$ are adjacent if and only if there exists a generator $s \in S$ such that $g = hs$.

If we denote by $d_{G,S}$ the length metric as defined in Subsection 1.2.1, which we call the *word metric* associated to the generating set S, we see that for every $g, h \in S$ their distance is

$$d_{G,S}(g, h) = \min\{k \in \mathbb{N} | \exists s_1, \ldots, s_k \in S \text{ s.t. } g = hs_1 \ldots s_k\}.$$

This distance is always finite, since $h^{-1}g$ is a finite product of generators. This shows that $\mathrm{Cay}(G, S)$ is a connected graph, and therefore a geodesic metric space.
The word metric is invariant under left multiplication by elements of G, or *left-invariant* for short, since if $g = hs$ then for every $f \in G$ we have that $fg = (fh)s$. In other words, the group acts on its Cayley graph by left multiplication, and every element acts as a simplicial automorphism.
For every element $g \in G$ we can define its *norm* as $\|g\|_{G,S} = d_{G,S}(1, g)$, where $1 \in G$ is the trivial element of the group. Notice that, for every $g, h \in G$, we have an estimate on the norm of their product:

$$\|gh\| = d(gh, 1) \leqslant d(gh, g) + d(g, 1) = d(h, 1) + d(g, 1) = \|h\| + \|g\|.$$

Now, as a graph $\mathrm{Cay}(G, S)$ depends both on the group and on the chosen generating set. However:

Lemma 5.1.5. *Let G be a finitely generated group. If S, S' are two symmetric generating set, then the identity map $id_G : G \to G$ extends to a quasi-isometry $id_G : Cay(G, S) \to Cay(G, S')$. Hence "the" Cayley graph $Cay(G)$ is well-defined up to quasi-isometry.*

Proof. Since S is a generating set and S' is finite, there exists a constant $K = \max_{s' \in S'} \|s'\|_{G,S}$. Now, every element $g \in G$ can be written as $g = s'_1 \ldots s'_l$ for some $s'_i \in S'$. Suppose that l is minimal with this property, so that $l = \|g\|_{G,S'}$. Then

$$d(g, 1)_{G,S'} = \|g\|_{G,S'} \leqslant \|s'_1\| + \ldots + \|s'_l\| \leqslant Kl = K\|g\|_{G,S'} = Kd(g, 1)_{G,S}.$$

Hence, using that both metrics are left-invariant, we see that the identity map is K-Lipschitz. Swapping S and S' in the previous argument we see that the identity map is a quasi-isometric embedding, and therefore a quasi-isometry since it is clearly surjective. □

Remark 5.1.6. An important property of the Cayley graph is that it is a *proper* metric space, meaning that its closed balls are compact. In fact, it is easy to see that a metric graph is proper if and only if every link of a vertex is finite, which happens in our case (all links are isomorphic via the action of G, which is transitive, and the link of the trivial element is the finite set S).

Defining the Cayley graph is the first step towards Geometric Group Theory, which is the study of algebraic properties of groups with geometric tools. The second step is establishing when an action of a group G on a metric space X induces a quasi-isometry between $\mathrm{Cay}(G)$ and X, so that the large scale geometry of the group can be studied by looking at its "nice" actions. We first need to define what "nice" means:

Definition 5.1.7. Let (X, d) be a metric space and G be a group acting on X by isometries. The action is said to be:

- *proper* if for any $x \in X$ and any ball $B \subseteq X$ there are only finitely many elements of G that map x inside B;
- *cobounded* if there exists a ball B of finite radius whose G-translates cover the whole X, that is to say $G \cdot B = X$;

We say that G acts *geometrically* on a metric space X if the action is proper, cobounded and by isometries.

It is easy to see that, if G is finitely generated, the action $G \circlearrowleft \mathrm{Cay}(G)$ is geometric. The following, which is sometimes called the fundamental lemma of Geometric Group Theory, show that geometric actions are the "nice" ones we were looking for:

Lemma 5.1.8 (Milnor-Švarc, [Š55, Mil68])**.** *Let G be any group, acting geometrically on a geodesic metric space X. Then*

- *G is finitely generated;*
- *$Cay(G)$ is quasi-isometric to X, via the map $g \to gx_0$ for any given choice of $x_0 \in X$. This quasi-isometry is G-equivariant, meaning that it commutes with the G-actions.*

Proof. See e.g [BH99, Proposition I.8.19] for a proof. □

We conclude this Subsection with the observation that a group morphism $\phi : G \to H$ can be viewed as a map $\mathrm{Cay}(G) \to \mathrm{Cay}(H)$. Then it is easy to show the following:

Lemma 5.1.9. *A group morphism $\phi : G \to H$ is a quasi-isometry if and only if it has finite kernel and $Im(\phi)$ has finite index in H. In particular, group isomorphisms are quasi-isometries.*

5.1.3 Gromov-hyperbolic spaces and groups

Given three points $x, y, z \in X$ in a geodesic metric space, a *geodesic triangle* will be the union of three geodesic segments $[x, y] \cup [y, z] \cup [z, x]$.

Definition 5.1.10 (Gromov, [Gro87])**.** A geodesic metric space (X, d) is *Gromov-hyperbolic*, or simply hyperbolic, if there exists a constant δ, called the hyperbolicity constant, such that the following holds. For every geodesic triangle $[x, y] \cup [y, z] \cup [z, x]$ and any $p \in [x, y]$ there exists some $q \in [y, z] \cup [z, x]$ with $d(p, q) \leqslant \delta$.

A triangle satisfying Definition 5.1.10 is said to be *δ-thin*, and looks like in Figure 5.1:

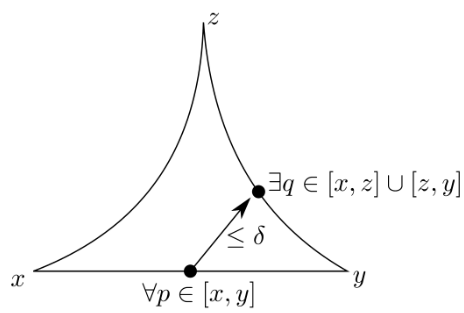

Figure 5.1: A δ-thin triangle. Every side is in the δ-neighborhood of the other two.

The next lemma shows that not only triangles, but also *quadrangles* (that is, the union of four consecutive geodesic segments) are slim:

Lemma 5.1.11. *Let $Q = [w, x] \cup [x, y] \cup [y, z] \cup [z, w]$ be a geodesic quadrangle inside a δ-hyperbolic space X. Then for every $p \in [w, x]$ there exists $q \in [x, y] \cup [y, z] \cup [z, w]$ with $d(p, q) \leqslant 2\delta$.*

Proof. Choose a geodesic $[w, y]$ and consider the two triangles $T = [w, x] \cup [x, y] \cup [y, w]$ and $T' = [w, y] \cup [y, z] \cup [z, w]$. Then in the first triangle p is δ-close to a point $r \in [x, y] \cup [y, w]$. If $r \in [y, w]$ we set $q = r$ and we are done, otherwise r itself is δ-close to a point $q \in [y, z] \cup [z, w]$. □

An important fact about hyperbolicity is that it is preserved under quasi-isometries, though the actual hyperbolicity constant δ may change:

Lemma 5.1.12. *Let X, Y be geodesic metric spaces. If there exists a quasi-isometric embedding $f : X \to Y$ and Y is hyperbolic, then so is X. In particular, being hyperbolic is a quasi-isometry invariant*

Proof. See e.g. [BH99, Theorem III.1.9] for a proof. ☐

Thanks to Lemma 5.1.12, we can say that a finitely generated group G is hyperbolic if one (hence any) of its Cayley graphs is hyperbolic. Equivalently, by Lemma 5.1.8 a group is hyperbolic if and only if it acts geometrically on a hyperbolic space.
For a hyperbolic metric space it is possible to define a notion of boundary, whose points represent the "directions at infinity" of the space. First, we recall the following notation. Given a point x in a metric space X, let $B(x, \varepsilon)$ be the ball of radius ε centered at x. Moreover, if $A \subset X$, we define its *ε-neighborhood* as $N_\varepsilon(A) = \bigcup_{a\in A} B_\varepsilon(a)$. Finally, given two subsets $A, B \subset X$ of a metric space, their *Hausdorff distance* is defined as

$$d_{Haus}(A, B) = \inf \{\varepsilon > 0 \mid B \subseteq N_\varepsilon(A), A \subseteq N_\varepsilon(B)\}.$$

Let $\mathcal{R}$ be the set of quasigeodesic rays $[0, +\infty) \to X$ inside X, which might be empty. We say that two rays γ, λ are asymptotic, and we write $\gamma \sim \lambda$, if their Hausdorff distance is finite. Intuitively, this means that the rays are "parallel" and co-oriented.

Definition 5.1.13. Let X be a δ-hyperbolic metric space. The *Gromov boundary ∂X* is the set of $\sim$-equivalence classes of quasigeodesic rays.

For every $p \in \partial X$, we say that every quasigeodesic ray γ which belongs to the class p is *asymptotic* to p, and we write $\gamma(+\infty) = p$.

Lemma 5.1.14. *There exists a constant $K(\delta)$ such that the following holds. For every $x \in X$ and every $p \in \partial X$ there exists a (K, K)-quasigeodesic ray $\gamma : [0, +\infty) \to X$ such that $\gamma(0) = x$ and $\gamma(+\infty) = p$. Moreover, for every two boundary points $p, q \in \partial X$ there exists a (K, K)-quasigeodesic line $\gamma : (-\infty, +\infty) \to X$ such that $\gamma(-\infty) = q$ and $\gamma(+\infty) = p$.*

Proof. See e.g. [KB02, Remark 2.16] for a proof. ☐

Roughly speaking, Lemma 5.1.14 says that the whole boundary is "visible" from any point $x \in X$, because every point in the boundary can be reached from x following a suitable quasigeodesic. The "moreover" part asserts that pairs of boundary points correspond to bi-infinite quasigeodesics, which will be very useful later. Notice that the constants of these quasigeodesics are the same for every x and every p in the boundary.

5.1.4 Acylindrical hyperbolicity

Definition 5.1.15. 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 $\mathbb{Z} \to \mathcal{S}$ given by $k \mapsto g^k(x)$ is a quasi-isometric embedding.

We call the orbit $\{g^k(x)\}$ a *(quasi)-axis* for the loxodromic element. If one precomposes the map in the definition with a quasi-isometry $\mathbb{R} \to \mathbb{Z}$ then one gets a bi-infinite quasigeodesic $\mathbb{R} \to \mathcal{S}$, whose endpoints at infinity are denoted $g^\pm \in \partial\mathcal{S}$. This pair does not depend on the choice of x, since if y is another point then $d_\mathcal{S}(g^k(x), g^k(y)) = d_\mathcal{S}(x, y)$ is constant, hence the two quasi-axis are at bounded Hausdorff distance and have the same endpoints in the boundary.

Definition 5.1.16 (Bestvina-Fujiwara [BF02]). In the same setting of Definition 5.1.15, 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

$$\left|\left\{h \in G\right| \max\left\{d_{\mathcal{S}}\left(x, h(x)\right), d_{\mathcal{S}}\left(g^N(x), hg^N(x)\right)\right\} \leqslant \varepsilon\right\}\right| < +\infty.$$

We denote by $\mathcal{L}_{WPD}$ the set of loxodromic WPD elements. It follows from the definitions that a power of a loxodromic and WPD element remains loxodromic and WPD.

Definition 5.1.17. A group G is *virtually-cyclic* if it has a finite index subgroup isomorphic to $\mathbb{Z}$. In particular, G and $\mathbb{Z}$ are quasi-isometric.

Definition 5.1.18 (Osin [Osi16]). A finitely-generated group G is *acylindrically hyperbolic* if it is not virtually-cyclic and it acts by isometries on a hyperbolic space $\mathcal{S}$, so that the action admits loxodromic WPD elements.

Remark 5.1.19. It follows from [Osi16, Theorem 1.1] that if G and $\mathcal{S}$ are as in Definition 5.1.18 then there exist at least two loxodromic elements g, h which are *independent*, meaning that the sets $\{g^{\pm}\}$ and $\{h^{\pm}\}$ are disjoint. Hence in particular $\partial\mathcal{S}$ has at least four points (and one can show that it is actually infinite).

5.1.5 A link to the past: more properties of the curve graph

Now that we have some more notions of coarse geometry we can say more about the curve graph and the action of the corresponding mapping class group. First we recall a celebrated theorem of Masur and Minsky [MM99, Min96]:

Theorem 5.1.20. *There exists a constant δ such that the curve graph of a surface S of complexity at least 1 (hence including the Farey complex) is δ-hyperbolic.*

Next, we identify the loxodromic WPD elements of this action:

Lemma 5.1.21 ([MM99, BF02]). *If S has complexity at least 2 then the action of a pseudo-Anosov element on the curve graph is loxodromic and WPD.*

Finally, the mapping class group of surfaces of complexity at least 2 is not virtually-cyclic. This is because whenever two curves are disjoint the corresponding Dehn twists commute and generate a $\mathbb{Z}^2$ subgroup, but a virtually-cyclic group does not contain a copy of $\mathbb{Z}^2$ (simply because for every non-trivial element $a \in \mathbb{Z}^2$ we can find some $b \in \mathbb{Z}^2$ such that no power of b is in $\langle a \rangle$). Combining these results we get:

Theorem 5.1.22. *The mapping class group of a surface of complexity at least 2 is non-virtually-cyclic and acts on its curve graph with loxodromic WPD elements, hence it is acylindrically hyperbolic.*

A similar result holds also for $MCG(S_4)$, but we will not need it.

Corollary 5.1.23. *The curve graphs of both S_4 and S_5 have at least four points at infinity.*

Proof. In the case of an S_4, it is easy to see that the boundary of the Farey complex is in bijection with the accumulation points of the endpoints of the arcs in Figure 1.3, which is dense in the boundary $\mathbb{S}^1$ of the disk. Regarding S_5, we know that $MCG(S_5)$ is non-virtually-cyclic, and it acts on $\mathcal{C}(S_5)$ with loxodromic WPD elements. Therefore the thesis follows by Remark 5.1.19. □

5.2 HHS background

This section is devoted to recover the useful facts about *hierarchically hyperbolic spaces and groups*, first defined by Behrstock, Hagen and Sisto in [BHS17]. Most of the time we will not go into the actual details, since a complete knowledge of the (rather lengthy) definitions will not be necessary for our purposes. It will suffice to have in mind what the various notions involved in the definitions are for mapping class groups, whose hierarchical features were already studied by Masur and Minsky in [MM00] (see Remark 5.2.5 below).

Definition 5.2.1. Roughly speaking, a *hierarchically hyperbolic space* (HHS) is the data of:

- a metric space X;
- an index set $\mathfrak{S}$ with a symmetric relation called *orthogonality* ($\perp$) and a partial order, called *nesting* ($\sqsubseteq$), with a unique maximal element $S \in \mathfrak{S}$; when two indices U, V are neither orthogonal nor $\sqsubseteq$-related then they are said to be *transverse* ($\pitchfork$);
- for every index $U \in \mathfrak{S}$, a space $\mathcal{C}U$, called the *curve graph* of U, which is uniformly hyperbolic (meaning that the hyperbolicity constant δ does not depend on U), and a uniformly coarsely Lipschitz map $\pi_U : X \to \mathcal{C}U$ (meaning that the constants of the coarsely Lipschitz property are the same for every U).

It is customary to denote a HHS simply by $(X, \mathfrak{S})$.

We may think of the projections $\pi_U : X \to \mathcal{C}U$ as coordinates, meaning that every point $x \in X$ is uniquely determined by the tuple $\{\pi_U(x)\}_{U \in \mathfrak{S}}$ of its coordinates.
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

$$d_X(x, y) \asymp_{D,D} \sum \{\{d_Y(x, y)\}\}_s \tag{5.1}$$

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 and is a HHS itself "of lower complexity" with index set $\mathfrak{S}_U = \{V \in \mathfrak{S} : V \sqsubseteq U\}$. A point $x \in F_U$ has all coordinates of the form $\pi_V(p)$ prescribed whenever $U \perp V$, $U \subsetneq V$ and $U \pitchfork V$, thus the indices in $\mathfrak{S}_U$ are the only coordinates needed to describe x.
Now, 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}. The product region $P_{\{U_i\}}$ has the property that for all U with $U_i \subsetneq U$ or $U \pitchfork U_i$ for some i, we have that $\pi_Y(P_{\{U_i\}})$ is a uniformly bounded set, that we denote $\rho_V^{\{U_i\}}$.
The following is [BHS21, 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 5.2.2. *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*

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

where the sum is made on all W such that $U_i \subsetneq W$ or $U_i \pitchfork W$ for some U_i.

From this Lemma we obtain a criterion for two product regions to intersect in a bounded set, in the following "coarse" sense:

Definition 5.2.3. Given two subsets A, B of a metric space X, if there exists R_0 so that for any $R, R' \geqslant R_0$, we have $d_{Haus}(N_R(A), N_{R'}(B)) < +\infty$ then the *coarse intersection* of A and B, which we denote $A \tilde{\cap} B$, is any subspace of X within bounded Hausdorff distance of all $N_{R_0}(A) \cap NR_0(B)$.

Corollary 5.2.4 (Intersection of different product regions)**.** *If $\{U_i\}$ and $\{V_j\}$ are maximal collections of pairwise orthogonal indices, such that $\{U_i\} \cap \{V_j\} = \varnothing$ and every V_j is $\sqsubseteq$-minimal, then $P_{\{U_i\}} \tilde{\cap} P_{\{V_j\}}$ is coarsely a point (that is, it is bounded).*

Proof. It is enough to apply Lemma 5.2.2 with $x \in P_{\{V_j\}}$. In fact, for every $W \notin \{V_j\}$ there is a V_j such that $V_j \pitchfork W$ or $V_j \subsetneq W$, since $\{V_j\}$ is a maximal collection of pairwise orthogonal, $\sqsubseteq$-minimal indices. Moreover, since $x \in P_{\{V_j\}}$, the projection $\pi_W(x)$ is contained in the uniformly bounded set $\rho_W^{\{V_j\}}$. Hence the index W contributes to the sum with a term of the form

$$\left\{d_W\left(\rho_W^{\{V_j\}}, \rho_W^{\{U_i\}}\right)\right\}_s$$

that does not depend on x, but only on $\{U_i, V_j\}$ and W. Now, if we fix a threshold s_0 there is a finite number of terms that contribute to the sum, since the sum is bounded above in terms of $d_X(x, P_{\{V_j\}})$. Therefore, if we choose a big enough threshold s we can erase this terms from the sum, and 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 $\mathcal{C}V_j$. Since all coordinates for $W \notin \{V_j\}$ are prescribed we must have that all these points have the same coordinates (up to a finite error), and therefore they are close in X by the Distance Formula 5.1. $\square$

Remark 5.2.5 (The motivating example). The mapping class group of a surface S is a hierarchically hyperbolic space, by combining results of Masur, Minsky, Behrstock, Kleiner and Mosher [MM99, MM00, Beh06, BKMM12]. More precisely:

- the index set $\mathfrak{S}$ is the set of essential subsurfaces (up to isotopy);
- nesting corresponds to containment of subsurfaces, and the maximal element is the whole surface S;
- orthogonality is disjointness of subsurfaces;
- $\mathcal{C}U$ are the "real" curve graphs, which are uniformly hyperbolic by Theorem 5.1.20;
- $\pi_U : MCG \to \mathcal{C}U$ is defined using subsurface projection;

- if $U \subsetneq V$ or $U \pitchfork V$ then ρ^U_V is the subsurface projection of ∂U to $\mathcal{C}V$;
- F_U can be thought of as the mapping class group of U, and has a hierarchical structure whose index set are the subsurfaces of U;
- if $U_1, \ldots, U_k$ are pairwise disjoint subsurfaces the corresponding mapping class groups commute, and $P_{\{U_i\}}$ can be thought of as the product of these subgroups;
- the distance formula is a result of [MM00].

5.2.1 Quasiflats in HHS

This section collects some results from [BHS21] about how quasi-isometric copies of $\mathbb{R}^n$ can be arranged in a HHS. First, a definition.

Definition 5.2.6. Let X be a metric space. A *k-dimensional quasiflat* is (the image of) a quasi-isometric embedding $\mathbb{R}^k \to X$. Similarly, a *k-dimensional orthant* is (the image of) a quasi-isometric embedding $\mathbb{R}^k_+ \to X$, where $\mathbb{R}_+ = [0, +\infty)$.

Notice that one-dimensional quasiflats and orthants are just quasigeodesic lines and rays, respectively.

Definition 5.2.7 (Hierarchy paths)**.** For $D \geqslant 1$, a (D, D)-quasigeodesic γ in X is a *D–hierarchy path* if for each $W \in \mathfrak{S}$, the path $\pi_W(\gamma) \subset \mathcal{C}W$ is an unparameterized (D, D)–quasigeodesic (meaning that it admits a reparameterization as a (D, D)–quasigeodesic).

In other words, hierarchy paths are quasigeodesics which are "compatible" with the HHS structure, since the "coordinates" of γ are themselves quasigeodesics with the same constants (up to choosing a different parameterization). One can similarly define hierarchy rays and lines. The following is [BHS19, Theorem 4.4]:

Lemma 5.2.8. *There exists $D_0 \geqslant 1$ such that, for every $x, y \in X$ there exists a D_0-hierarchy path joining them.*

Notice that, since for every $U \in \mathfrak{S}$ the space F_U is itself a HHS, it makes sense to talk about hierarchy paths in F_U. Now, given a product region $P_{\{U_i\}} = \prod_i F_{U_i}$, if we choose a hierarchy line (resp. ray) γ_i inside F_{U_i} for $i \leqslant k$ and a point $p_i \in F_{U_i}$ for $i > k$ we can consider their product $\prod_{i \leqslant k} \gamma_i \times \prod_{i>k}\{p_i\}$, which is a k-dimensional quasiflat (resp. orthant). This is what we will refer to as *standard k-flats (resp., orthants)*. 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. We say that ν is the *rank* of our HHS. Now we claim that, if U is in a complete support set and $\partial\mathcal{C}U$ is non-empty, we can prescribe the asymptotic behavior of hierarchy rays and lines in F_U. This requires the following technical condition, which is [BHS21, Definition 1.14]:

Definition 5.2.9 (Asymphoricity)**.** We say that the HHS $(X, \mathfrak{S})$ of rank ν is *asymphoric* if there exists a constant A with the property that there does not exist a set of $\nu + 1$ pairwise orthogonal elements U of $\mathfrak{S}$ where each $\mathcal{C}U$ has diameter at least A.

From now on we will always assume without saying that all HHS are asymphoric. We will moreover assume that they are *normalized*, meaning that every projection $\pi_U : X \to \mathcal{C}U$ is uniformly coarsely surjective.

Lemma 5.2.10. *Let U be an element of a complete support set. For every $p \in \partial\mathcal{C}U$ there exists a hierarchy ray $\gamma \subset F_U$, such that $\pi_U(\gamma)$ is a quasi-ray which is asymptotic to p, while $\pi_V(\gamma)$ is uniformly bounded for every $V \sqsubseteq U$.*
Moreover, for every pair of distinct points $p^\pm \in \partial\mathcal{C}U$ there exists a hierarchy line $l \subset F_U$, such that $\pi_U(l)$ is a quasigeodesic which is asymptotic to $p^\pm$, while $\pi_V(l)$ is uniformly bounded for every $V \sqsubseteq U$.

Proof. The asymphoricity assumption implies that, whenever U is an element of a complete support set and $V, W \sqsubsetneq U$, if $V \perp U$ then either $\operatorname{diam}(\mathcal{C}V) \leqslant A$ or $\operatorname{diam}(\mathcal{C}W) \leqslant A$. By [BHS21, Corollary 2.16] this is equivalent to the fact that F_U is δ'-hyperbolic, for some constant δ' depending only on the hyperbolicity constant δ of $\mathcal{C}U$, which is uniform for all U, and on the HH structure.
Now we show that the first statement holds, since the “moreover” part has a similar proof. The proof will be similar to that of [KB02, Remark 2.16].
Let $\eta : [0, +\infty) \to \mathcal{C}U$ be a (K, K)-quasi-ray converging to p, where K is the constant from Lemma 5.1.14. For every $n \in \mathbb{N}$ let $x_n \in F_U$ be a point whose projection $\pi_U(x_n)$ is uniformly close to $\eta(n)$, which exists by the normalization assumption. Notice that, since π_U is coarsely Lipschitz, for every fixed $m \in \mathbb{N}$ we have $\lim_{n\to+\infty} d(x_m, x_n) = +\infty$.
For every $n \in \mathbb{N}$ let h_n be a hierarchy path from x_0 to x_n. We claim that for every $k \in \mathbb{N}$ there exists $N_0 = N_0(k)$ such that for every $n, m \geqslant N_0$, the initial segment of h_n of length k and the initial segment of h_m of length k are uniformly close in the Hausdorff distance. To see this, consider a hierarchy path l_{mn} from x_m to x_n, so that $T = h_n \cup h_m \cup l_{mn}$ is a triangle of (D, D)-quasigeodesics. By e.g [BH99, Corollary III.1.8] there exists δ'', depending only on D and δ', such that this triangle is δ''-slim. Moreover, since l_{mn} is a hierarchy path it projects to an unparameterized (D, D)-quasigeodesic from $\pi_U(x_m)$ to $\pi_U(x_n)$, and by stability of quasi-geodesics in hyperbolic spaces (see e.g. [BH99, Theorem III.1.7]) there is an M, depending only on D and δ, such that $\pi_U(l_{mn})$ and $\eta_{[m,n]}$ are M-Hausdorff-close. Hence the projection of l_{mn} is also arbitrarily distant from $\pi_U(x_0)$ when $m, n \to +\infty$, and therefore we can assume that every point in $l_{m,n}$ is at distance at least $k + 100\delta''$ from x_0. Now pick a point z on h_m at distance at most k from x_0. Since the triangle T is δ''-slim there must be some $q \in h_n \cup l_{mn}$ which is δ''-close to p. However, q cannot lie on l_{mn}, because otherwise $d(x_0, l_{mn}) \leqslant k + \delta''$. Therefore we proved that a segment of length k of h_m is δ''-close to an initial segment of fixed length of h_n, and symmetrically the same holds if we swap m and n.
Now, arguing as in [KB02, Remark 2.16] we can find a quasigeodesic ray γ whose constants are bounded in terms of δ and D and whose projection to $\mathcal{C}U$ is uniformly close to η. To show that γ is a hierarchy path it suffices to prove that, whenever $V \sqsubsetneq U$, we have that $\pi_V(\gamma)$ is bounded. This is true since by the second of the consistency inequality [BHS19, Definition 1.1.(3)] the projection of γ to $\mathcal{C}V$ is the union of the projection of a bounded initial segment (which is bounded) and a piece which is uniformly close to $\rho^U_V(\eta)$. In turn, the latter is uniformly bounded by the bounded geodesic image axiom [BHS19, Definition 1.1.(7)]. □

Now suppose that all curve graphs have non-empty Gromov boundary, as will be the case for us. In this case the cardinality ν of a complete support set is the maximal dimension of a standard orthant, because we can choose a hierarchy ray h_i in every F_{U_i} and take their product. Moreover, thanks to Lemma 5.2.10, we can also ask that the projection $\pi_{U_i}(h_i)$ is asymptotic to a prescribed boundary point $p_i \in \partial\mathcal{C}U_i$. Similarly, 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 a hierarchy line in each F_{U_i} whose projection to $\mathcal{C}U_i$ has endpoints $p_i^\pm$.

5.2.2 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 in the sense of Definition 5.2.3.

Remark 5.2.11. The coarse intersection of any two standard flats $\mathfrak{F}_1, \mathfrak{F}_2$ is well-defined by [BHS21, Lemma 4.12]. In fact, the lemma gives a description of the coarse 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 quasigeodesic lines with the same endpoints in the Gromov boundary.

Definition 5.2.12 (Hinge graph)**.** From now on let X be an asymphoric and normalized HHS. As in [BHS21, Definition 5.2], let $\mathbf{Hinge}(\mathfrak{S})$ be the set of *hinges*, that is, pairs (U, p) where U is contained in a complete support set and $p \in \partial\mathcal{C}U$. 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 $\mathbf{Hinge}(\mathfrak{S})$ a metric graph structure by declaring two hinges to be adjacent if and only if they are compatible.

Remark 5.2.13. As in [BHS21, Definition 5.3], one can associate to a hinge $\sigma = (U, p)$ a standard 1-orthant, denoted h_σ. This is the hierarchy ray constructed in Lemma 5.2.10, whose projection to $\mathcal{C}(U)$ is a quasigeodesic ray asymptotic to p, and whose projections to all other curve graphs are uniformly bounded. Another property of h_σ, stated in [BHS21, Remark 5.4], is that if $\sigma \neq \sigma'$ are two different hinges then $d_{Haus}(h_\sigma, h_{\sigma'}) = +\infty$. In fact, if $\sigma = (U, p)$ and $\sigma' = (U, p')$ then $\pi_U(h_\sigma)$ and $\pi_U(h_{\sigma'})$ are two quasi-rays that must diverge; if $\sigma = (U, p)$ and $\sigma' = (U', p')$ with $U \neq U'$ then $\pi_{U'}(h_\sigma)$ is bounded while $\pi_{U'}(h_{\sigma'})$ is not. Either way there is a curve graph where $\pi_U(h_\sigma)$ and $\pi_U(h_{\sigma'})$ are at infinite Hausdorff distance, and therefore $d_{Haus}(h_\sigma, h_{\sigma'}) = +\infty$ by the Distance Formula (Equation 5.1).

Lemma 5.2.14. *Let $(X, \mathfrak{S})$ be an asymphoric and normalized 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 $h_{\sigma'}$.*

Proof.
1. Choose another ideal point $q \in \partial\mathcal{C}U \backslash \{p\}$, 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 5.2.11.
2. The proof is very similar, in this case fixing two hierarchy lines rather than one.

□

Theorem 5.2.15. *Let $(X, \mathfrak{S})$, $(Y, \mathfrak{T})$ be asymphoric and normalized 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 f_{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 [BHS21, Theorem 5.7], which applies to HHSs satisfying three additional assumptions, stated in [BHS21, Section 5] but not satisfied in our case. An inspection of the proof shows that each use of the three assumptions can be replaced by our Lemma 5.2.14. □

The following is the analogue of [BHS21, Lemma 5.9], which is a corollary of [BHS21, 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 induce map f_{hin} on hinge graphs. Crucially for applications, the bound is only in terms of the HHSs and the quasi-isometry constants.

Lemma 5.2.16 (Flats go to flats)**.** *Let $(X, \mathfrak{S})$, $(Y, \mathfrak{T})$ be asymphoric and normalized 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 $i = 1, \ldots, \nu$ we have $f_{hin}(U_i, p_i^{\pm}) = (V_i, q_i^{\pm})$. Then, $d_{Haus}(f(\mathfrak{F}_{\{(U_i, p_i^{\pm})\}}), \mathfrak{F}_{\{(V_i, q_i^{\pm})\}}) \leqslant C_0$.*

Chapter 6

Quasi-isometric rigidity of pants graphs of spheres

For the rest of the thesis let $S = S_b$ be a sphere with $b \geqslant 7$ punctures.

Definition 6.0.1 (Hatcher-Thurston [HT80]). The *pants graph* $\mathbb{P}(S_b)$ is the simplicial graph such that:

- there is a vertex for every pants decomposition $\Delta = \{\gamma_1, \ldots, \gamma_{b-3}\}$, that is, for every maximal (unordered) simplex of $\mathcal{C}(S_b)$;
- two pants decompositions Δ and Δ' are joined by and edge if and only if they share a facet P, such that $\Delta = \gamma \star P$, $\Delta' = \gamma' \star P$ and $i(\gamma, \gamma') = 2$ (that is, the two curves have minimal intersection in the S_4 that P cuts out).

By combining a theorem of Brock [Bro03] with [MM99, MM00, Beh06, BKMM12], we get that $\mathbb{P}(S_b)$ is hierarchically hyperbolic with the following structure:

- $\mathfrak{S}$ is the set of isotopy classes of essential, non-annular subsurfaces;
- $\perp$ is disjointness and $\sqsubseteq$ is inclusion;
- for every $U \in S$, $\mathcal{C}U$ is the curve graph;
- Projections are defined using subsurface projections. More precisely, for every pants decomposition Δ, which we see as a maximal simplex of $\mathcal{C}(S)$, its projection to a subsurface U is $\bigcup_{\gamma \in \Delta} \pi_U(\gamma)$, which is uniformly bounded since subsurface projections are Lipschitz 1.3.1.

This structure is very similar to that of the mapping class group, described in Remark 5.2.5. The only important difference is that the index set for the pants graph does not contain annuli, thus if a connected subsurface U is $\sqsubseteq$-minimal then it is a four-holed sphere. As we will see in Remark 8.0.2, the quotient MCG/DT_K will have a similar structure, whose minimal indices are (DT_K-orbits of) four-holed spheres.
The goal of this Chapter is to prove a quasi-isometric rigidity result for pants graphs of punctured spheres, as a blueprint for the case of MCG/DT_K. We will partly recover a result of Bowditch [Bow20, Theorem 1.4] with our machinery, though the proof will not be completely new since we will still rely on some theorems of the named author from [Bow20, Sections 6 and 7] and on the main theorem of [Bow16].

Theorem 6.0.2 (Quasi-isometric rigidity of the pants graph)**.** *Let S_b be a punctured sphere, with $b \geqslant 7$. Any self-quasi-isometry f of the pants graph $\mathbb{P}(S_b)$ is at uniformly bounded distance from an element of the extended mapping class group.*

The "source" of rigidity here will be the fact that automorphisms of a certain graph, the graph of 1-separating curves, can only be extended mapping classes [Bow20, Bow16]. Hence our goal is to show that automorphisms of the hinge graph induce automorphisms of this other graph, as we shall do in Corollary 6.2.3, and in order to show this the key thing to do is roughly the following. Since hinges are not just subsurfaces, but rather pairs (U, p) where U is a subsurface and p is a point in the boundary of its curve graph, we have to be able to determine which hinges have the same support subsurface in a "combinatorial" way, that is, just by looking at the hinge graph. This is not always possible, but we will be able to do so for "enough" hinges.

6.1 Unambiguous subsurfaces

We recall that a surface is said to be *odd* if its complexity is odd, otherwise it is *even* (in our case, ironically S_b is odd iff b is even). Bowditch [Bow20, Section 6] pointed out that if a subsurface U belongs to a complete support set 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 from S. If S is odd then U is an S_4 and each connected component of the complement $S \backslash U$ is odd and meets U in exactly one curve. If S is even then U is either:

1. an S_4 with all but one of the complementary components odd;
2. an S_5 with all complementary components odd.

First we ensure that we are in the hypotheses of Section 5.2.2.

Lemma 6.1.1. *The HHS structure of the pants graph is asymphoric and normalized.*

Proof. The projection $\pi_U : \mathbb{P}(S) \to \mathcal{C}U$ is surjective, since every curve on U can be completed to a pants decomposition of S. Regarding asymphoricity, it suffices to notice that the curve graph of an essential, non-annular subsurface U is unbounded by Corollary 1.4.8, thus every collection of pairwise disjoint subsurfaces has cardinality less than the rank. □

Now, by Corollary 5.1.23, both $\mathcal{C}(S_4)$ and $\mathcal{C}(S_5)$ have at least four boundary points, therefore we are always in the assumptions of Theorem 5.2.15. Thus a quasi-isometry f of $\mathbb{P}(S_b)$ induces an automorphism f_{hin} of the hinge graph, and we want to show that f_{hin} maps hinges with the same underlying subsurface to hinges with the same underlying subsurface, at least in the vast majority of cases.
Given a hinge $\sigma = (U, p)$ 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 6.1.2. Two hinges σ, σ' are said to be *equally completable*, or to have the same completions, if $\mathrm{Compl}(\sigma) = \mathrm{Compl}(\sigma')$.

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 are equally completable then they have the same underlying subsurface. However this is not always true: in the even case, an S_4 whose even complementary component is a pair of pants P has the same completions as the S_5 given by the union of S_4 and P. See Figure 6.1 to understand the situation. Luckily we will see that this is the only problem that could arise.

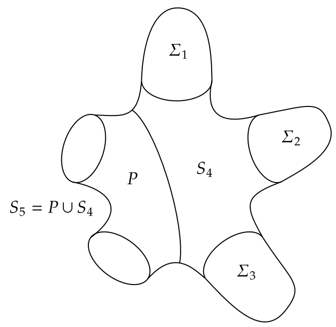

Figure 6.1: The S_4 and the S_5 in the Figure cannot be distinguished just by their completions, all of which must be supported in the union of the Σ_i's.

Definition 6.1.3. A hinge σ is *unambiguous* if any other equally completable hinge σ' has the same support. A support U is unambiguous is every hinge supported in U is unambiguous.

We need another definition that characterizes (some) unambiguous surfaces and can be recognized inside the hinge graph.

Definition 6.1.4. A hinge σ is *minimal* if $\mathrm{Compl}(\sigma)$ is maximal by inclusion, among completions.

Lemma 6.1.5. *Let $S = S_b$ be a sphere with $b \geqslant 7$ punctures. If the complexity is odd, every hinge is minimal and unambiguous. If the complexity is even, a hinge is minimal if and only if it is unambiguous and its support is a four-holed sphere.*

Proof. Let $\sigma = (U, p)$ be a hinge. We examine all possible shapes of U to determine in which cases σ is minimal and/or unambiguous. In the odd case U is always an S_4 and each component of the complement is odd. Then we may find a completion $\{U_i\}$ whose union is $S \backslash U$: it suffices to choose a pants decomposition of every complementary component Σ_i, which must contain an even number of pairs of pants (since Σ_i must be odd), and then match these pants in couples to get some four-holed spheres whose union covers Σ_i. Now U must be unambiguous and minimal, since any other V completing $\{U_i\}$ must be inside $S \backslash \bigcup U_i = U$ and therefore coincide with U which has already minimal complexity. In the even case there are two possibilities:

1. Suppose U is an S_5. Since $b \geqslant 7$ we know that U does not coincide with S, so set $S \backslash U = \bigsqcup_{j=1}^{k} \Sigma_j$ with $k \leqslant 5$. Cover every Σ_j with four-holed spheres. Moreover, choose a pair of pants $P \subset U$ whose boundary touches some four-holed sphere $W \subseteq \Sigma_1$, and let $U' = U \backslash P$. Then U is not minimal: any completion of U works also for U', but if we replace W with $W' = W \cup P$ we get a completion for U' but not for U. See Figure 6.2 to understand the situation.

2. Suppose U is an S_4. Cover the odd complementary components with four-holed spheres. If the even component is not a pair of pants then we can find a way to cover it with a single S_5 and some S_4, and argue that U must be minimal and unambiguous as in the odd case. Otherwise U is ambiguous, since we are precisely in the case described in Figure 6.1. This also means that U has the same completions of a five-holed sphere, which is not minimal as we have already showed.

□

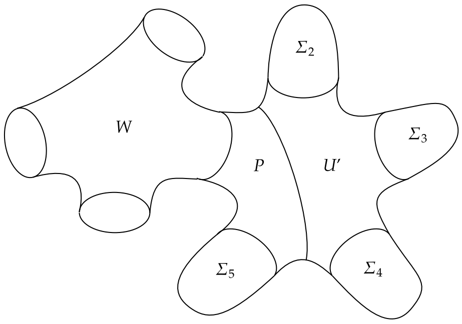

Figure 6.2: Replacing W with $W' = W \cup P$ we get a completion for U' but not for $U = U' \cup P$.

Remark 6.1.6. 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 U is a minimal support, we can define $f_{supp}(U)$ as the support of $f_{hin}(U, p)$ for any hinge (U, p) supported in U.

Corollary 6.1.7 (Products go to products)**.** *Let $b \geqslant 7$ and let f be a self-quasi-isometry of the pants graph $\mathbb{P}(S_b)$. There exists C, depending only on the quasi-isometry constants of f, with the following property. Let $\{U_i\} \subseteq \mathfrak{S}$ be a complete support set made of minimal supports, and let $f_{supp}(U_i) = V_i$ for all i. Let $P_{\{U_i\}}$ and $P_{\{V_i\}}$ be the standard product regions defined by $\{U_i\}$ and $\{V_i\}$, respectively. Then $d_{Haus}\left(f\left(P_{\{U_i\}}\right), P_{\{V_i\}}\right) \leqslant C$.*

Proof. Another way of stating Remark 6.1.6 is that, if (U, p^+) and (U, p^+) are two hinges with the same minimal support and $f_{hin}(U, p^+) = (V, q^+)$, then there exists $q^- \in \partial\mathcal{C}V$ such that $f_{hin}(U, p^-) = (V, q^-)$. Thus the Flats to Flats Lemma 5.2.16 says that, if $\{(U_i, p_i^\pm)\}$ is a complete support set made of minimal supports, with a choice of two points in every $\partial\mathcal{C}U_i$, and if we set $(V_i, q_i^\pm) := f_{hin}(U_i, p_i^\pm)$, then $d_{Haus}(f(\mathfrak{F}_{\{(U_i,p_i^\pm)\}}), \mathfrak{F}_{\{(V_i,q_i^\pm)\}}) \leqslant C_0$ for some constant C_0 depending only on the quasi-isometry constants of f. Hence

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

where $\bigcup \mathfrak{F}_{\{(U_i,*)\}}$ is the union of all standard flats supported in $\{U_i\}$.
Now, the thesis follows if we prove the existence of some constant C_1 such that, for every complete support set $\{U_i\}$ made of minimal supports,

$$d_{Haus}\left(\bigcup \mathfrak{F}_{\{(U_i,*)\}}, P_{\{U_i\}}\right) \leqslant C_1.$$

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

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

where C_2 is some constant depending only on C_1 and the quasi-isometry constants of f. In turn, since $P_{\{U_i\}} = \prod F_{U_i}$ with the product metric, 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 6.1.8. ☐

Lemma 6.1.8. *There exists $C_3 \geqslant 0$ such that whenever $U \in \mathfrak{S}$ is a minimal domain and $x \in F_U$, there exists a hierarchy line $\gamma \subset F_U$ such that $d_{F_U}(x, \gamma) \leqslant C_3$.*

Proof. Notice that, by the distance formula (Equation 5.1) and the fact that U is an S_4, and therefore a $\sqsubseteq$-minimal support, we have that F_U is uniformly quasi-isometric to $\mathcal{C}U$. Moreover, a hierarchy line $\gamma \subset F_U$ simply corresponds to a bi-infinite quasigeodesic (with certain constants), since for every $V \neq U$ the coordinate $\pi_V(\gamma)$ is prescribed, and therefore constant. Finally recall that $MCG(U)$ acts on $\mathcal{C}U$ coboundedly and by isometries, hence every point can be moved within uniformly bounded distance from a fixed quasigeodesic (which exists by the "moreover" part of Lemma 5.1.14). $\square$

6.2 Automorphism of terminal subsurfaces

Definition 6.2.1. A support U is *terminal* if it has complexity 1 and it is cut out by a single curve. We say that the boundary curve of a terminal support is 1-*separating*, and we denote by $\mathcal{C}^1$ the full subgraph of $\mathcal{C}$ spanned by 1-separating curves.

Lemma 6.2.2. *A hinge $\sigma = (U, p)$ has terminal support if and only if it is minimal and there exists a hinge (V, q), compatible with σ, such that any complete support set containing (V, q) must contain some σ' which has the same completions as σ (i.e., it has the same support). In particular, having terminal support is preserved by automorphisms of the hinge graph.*

Proof. This is just a restatement in our context of [Bow20, Lemma 6.3]. We just have to notice that a terminal hinge is minimal, which follows from the discussion of Lemma 6.1.5 since a terminal support has a single complementary component. $\square$

Corollary 6.2.3. *If $b \geqslant 7$, any automorphism of the hinge graph induces an automorphism of $\mathcal{C}^1$.*

Proof. Any automorphism must map terminal supports to terminal supports and must preserve compatibility (which, for terminal subsurface, is equivalent to disjointness of their boundaries). $\square$

Now we are ready to prove Theorem 6.0.2, that is, that any self-quasi-isometry f of the pants graph $\mathbb{P}(S_b)$ is at uniformly bounded distance from an element of the extended mapping class group.

Proof of Theorem 6.0.2. The previous discussion shows that f induces an automorphism of $\mathcal{C}^1$, which is the restriction of some extended mapping class $g \in MCG^{\pm}$ when $b \geqslant 7$ by Bowditch's extension of Ivanov's Theorem, which is Theorem 1.4.11. In other words, f_{supp} and g agree on terminal subsurfaces, and we want to show that they coincide on every minimal support U. We recall that, as showed in the proof of Lemma 6.1.5, every complementary components Σ of U has complexity at least 1, i.e., U does not cut out any pair of pants. Now, if Σ has complexity 1 then it is a terminal subsurface; otherwise Σ has complexity at least 2, and therefore there exist two terminal subsurfaces inside Σ whose boundary curves fill Σ (for example, we can choose a 1-separating curve $\alpha \in \mathcal{C}(\Sigma)$ and apply a high enough power of a pseudo-Anosov element of $MCG(\Sigma)$ to get a pair of filling curves, as in Lemma 1.4.7). Thus U is the unique minimal support that is disjoint from all these terminal subsurfaces, and therefore f_{supp} and g must agree on U since they both preserve disjointness.
By Corollary 6.1.7 if $\{U_i\}$ is a complete support set made of minimal supports then f maps the corresponding product region $P_{\{U_i\}}$ within uniformly bounded Hausdorff distance from $g(P_{\{U_i\}})$, since they are both uniformly Hausdorff close to $P_{f_{supp}\{U_i\}} = P_{g\{U_i\}}$.
We are left to prove that every point $x \in \mathbb{P}(S_b)$ 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 $g(P)$ and $g(P')$, which is coarsely $g(x)$. In order to prove this, it is enough to show that there is some point $x \in \mathbb{P}(S_b)$ which is the coarse intersection of two standard product regions with minimal supports, since the mapping class group acts cofinitely on $\mathbb{P}(S_b)$ by Lemma 1.4.1. In turn, by Corollary 5.2.4 we are left to prove that there exist two complete support sets $\{U_i\}$ and $\{V_i\}$, with minimal supports, whose indices are pairwise distinct (notice that we can apply Corollary 5.2.4 since minimal indices are S_4, and therefore they are also minimal with respect to nesting). If the surface is odd choose any complete support set $\{U_i\}$, which is already minimal. In the even case let $\tilde{U}$ be an S_5 cut out by a single curve, and let $\{U_i\}$ be a support set that covers the complement of $\tilde{U}$. Since $\{\tilde{U}\} \cup \{U_i\}$ covers the surface, every U_i must be minimal (this is the same argument as in the proof of Lemma 6.1.5). Now if we replace $\tilde{U}$ with a terminal S_4 contained in $\tilde{U}$, call it U_1, we get a complete support set $\{U_i\}$ made of minimal supports. In both cases, we can choose a pseudo-Anosov mapping class ϕ such that every boundary curve of $\{U_i\}$ crosses every boundary curve of $\{V_i\} = \phi\{U_i\}$. Therefore $U_i \neq V_j$ for every choice of i and j. $\square$

Chapter 7

Extending automorphisms of some subgraphs of the curve graph

Recall that we denoted by $\mathcal{C}^1$ the subgraph of $\mathcal{C}$ spanned by 1-separating curve, as in Definition 6.2.1. Moreover we recall the following definition from [Bow16]:

Definition 7.0.1. 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 $\mathcal{C}/DT_K$.
A hidden but fundamental piece of the proof of Theorem 6.0.2 was the fact that every automorphism of $\mathcal{C}^1$ is induced by an extended mapping class. Therefore, if we want to emulate the same proof for $MCG^{\pm}/DT_K$, we must establish the following analogue for $\mathcal{C}^1/DT_K$. For short, for the rest of the thesis 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 7.0.2. *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 4.4.1 we get the following:

Theorem 7.0.3. *For all $b \geqslant 7$ and for all large multiples K, any automorphism of $\mathcal{C}^1/DT_K$ is the restriction of an element of $MCG^{\pm}(S_b)/DT_K$.*

Following the roadmap of [Bow16, Bow20] we split the proof of Theorem 7.0.2 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 $\mathcal{C}/DT_K$.

7.1 From 1-separating to strongly separating

For the rest of the Chapter, 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 4. We first recall some definitions:

Definition 7.1.1. 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 3.1.1.
The following definition is from [Bow20, Section 7]:

Definition 7.1.2. 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 [Bow20, Lemma 7.2] and the discussion following [Bow20, Lemma 7.5], any division of $\mathcal{C}^1$ is induced by a unique strongly separating curve $\delta \in \mathcal{C}^{ss} \backslash \mathcal{C}^1$: 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 $(\mathcal{C}^{ss}/DT_K)\backslash(\mathcal{C}^1/DT_K)$.

Definition 7.1.3. 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. $P^+ = \mathrm{Lk}_\Gamma(\alpha^-) \cap \mathrm{Lk}_\Gamma(\beta^-)$, 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 will not consider divisions which do not allow a filling).

Remark 7.1.4. 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 subsurfaces 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 7.1.11, and the situation will be depicted in Figure 7.2.

Here we give another, simpler definition which does not explicitly involve the underlying division, and then prove that it is just a restatement of Definition 7.1.3.

Definition 7.1.5. 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 $d_\Gamma(\gamma, \gamma') = 1$.

Lemma 7.1.6. *Definitions 7.1.3 and 7.1.5 are equivalent.*

Proof. It is obvious that if $(P^\pm, \alpha^\pm, \beta^\pm)$ is a filled division in the sense of Definition 7.1.3 then the four curves satisfy Definition 7.1.5. Conversely, let $(\alpha^\pm, \beta^\pm)$ be two pairs of curves satisfying Definition 7.1.5, and we show that $P^\pm := \mathrm{Lk}_\Gamma(\alpha^\mp) \cap \mathrm{Lk}_\Gamma(\beta^\mp)$ satisfy Definition 7.1.3. By Conditions 2 and 3 of Definition 7.1.5 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. $\square$

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

Definition 7.1.7. 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 one into the other if they are disjoint.

Now we want to assign to every $\overline{\delta} \in (\mathcal{C}^{ss}/DT_K)\backslash(\mathcal{C}^1/DT_K)$ some filled division $(\overline{P}^\pm, \overline{\alpha}^\pm, \overline{\beta}^\pm)$. First, choose representatives $\delta_1, \ldots, \delta_k \in \mathcal{C}^{ss}\backslash\mathcal{C}^1$ for every homeomorphism type of strongly separating, non-1-separating curves (we know that there are finitely many representatives, because of Lemma 1.4.1). For every $i = 1\ldots, k$ 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 $\delta \in \mathcal{C}^{ss}\backslash\mathcal{C}^1$ choose some mapping class f such that $\delta = f(\delta_i)$ for some $i \in \{1, \ldots, k\}$, and set $(\alpha^\pm(\delta), \beta^\pm(\delta)) = f(\alpha_i^\pm, \beta_i^\pm)$, which are again two pairs of filling curves for the sides of δ. Now, choose Θ big enough that, for every $i = 1, \ldots, k$ and for every $s \in \mathcal{C}$, we have that $\max\{d_s(\alpha_i^\pm, \beta_i^\pm), d_s(\alpha_i^\pm, \beta_i^\mp)\} \leqslant \Theta$ (this is always possible, as we showed in the proof of Lemma 2.3.8). 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 $\delta \in \mathcal{C}^{ss}\backslash\mathcal{C}^1$ and for every $s \in \mathcal{C}$,

$$\max\{d_s(\alpha^\pm(\delta), \beta^\pm(\delta)), d_s(\alpha^\pm(\delta), \beta^\mp(\delta))\} \leqslant \Theta. \tag{7.1}$$

This way, if K is a large enough multiple with respect to this constant Θ we will be able to invoke Lemma 2.3.7 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 thesis.
Now let $(\overline{\alpha}^\pm, \overline{\beta}^\pm) := \pi(\alpha^\pm(\delta), \beta^\pm(\delta))$.

Lemma 7.1.8. *The two pairs $(\overline{\alpha}^\pm, \overline{\beta}^\pm)$ give a filled division.*

Proof. We show that the conditions of Definition 7.1.5 are satisfied. First notice that, since by construction all projections between $\alpha^\pm$ and $\beta^\pm$ are short (see Equation (7.1)), Lemma 2.3.7 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 Definition 7.1.5. 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 2.3.4. 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 2.3.12. 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 since g sends the terminal surface bounded by x to a terminal curve, thus $g^n(x) \in P^+$. Now we can proceed as in Corollary 2.3.9 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 by Lemma 1.4.1 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 2.3.5. 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 generalized square whose "vertices" are $\overline{\gamma}$, $\overline{\alpha}^+$, $\overline{\Delta}^+$, $\overline{\beta}^+$ as in Figure 7.1.

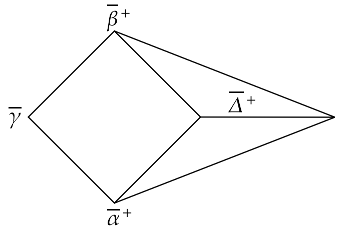

Figure 7.1: The generalized square from Lemma 7.1.8. The simplex $\overline{\Delta}^+$ is represented as a segment.

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 2.2.6. In particular we may assume 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 2.3.10. 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 2.3.4 we get that $d_{\mathcal{C}^1/DT_K}(\overline{\gamma}, \overline{\eta}) = 1$, as required. □

Now we want to define a bijection between vertices of $\mathcal{C}^{ss}/DT_K$ and slices of $\mathcal{C}^1/DT_K$. Let $\overline{\delta}$ be in $\mathcal{C}^{ss}/DT_K$. 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 7.1.9. *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 $(\mathcal{C}^{ss}/DT_K) \backslash (\mathcal{C}^1/DT_K)$. If this is the case then clearly it maps adjacent vertices to nested slices and vice versa, and the proof of Theorem 7.1.9 is complete. We break the proof into a series of lemmas.

Lemma 7.1.10. *The map S is well-defined. In other words, the underlying 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 7.1.8 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.

Lemma 7.1.11. *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 underlying 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 7.2: 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 characterized 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 7.2. Now, since we can choose $\gamma^{\pm}$ to induce different puncture separations, their projections remain at distance at least 2 by Lemma 2.3.4, hence they contradict Condition 3 for $(\overline{\alpha}^{\pm}, \overline{\beta}^{\pm})$.

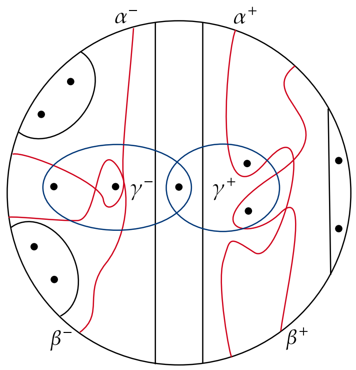

Figure 7.2: The red curves fill the subsurfaces whose boundaries are the black curves, and may cut out some punctured disks. If the annulus in between contains some punctures then the projections of the blue curves remain at distance at least 2.

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 we have that $\pi(P^{\pm}) = S(\pi(\delta))$, and in particular it is a maximal join. Hence $S(\pi(\delta)) = \overline{P}^{\pm}$ by maximality. $\square$

Lemma 7.1.12. *For all large multiples K, the vertex $\overline{\delta}$ from Lemma 7.1.11 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 thesis.

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 7.1.11, 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 7.3.
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 7.3 such that $\alpha_1^{+} = \alpha_2^{+}$ and $\beta_1^{+} = \beta_2^{+}$. If g is not the identity let (s, γ_s) be as in Proposition 2.2.1. 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 2.2.2 (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 sub-paths must intersect the star of s, because otherwise we could find a geodesic sub-path 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 curves in the respective DT_K-orbits. But now $(\alpha_1^{+}, \beta_1^{+}) = (\alpha_2^{+}, \beta_2^{+})$, and therefore $\delta_1 = \delta_2$ since they are both characterized as the only boundary component of the surface filled by $\alpha^{+} \cup \beta^{+}$ which lies in $\mathcal{C}^{ss}$. $\square$

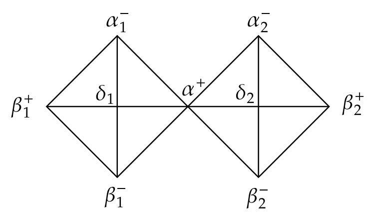

Figure 7.3: The two squares inside $\mathcal{C}^{ss}$ described in the proof of Lemma 7.1.12. We may independently choose a point for every column and find a path from β_1^{+} to β_2^{+} passing through those points. Every such path must pass through the star of s, i.e., must contain a vertex fixed by γ_s.

Now the proof of Theorem 7.1.9 is complete if we show that:

Corollary 7.1.13. *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 7.1.11 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 7.1.11 does not depend on the choice of filled divisions for $\overline{P}^{\pm}$. $\square$

7.2 From strongly separating to all curves

For this whole section we follow the footsteps of [Bow16], except we focus on the case of punctured spheres. We recall that, as in Definition 7.0.1, 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 7.2.1. *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 [Bow16], 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 7.2.2. 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 7.4.

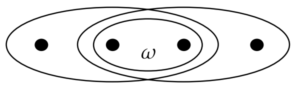

Figure 7.4: A surrounding pair and the minimal curve ω it surrounds.

Definition 7.2.3. 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 7.5.

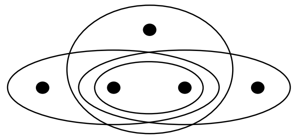

Figure 7.5: A surrounding triple.

The following is a restatement of [Bow16, Lemmas 4.2 and 4.3], which are core results of that paper:

Lemma 7.2.4. *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 recognized 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 7.2.4 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 [Bow16] 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.

7.2.1 7 punctures

We will say that an n-cycle inside a graph Γ is a closed path of length n which does not cross itself, that is, its vertices are all distinct. The following lemma, which summarizes [Bow16, Sections 3 and 4], shows that surrounding pairs and triples can be recognized inside $\mathcal{C}^{ss}(S_7)$ (which actually coincides with $\mathcal{C}^1(S_7)$) by using 7-cycles:

Lemma 7.2.5. *In a seven-punctured sphere:*

- *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.*

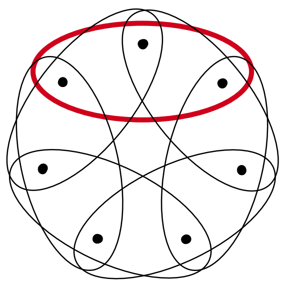

Figure 7.6: These curves form an isometrically embedded heptagon in $\mathcal{C}^1(S_7)$. The curves are all obtained from the same curve, say the red one at the top, by rotating. Notice that all curves induce different puncture separations.

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

Definition 7.2.6. 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 7.2.7. *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 7.7.

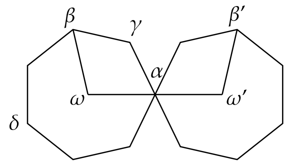

Figure 7.7: The two heptagons from the proof of Lemma 7.2.7.

Let (s, γ_s) be as in Proposition 2.2.1. 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 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 7.1.12 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 characterized 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 7.2.5 there is a unique heptagon in $\mathcal{C}^1(S_7)$ up to the action of the mapping class group, the vertices of T are seven curves that look like in Figure 7.6, and therefore induce different puncture separations. Thus by Lemma 2.3.2 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 characterize the class of pairs that surround the same $\overline{\omega}$. Again, our definition is inspired by Lemma 7.2.5.

Definition 7.2.8. 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 7.2.9. *A surrounding triple $\alpha, \beta, \gamma \in \mathcal{C}^{ss}$ projects to a surrounding triple $\overline{\alpha}, \overline{\beta}, \overline{\gamma}$ inside $\mathcal{C}^{ss}/DT_K$.*

Proof. We have already seen 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 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 2.3.4. 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 7.2.10. *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 7.2.5 and 7.2.9. □

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 7.2.11. *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 7.8.

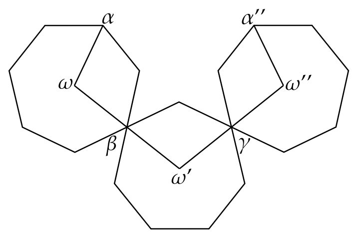

Figure 7.8: The three heptagons from the proof.

Arguing as in Lemma 7.2.7 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 $\overline{\delta}$, thus violating the fact that $\overline{\alpha}, \overline{\beta}, \overline{\gamma}$ is a surrounding triple. This implies that α, β, γ form a surrounding triple, and therefore $\omega = \omega' = \omega''$. □

Corollary 7.2.12. *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:

Lemma 7.2.13. *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 7.2.1:

Corollary 7.2.14. *If $b = 7$, for all large multiples K every automorphism of $\mathcal{C}^{ss}/DT_K$ extends to an automorphism of $\mathcal{C}/DT_K$.*

7.2.2 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 7.2.15. Let Γ be (a subgraph of) either $\mathcal{C}$ or $\mathcal{C}/DT_K$. A vertex x *splits* Γ if $\mathrm{Lk}^*_\Gamma(x)$ is connected.

Notice that we can recognize $\mathcal{C}^1$ 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 7.2.16. *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 7.2.16 is implied by the following:

Lemma 7.2.17. *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 $\overline{\gamma}_0 = \overline{\alpha}, \overline{\gamma}_1, \ldots, \overline{\gamma}_k = \overline{\beta} \in \mathrm{Lk}(\overline{\delta})$ and choose some lifts $\gamma_0 = \alpha, \gamma_1, \ldots, \gamma_k = \beta \in \mathrm{Lk}(\delta)$. Notice that, since the projection map is 1-Lipschitz, any two consecutive curves must be at distance at least 2 in $\mathcal{C}$. Hence 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 since $\delta \neq \alpha$. 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 7.9. Then $\overline{\gamma}$ must be at distance at least 2 from both $\overline{\alpha}$ and $\overline{\beta}$ by Lemma 2.3.2, and therefore $\overline{\delta}$ does not separate $\overline{\alpha}$ and $\overline{\beta}$. □

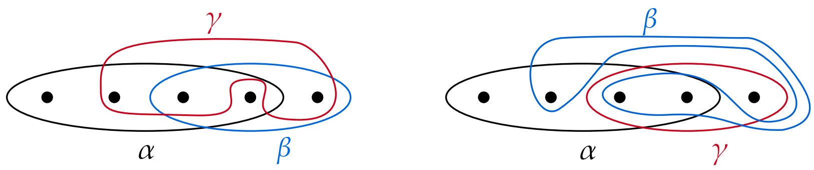

Figure 7.9: The shape of a possible curve γ depends on whether one of the discs contains the punctures of the other.

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 7.2.18 (Bizarre simplices). *Let $b \geqslant 9$. If $\Delta = (\delta_3, \ldots \delta_{b-6}) \subseteq \mathcal{C}^{ss}$ 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 7.10. 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, Lk(Δ) 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). $\square$

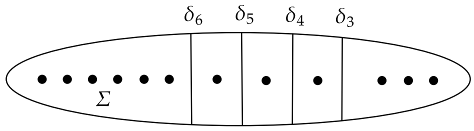

Figure 7.10: Starting from δ_3, Δ is constructed by consecutively cutting out once-punctured annuli.

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 7.2.17. Hence we get the following:

Corollary 7.2.19. *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:*

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 recognize $\mathcal{C}^{ss}(S_7)$, 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 recognize $\pi(\mathcal{C}^{ss}(S_7))$ inside $\mathcal{C}^{ss}/DT_K$ in the same way. Notice that, as discussed in Corollary 2.3.12, $\pi(\mathcal{C}^{ss}(S_7))$ 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 7.4 and 7.5. The following is a summary of [Bow16, Lemmas 6.1 and 6.3]:

Lemma 7.2.20. *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 flavor 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 generalized heptagon, thus we can find lifts Δ and $T \subset \mathrm{Lk}_{\mathcal{C}^{ss}}(\Delta)$ by Theorem 2.2.6.
- To prove the equivalent of Lemma 7.2.7 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 7.2.9 follows.
- Lemma 7.2.13 deals with disjointness, and the same statement works in our case.

We are left to prove an analogous to Lemma 7.2.11, which needs a bit more care.

Lemma 7.2.21. *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 7.2.11 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 7.11.

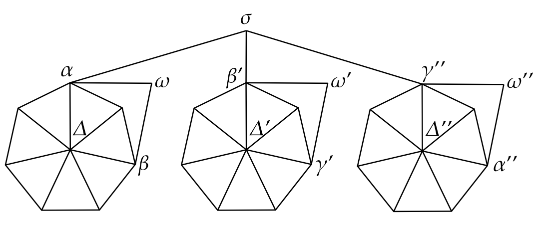

Figure 7.11: The various lifts from the proof of Lemma 7.2.21.

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 2.2.1. Arguing as in the proof of Lemma 7.1.12 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 7.2.7, 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 7.12.

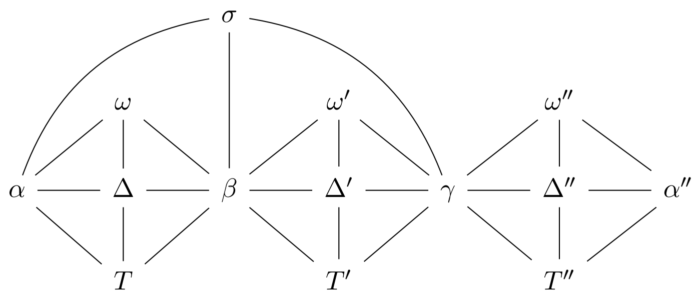

Figure 7.12: This graph represents all possible "paths" from α to α'', and elements on the same column correspond to the same cut set. Notice that every heptagon actually represents two paths.

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 2.2.1. 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 7.2.7, 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 7.2.1:

Corollary 7.2.22. *For every $b \geqslant 9$ and all large multiples K, every automorphism of $\mathcal{C}^{ss}/DT_K$ extends to an automorphism of $\mathcal{C}/DT_K$.*

7.2.3 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 recognize surrounding pairs. Define $\mathfrak{O}$ as the graph obtained by adding the four longest diagonal to an octagon, as in Figure 7.13. A possible realization of this graph inside $\mathcal{C}^1(S_8)$ is given by the eight curves in the same Figure, and [Bow16, 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).

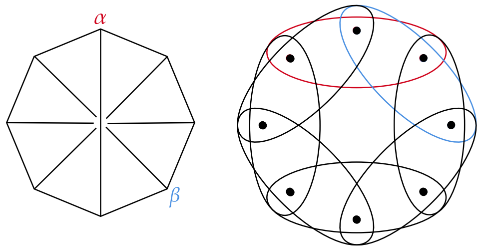

Figure 7.13: The graph $\mathfrak{O}$ and its realization with 1-separating curves, each of which surrounds three consecutive punctures. Notice that two vertices correspond to a surrounding pair (as the red curve and the blue curve) if and only if they are connected by exactly two paths of length 2 inside $\mathfrak{O}$ (as the vertices α and β).

In [Bow16] 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 $\mathcal{C}^1$;
- 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 $\delta \in \mathcal{C}^{ss} \backslash \mathcal{C}^1$ 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 7.2.23. *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 3.3.3 with X_b replaced by $\mathfrak{O}$. □

Lemma 7.2.24. *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 3.1.6). Now, it is useful to see $\mathfrak{O}$ as the 1-skeleton of a Möbius band made of four squares, as in Figure 7.14.
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 2.2.1, applied to $x = 1$. If

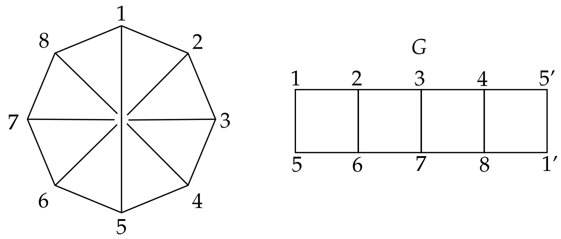

Figure 7.14: $\mathfrak{O}$ is obtained from the strip G by gluing 1 to 1′ and 5 to 5′.

$d_{\mathcal{C}}(s,1) \leqslant 1$ 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. 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 of one of the squares). 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 in 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 7.2.25. *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 7.15.

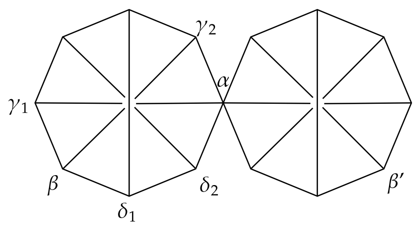

Figure 7.15: The two octagons from the proof of Lemma 7.2.25.

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 2.2.1. 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 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 7.16.
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

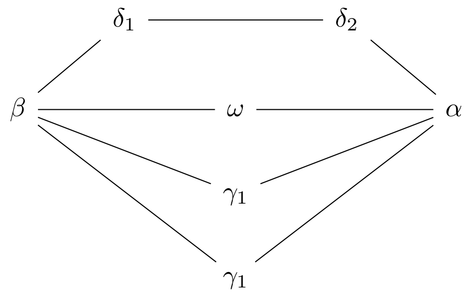

Figure 7.16: A schematic representation of $\{\omega\} \cup \mathrm{Lk}_{\mathfrak{D}}(\alpha) \cup \mathrm{Lk}_{\mathfrak{D}}(\beta)$ from Lemma 7.2.25.

γ_s to $\mathfrak{D}' \cup \mathrm{Lk}(\alpha)$. In both cases ω remains the curve surrounded by β and $\gamma_s(\alpha)$, as argued in Lemma 7.2.7; therefore we can proceed by induction. $\square$

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

- Arguing exactly as in Lemma 7.2.9 we see that surrounding triples project to surrounding triples.
- Adapting the proof of Lemma 7.2.11 with the cut set arguments from Lemma 7.2.25 we get that the vertex $\overline{\omega}$ surrounded by a surrounding triple is well-defined.
- Again, the conclusion of Lemma 7.2.13 holds, so we can recognize disjointness.

Thus we get the final piece of the proof of Theorem 7.2.1:

Corollary 7.2.26. *If $b = 8$, for all large multiples K every automorphism of $\mathcal{C}^{ss}/DT_K$ extends to an automorphism of $\mathcal{C}/DT_K$.*

Chapter 8

Quasi-isometric rigidity

This Chapter is devoted to the proof of Theorem 1, following the same steps as for the pants graph. Let us start from the HHS structure of $MCG^{\pm}/DT_K$.

Definition 8.0.1. For Δ, Δ' simplices of a graph G, we write $\Delta \sim \Delta'$ to mean $\mathrm{Lk}(\Delta) = \mathrm{Lk}(\Delta')$. 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.

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

Remark 8.0.2. As shown in [BHMS20] (see Theorem 7.1 and Proposition 8.13 there), MCG/DT_K has the following HHS structure:

(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 $\overline{U}, \overline{V} \in \mathfrak{S}$ are orthogonal (resp. nested) if $j(\overline{U})$ and $j(\overline{V})$ admit representatives, which we also call *lifts*, that are disjoint (resp. nested);

(iv) If $\overline{\Delta}$ is not a facet, then $\mathcal{C}(\overline{\Delta}) := \mathcal{C}([\overline{\Delta}])$ is quasi-isometric to $\mathrm{Lk}(\overline{\Delta})$ (see [BHMS20, Claim 6.11]). Otherwise the pointwise stabilizer 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 the notion of *convex-cocompact subgroups*, introduced by Farb and Mosher in [FM02], and the following two properties these subgroups enjoy. Recall that a group is *free* on two generators if it admits a presentation of the form $\mathcal{F}_2 = \langle a, b \rangle$.

Lemma 8.0.3 ([FM02, Theorem 1.4])**.** *For every subgroup $H \leqslant MCG^{\pm}$ of finite index there exists a copy $\mathcal{F}_2 \leqslant H$ of a free group on two generators which is convex-cocompact.*

Lemma 8.0.4 (Kent-Leiniger, [KL08])**.** *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 $x, s \subset U$ 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 8.0.5. *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 [BHMS20, Theorem 7.1.iii]. □

Now we turn to the proof of quasi-isometric rigidity. In order for our machinery to work we must ensure that we are in the assumptions of Theorem 5.2.15. We split this passage in Remark 8.0.6 and Lemma 8.0.7.

Remark 8.0.6 (Asymphoricity and normalization for HHGs). MCG/DT_K is an example of what is called a *hierarchically hyperbolic group* (HHG), that is, a group whose Cayley graph has a HHS structure which is "compatible" with the action by left multiplication. The precise definition is [BHS19, Definition 1.21], which also requires the group to act cofinitely on the set of curve graphs, such that two curve graphs in the same orbit are isometric. Hence the diameters of the curve graph of a HHG may assume finitely many different values in $[0, +\infty]$, and this implies that every HHG is asymphoric, as in Definition 5.2.9.
Also, as explained in [DHS17, Proposition 1.16] and [DHS20, Remark 2.1] we can and shall assume that the structure is normalized, up to replacing every $\mathcal{C}U$ with (a uniform neighborhood of) $\pi_U(X)$.

Lemma 8.0.7. *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, by Remark 8.0.2 and the Milnor-Şvarc Lemma 5.1.8, $\mathcal{C}(\overline{\Delta})$ 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 $\mathcal{F}_2$ 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 (see e.g [KB02, Theorem 2.28]).
Let $\Delta \subset \mathcal{C}$ be a lift of $\overline{\Delta}$. By [BHMS20, 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 $\mathcal{F}_2$ 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 Δ must be a pants decomposition for $S \backslash U$, 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 $PMCG(S_4)$.
Notice that $PMCG(S_4)$ has finite index inside $MCG(S_4)$, since it is the kernel of the action of the mapping class group on the punctures. Hence by Lemma 8.0.5 we may find a convex-cocompact copy of $\mathcal{F}_2$ inside $PMCG(S_4)$, call it Q. Suppose by contradiction that $\pi|_Q$ is not injective, which means that there is some $h \in Q \cap DT_K \backslash \{1\}$. 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 2.2.1. 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 2.2.7 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 such that $d_s(x, h(x)) \leqslant D$. Notice that,

since there are finitely many subsurfaces homeomorphic to S_4, up to the action of the mapping class group, therefore we may find a D which works for all possible U. This is a contradiction if we choose $\Theta > D$ and a large enough multiple 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 again by Remark 8.0.2 $\mathcal{C}(\overline{\Delta})$ 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 then as pointed out in Remark 1.2.11 $\mathcal{C}(U)$ 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 2.3.12.
Then 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 8.0.5: if we choose a convex-cocompact copy of $\mathcal{F}_2$ inside $MCG(U)$ then its projection is still a copy of $\mathcal{F}_2$ whose orbit maps to $\mathcal{C}(U)/DT_K(U)$ are quasi-isometric embeddings, and we are done. □

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 DT_K-classes of subsurfaces (thus omitting the bijection j whenever possible).

Lemma 8.0.8. *The following holds for all large multiples K. Let $\left\{\overline{U}_i\right\}_{i=1}^{r}$ 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 $\left\{\overline{U}_i\right\}_{i=1}^{r}$ 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 proof can be any equivalence classes of subsurfaces, not necessarily elements of a complete support set. However, the main reason we want to establish the Lemma is the following:

Corollary 8.0.9. *For every complete support set $\{\overline{U}_i\}$ for $MCG(S_b)/DT_K$ there exist representatives $\{U_i\}$ which form a complete support set for $\mathbb{P}(S_b)$.*

Proof of Lemma 8.0.8. 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:

$$\begin{tikzcd} C_1 \arrow[r, dash] \arrow[rrr, bend right, "g"'] & C_2 \arrow[r, dash] & C_3 \arrow[r] & C_1' \end{tikzcd}$$

We proceed by induction on the complexity of g. If g is the identity we are done; otherwise fix a curve $x \in C_1$ and let (s, γ_s) be as in Proposition 2.2.1. 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 2.2.2 (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 $U_1, \ldots, U_k$ 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 $C_1, \ldots, C_{k+1}, C_1'$ 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 $i = 3, \ldots, k$, as shown in this schematic picture:

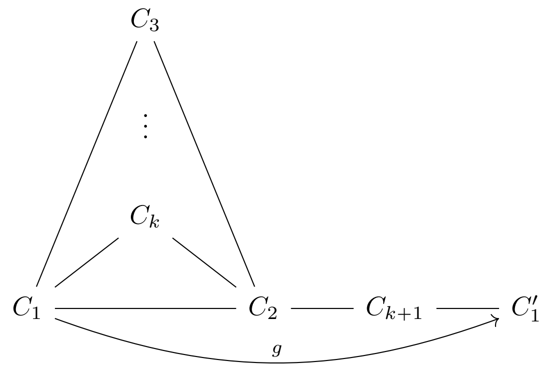

The same argument as before shows that we can glue C_1 to C_1', this time without breaking the "triangles": in every inductive step, γ_s fixes either C_1 (and we apply γ_s to the whole data) or one between C_2 and C_{k+1} (and we can apply γ_s beyond these curves, without moving the triangles). If we do this for $k = 3, \ldots, r-1$ 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. $\square$

In the previous proof 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 8.0.10. *For all large multiples K the following holds. Let $\left\{\overline{U_i}\right\}_{i=1}^r$ 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 $i = 1, \ldots, r-1$. Now let $h \in DT_K$ be an element that maps U_r to U_r', and let $C_1, \ldots, C_r, C_r'$ 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 2.2.1. 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 $C_1, \ldots, C_{r-1}$ pointwise, so we can apply γ_s to the "simplex" $\{C_1, \ldots, C_{r-1}, C_r'\}$ and proceed by induction. $\square$

From now on we can almost repeat the arguments of Section 6.1. First we define minimal and unambiguous hinges in the same way and prove the following:

Lemma 8.0.11. *For all large multiples K, a hinge $(\overline{U}, p)$ for MCG/DT_K is minimal iff it is unambiguous and its support is the class of a four-holed sphere.*

Proof. Let U be a lift of $\overline{U}$. Again, we look at all possible shapes of U, according to Lemma 6.1.5.

- Suppose U is an S_5, and choose some minimal $V \subsetneq U$ (for example, choose a pair of pants $P \subset U$ whose boundary touches a connected component of $S \backslash U$, and let $V = U \backslash P$, which is minimal by the discussion of Lemma 6.1.5). Every completion for $\overline{U}$ admits representatives that are disjoint from U by Lemma 8.0.8, 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 (this is always possible, since the complementary components of V are not pairs of pants and therefore admit pseudo-Anosov elements), and let $\{\overline{V}_i\}$ be their classes. Now recall that, as a consequence of Lemma 2.3.7, 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 ambiguous S_4. Let V be the S_5 given by the union of U and the pair of pants cut out by U, again as in Figure 6.1. Then slightly abusing notation we have that $\mathrm{Compl}(\overline{U}) = \mathrm{Compl}(\overline{V})$, because they are both the projection of $\mathrm{Compl}(U) = \mathrm{Compl}(V)$. Thus $\overline{U}$ is ambiguous, and it is not minimal since $\overline{V}$ is not.

- The only case left is when U is a minimal S_4. Choose a completion $\{U_i\}$ for U so that $U = S \backslash \bigcup U_i$, 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 by Lemma 8.0.10). Then again $V = U$ since U has minimal complexity, thus we proved that $\overline{U}$ is minimal and unambiguous.

□

Then the Products to products Corollary 6.1.7 holds also in the case of MCG/DT_K with the same proof, provided that we show the following analogue of Lemma 6.1.8:

Lemma 8.0.12. *There exists a constant C_4 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 $d_{F_{\overline{U}}}(x, \gamma) \leqslant C_4$.*

Proof. As in the proof of Lemma 6.1.8, since $\overline{U}$ is $\sqsubseteq$-minimal we just need to show that every point $x \in \mathcal{C}\overline{U}$ is within uniformly bounded distance from a bi-infinite quasigeodesic. Now, $\mathcal{C}\overline{U}$ has at least 2 points at infinity by Lemma 8.0.7, and therefore Lemma 5.1.14 grants the existence of a bi-infinite quasigeodesic γ. Moreover, setting $[\overline{\Delta}] = j^{-1}(\overline{U})$, Remark 8.0.2.(iv) tells us that the hyperbolic group $P(\overline{\Delta})$ acts cocompactly on $\mathcal{C}\overline{U}$. Hence every point $x \in \mathcal{C}\overline{U}$ lies within uniformly bounded distance from some image of γ under the action of $P(\overline{\Delta})$, which is again a bi-infinite quasigeodesic since the action is by isometries. □

The next step is to characterize terminal supports:

Lemma 8.0.13. *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, i.e., one (hence every) of its lifts U is terminal. Choose some V that cuts out U. 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 8.0.11 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 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 (again, as a consequence of Lemma 1.4.1). 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 $S\backslash V$ it belongs to, and therefore there must be a relative boundary curve of U inside Σ that we can choose as δ. Now let F be the finite union of all δs and ηs that arise from these possibilities. By Corollary 2.3.8 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 8.0.11). 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 $\{V_i\}$, and we are done. $\square$

Corollary 8.0.14. *For all large multiples K, any self-quasi-isometry f of MCG/DT_K induces an automorphism of $\mathcal{C}^1/DT_K$.*

Then this automorphism comes from some $\overline{g} \in MCG^{\pm}/DT_K$, by Theorem 7.0.3. 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 8.0.15. *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 $S\backslash U = \bigsqcup_{i=1}^{4} \Sigma_i$. 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 $i = 1, \ldots, 4$, 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 8.0.8 to lift the "simplex" $\{\overline{V}\} \cup \{\overline{\Sigma}_i\}_{i=1}^{4}$, and Lemma 8.0.10 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, which is an S_4, 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 $S\backslash\Sigma$, 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 2.3.10 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.
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:

$$\begin{tikzcd} C \arrow[r, dash] \arrow[rrrr, bend right, "g"] & \alpha \arrow[r, dash] & \Delta \arrow[r, dash] & \alpha' \arrow[r, dash] & C' \end{tikzcd}$$

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 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.
- If $d_{\mathcal{C}}(\alpha', s) \leqslant 1$ we can apply γ_s just to C', without touching α and 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 $\overline{V}$. □

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 8.0.16. *Let $S = S_{0,b}$ be a punctured sphere, with $b \geqslant 7$ punctures. For all large multiples K the following holds. For every $T > 0$ there exists a constant $D > 0$ 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 f_{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 7.2.1. Then f_{supp} and g agree on terminal supports, and therefore also on every minimal support by Lemma 8.0.15. Then the Products to products Corollary 6.1.7, which holds for MCG/DT_K as we already noticed, 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 $\overline{g}(P_{\{U_i\}})$. We are left to prove that every point $x \in MCG/DT_K$ is the (uniform) coarse intersection of two standard product regions $P \widetilde{\cap} P'$, coming from minimal complete support sets. With the same reduction as in the proof of Theorem 6.0.2 for the pants graph, it is enough to prove that there exist two complete support sets $\{\overline{U}_i\}$ and $\{\overline{V}_i\}$ with minimal, pairwise distinct supports (again, we can apply Corollary 5.2.4 since minimal supports are also $\sqsubseteq$-minimal). Let $\{U_i\}$ be a complete support set made of minimal support, whose existence we showed in the proof of Theorem 6.0.2. Choose a pseudo-Anosov mapping class ϕ such that every boundary curve of $\{U_i\}$ crosses every boundary curve of $\{V_i\} = \phi\{U_i\}$. 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 8.0.17. *For every $T > 0$ 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 this map is the identity then $\overline{g}$ can be chosen to be the identity, and the corollary follows. Now, Theorem 5.2.15 tells us that $d_{Haus}(h_{f_{hin}(\sigma)}, f(h_\sigma)) < +\infty$. Moreover $d_{Haus}(f(h_\sigma), h_\sigma) < +\infty$, because f moves every point of h_σ within uniformly bounded distance from the point itself. Hence 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 $d_{Haus}(h_\sigma, h_{\sigma'}) = +\infty$, as pointed out in Remark 5.2.13. ☐

Thus Theorem 1 is implied by the general statement below, which follows from standard arguments (see e.g. [Sch95, Section 10.4]):

Lemma 8.0.18. *Let H be a finitely generated group, and let $QI(H)$ be the group of its self-quasi-isometries up to bounded distance. Suppose that for every $T > 0$ there exists $D > 0$ 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.*

Then H is a finite extension of $QI(H)$, meaning that it surjects onto $QI(H)$ with finite kernel. 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. Fix a finite generating set S for H, and endow the group with the corresponding word metric. Set $\mu : H \to \mathrm{QI}(H)$ by mapping h to the left multiplication by h, which is clearly an isometry of H (and in particular a $(1,1)$-quasi-isometry). This map is surjective by the first hypothesis; 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 $d_H(h,1) = d_H(\mu(h)(1),1) \leqslant D$. Then we conclude since there are finitely many $h \in H$ within distance D from the identity, because Cayley graphs are proper as pointed out in Remark 5.1.6.
Regarding the second statement, let L be the finite normal subgroup 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 map $\psi : G \to \mathrm{QI}(H)$ by setting

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

We claim that this map is a group homomorphism. In fact, if $g, g' \in G$ and $h \in H$ we have that

$$d_H\left(\psi(gg')(h), \psi(g)(\psi(g')(h))\right) = d_H\left(\phi gg'\phi^{-1}(h)), \phi g\phi^{-1}\phi g'\phi^{-1}(h))\right).$$

Now, since ϕ is a quasi-isometry for some constants (C_1, C_1) we have that

$$d_H\left(\phi gg'\phi^{-1}(h)), \phi g\phi^{-1}\phi g'\phi^{-1}(h))\right) \leqslant C_1 d_G\left(gg'\phi^{-1}(h)), g\phi^{-1}\phi g'\phi^{-1}(h))\right) + C_1,$$

and since the left multiplication by g is an isometry

$$C_1 d_G\left(gg'\phi^{-1}(h)), g\phi^{-1}\phi g'\phi^{-1}(h))\right) + C_1 = C_1 d_G\left(g'\phi^{-1}(h)), \phi^{-1}\phi g'\phi^{-1}(h))\right) + C_1.$$

Now ϕ^{-1} is a quasi-inverse for ϕ, which means that there exists a constant C_2 such that

$$d_G\left(g'\phi^{-1}(h)), \phi^{-1}\phi g'\phi^{-1}(h))\right) \leqslant C_2.$$

Hence

$$d_H\left(\psi(gg')(h), \psi(g)(\psi(g')(h))\right) \leqslant C_1(C_2 + 1),$$

which means that $\psi(gg')$ and $\psi(g) \circ \psi(g')$ coincide up to bounded distance, and therefore are the same element of $\mathrm{QI}(H)$.
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 left multiplication is D-close to $\psi(g)$. Then by construction the following diagram commutes:

$$\begin{tikzcd} & H \arrow[d, "\mu"] \\ G \arrow[ur, "\theta"] \arrow[r, "\psi"] & \mathrm{QI}(H) \end{tikzcd}$$

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. $\square$

Chapter 9

Algebraic rigidity

This Chapter is devoted to the proof of Theorem 3, which is covered by Theorems 9.0.1 and 9.0.8 below. Recall that if G is a group, we denote by $\mathrm{Aut}(G)$ the group of group isomorphisms from G to itself. Among these are the inner automorphisms $\mathrm{Inn}(G)$, which are just the conjugation by elements of G. Then one might be interested in the quotient $\mathrm{Out}(G) := \mathrm{Aut}(G)/\mathrm{Inn}(G)$, which we call the group of *outer* automorphisms. In [Iva97] Ivanov showed that, except in some sporadic cases of low complexity, the homomorphism $MCG(S) \to \mathrm{Aut}(MCG^{\pm}(S))$ which sends every element to the corresponding conjugation is an isomorphism, and in particular $\mathrm{Out}(MCG^{\pm}(S))$ is trivial. Our goal is to prove the following analogue for the quotient:

Theorem 9.0.1. *Let S_b be a sphere with $b \geqslant 7$ punctures. For all large multiples K, every $\phi : H \to H'$ isomorphism between finite index subgroups of $MCG^{\pm}(S_b)/DT_K$ is the restriction of an inner automorphism. In particular $Out(MCG^{\pm}(S_b)/DT_K)$ is trivial.*

We recall a definition from [AMS16]:

Definition 9.0.2. Two elements h and g of a group G are *commensurable*, and we write $h \overset{G}{\approx} g$, if there exist $m, n \in \mathbb{Z}\backslash\{0\}$, $k \in G$ such that $kg^m k^{-1} = h^n$ (that is, if they have non-trivial conjugate powers).

Moreover recall that a subgroup $E \leqslant G$ is *normalized* by another subgroup H if for every $h \in H$ and every $g \in E$ we have that $hgh^{-1} \in E$. The following result is a special case of [AMS16, Theorem 7.1]. Roughly speaking, the theorem says that if H is a non-virtually-cyclic subgroup of G and both act with loxodromic WPD elements 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 9.0.3. *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 normalized by H, whose existence is proven in [AMS16, 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.*

Outline of the proof of Theorem 9.0.1. We just need to verify that, for all $b \geqslant 7$ and for all large multiples K, the hypotheses of Theorem 9.0.3 are satisfied for $G = MCG^{\pm}(S_b)/DT_K$ and any isomorphism $\phi : H \to H'$ between subgroups of finite index. By results of Dahmani, Hagen and Sisto [DHS21, Theorems 2.1 and 5.2] the quotient $\mathcal{C}/DT_K$ is a hyperbolic graph, G is non-virtually-cyclic, and the action of G on $\mathcal{C}/DT_K$ admits loxodromic WPD elements. In particular, every subgroup of finite index is not virtually-cyclic (otherwise G itself would be so) and it contains some power of every loxodromic WPD element, which remains loxodromic WPD. In Lemma 9.0.4 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 9.0.5, in order to prove that $E_G(H) = \{1\}$ it suffices to show that $E_G(G) = \{1\}$, i.e., that $MCG^{\pm}/DT_K$ has no nontrivial finite normal subgroups. This is done in Lemma 9.0.7. $\square$

Lemma 9.0.4. *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 any of the closest elements of H). Let $\Phi = \phi \circ f : G \to H'$, which coincides with ϕ on H and is a self-quasi-isometry of G because ϕ is an isomorphism between finite index subgroups (this is an easy consequence of Lemma 5.1.9). Then by Theorem 8.0.16 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 to cancel gk in the first distance. 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 [Hru10, 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-neighborhood 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, otherwise $N_{D'}(H_1 \cap H_2)$ would be a finite union of balls of radius D', and as pointed out in Remark 5.1.6 balls in Cayley graphs contain finitely many elements. Therefore there exist common powers $gh^m g^{-1} = \phi(h)^n$, as required. $\square$

Lemma 9.0.5. *Let S be a hyperbolic space, on which a non-virtually-cyclic group G acts coboundedly and with 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 $g \in G$ let $E_G(g)$ be the *virtual commutator* of the subgroup generated by g, that is,

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

It follows from the proof of [DGO17, 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 [AMS16, 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

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

Since H has finite index there exists $k \in \mathbb{N}_{>0}$ 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 9.0.6. *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 nontrivial finite normal subgroups. In other words $E_{MCG^{\pm}}(MCG^{\pm}) = \{1\}$.*

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 [Iva97, Theorem 1] and its extension to lower genera by Korkmaz [Kor99, 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 (see Lemma 1.1.3), 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

$$T_\gamma^{-2} T_{f(\gamma)}^{\pm 2} = g f^{-1}. \tag{9.1}$$

Now, referring to the table at the end of [FM12, 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 $\mathcal{F}_2$, hence the left-hand side of Equation (9.1) has infinite order. This is impossible, since the right-hand side is an element of the finite subgroup N. □

Lemma 9.0.7. *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 $MCG^{\pm}$. Recall that, as stated in Theorem 5.1.22, $MCG^{\pm}$ is acylindrically hyperbolic, hence by [DGO17, 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 2.3.9 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 $\{g^n(x)\}_{n \in \mathbb{Z}}$ by Lemma 2.3.7. 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 $n \in \mathbb{Z}$ to be determined later, 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 2.3.7. 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 2.2.6 there exists a lift Q of $\overline{Q}$, and by Lemma 2.2.9 we can assume that the lifts l', s' of $\overline{l}, \overline{s}$ are DT_K-translates of l, s respectively. Up to elements of DT_K we can also 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 [DGO17, Lemma 6.7] which says, roughly, that coarsely stabilizing a large segment of an axis is equivalent to stabilizing the whole axis (as noted in [DGO17], the lemma has the same proof as [BF02, Proposition 6], which has more restrictive hypotheses). Thus $\psi = g^m$ for some $m \in \mathbb{Z}$, and in the quotient $\overline{\phi} = \overline{\psi} = \overline{g}^m$. Then we must have that $m = 0$ since $\overline{\phi}$ has finite order while $\overline{g}$ has infinite order. □

As a consequence of Theorem 9.0.1 we can also describe the automorphism group of MCG/DT_K:

Theorem 9.0.8. *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 9.0.9. *For every $b \geqslant 7$ and for all large multiples K, $MCG(S_b)/DT_K$ has trivial center.*

Proof. Let $G = MCG^{\pm}(S_b)/DT_K$ and $H = MCG(S_b)/DT_K$. The center $Z(H)$ 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 9.0.5. □

Proof of Theorem 9.0.8. Again, let $G = MCG^{\pm}/DT_K$ and $H = MCG/DT_K$. Theorem 9.0.1 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 center by Lemma 9.0.9. 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 9.0.7. This proves that $G \cong \mathrm{Aut}(H)$, and since H has trivial center we also get

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

as required. □

Bibliography

[AL13] Javier Aramayona and Christopher J. Leininger. Finite rigid sets in curve complexes. *J. Topol. Anal.*, 5(2):183–203, 2013.

[AMS16] Yago Antolín, Ashot Minasyan, and Alessandro Sisto. Commensurating endomorphisms of acylindrically hyperbolic groups and applications. *Groups Geom. Dyn.*, 10(4):1149–1210, 2016.

[Art47] E. Artin. Theory of braids. *Ann. of Math. (2)*, 48:101–126, 1947.

[BE88] B. H. Bowditch and D. B. A. Epstein. Natural triangulations associated to a surface. *Topology*, 27(1):91–117, 1988.

[Beh06] Jason A. Behrstock. Asymptotic geometry of the mapping class group and Teichmüller space. *Geom. Topol.*, 10:1523–1578, 2006.

[BF02] Mladen Bestvina and Koji Fujiwara. Bounded cohomology of subgroups of mapping class groups. *Geom. Topol.*, 6:69–89, 2002.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[BHMS20] Jason Behrstock, Mark F. Hagen, Alexandre Martin, and Alessandro Sisto. A combinatorial take on hierarchical hyperbolicity and applications to quotients of mapping class groups, 2020.

[BHS17] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups. *Geom. Topol.*, 21(3):1731–1804, 2017.

[BHS19] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces II: Combination theorems and the distance formula. *Pacific J. Math.*, 299(2):257–338, 2019.

[BHS21] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *Duke Math. J.*, 170(5):909–996, 2021.

[BKMM12] Jason Behrstock, Bruce Kleiner, Yair Minsky, and Lee Mosher. Geometry and rigidity of mapping class groups. *Geom. Topol.*, 16(2):781–888, 2012.

[Bow16] Brian H. Bowditch. Rigidity of the strongly separating curve graph. *Michigan Math. J.*, 65(4):813–832, 2016.

[Bow20] Brian H. Bowditch. Large-scale rank and rigidity of the Weil-Petersson metric. *Groups Geom. Dyn.*, 14(2):607–652, 2020.

[Bro03] Jeffrey F. Brock. The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores. *J. Amer. Math. Soc.*, 16(3):495–535, 2003.

[Dah18] François Dahmani. The normal closure of big Dehn twists and plate spinning with rotating families. *Geom. Topol.*, 22(7):4113–4144, 2018.

[Deh38] M. Dehn. Die Gruppe der Abbildungsklassen. *Acta Math.*, 69(1):135–206, 1938. Das arithmetische Feld auf Flächen.

[Deh87] Max Dehn. *Papers on group theory and topology.* Springer-Verlag, New York, 1987. Translated from the German and with introductions and an appendix by John Stillwell, With an appendix by Otto Schreier.

[DGO17] F. Dahmani, V. Guirardel, and D. Osin. Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *Mem. Amer. Math. Soc.*, 245(1156):v+152, 2017.

[DHS17] Matthew Gentry Durham, Mark F. Hagen, and Alessandro Sisto. Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.*, 21(6):3659–3758, 2017.

[DHS20] Matthew Gentry Durham, Mark F. Hagen, and Alessandro Sisto. Correction to the article Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.*, 24(2):1051–1073, 2020.

[DHS21] François Dahmani, Mark F. Hagen, and Alessandro Sisto. Dehn filling Dehn twists. *Proc. Roy. Soc. Edinburgh Sect. A*, 151(1):28–51, 2021.

[FM02] Benson Farb and Lee Mosher. Convex cocompact subgroups of mapping class groups. *Geom. Topol.*, 6:91–152, 2002.

[FM12] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series.* Princeton University Press, Princeton, NJ, 2012.

[Fun99] Louis Funar. On the TQFT representations of the mapping class groups. *Pacific J. Math.*, 188(2):251–274, 1999.

[Gro87] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[Ham07] U Hamenstädt. Geometry of the mapping class group iii: Quasi-isometric rigidity. preprint. *arXiv preprint math.GT/051242*, 2007.

[Har81] W. J. Harvey. Boundary structure of the modular group. In *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978)*, volume 97 of *Ann. of Math. Stud.*, pages 245–251. Princeton Univ. Press, Princeton, N.J., 1981.

[Hru10] G. Christopher Hruska. Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.*, 10(3):1807–1856, 2010.

[HT80] A. Hatcher and W. Thurston. A presentation for the mapping class group of a closed orientable surface. *Topology*, 19(3):221–237, 1980.

[HT85] Michael Handel and William P Thurston. New proofs of some results of Nielsen. *Advances in Mathematics*, 56(2):173–191, 1985.

[Iva97] Nikolai V. Ivanov. Automorphisms of complexes of curves and of Teichmüller spaces. In *Progress in knot theory and related topics*, volume 56 of *Travaux en Cours*, pages 113–120. Hermann, Paris, 1997.

[Jor91] Camille Jordan. *Cours d'analyse de l'École polytechnique. Tome III.* Les Grands Classiques Gauthier-Villars. [Gauthier-Villars Great Classics]. Éditions Jacques Gabay, Sceaux, 1991. Équations différentielles. [Differential equations], Reprint of the third (1915) edition.

[KB02] Ilya Kapovich and Nadia Benakli. Boundaries of hyperbolic groups. In *Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001)*, volume 296 of *Contemp. Math.*, pages 39–93. Amer. Math. Soc., Providence, RI, 2002.

[KL08] Richard P. Kent, IV and Christopher J. Leininger. Shadows of mapping class groups: capturing convex cocompactness. *Geom. Funct. Anal.*, 18(4):1270–1325, 2008.

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

[Lic64] W. B. R. Lickorish. A finite set of generators for the homeotopy group of a 2-manifold. *Proc. Cambridge Philos. Soc.*, 60:769–778, 1964.

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

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

[Min96] Yair N. Minsky. A geometric approach to the complex of curves on a surface. In *Topology and Teichmüller spaces (Katinkulta, 1995)*, pages 149–158. World Sci. Publ., River Edge, NJ, 1996.

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

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

[MS22] Giorgio Mangioni and Alessandro Sisto. Rigidity of mapping class groups mod powers of twists, 2022.

[Nie44] Jakob Nielsen. Surface transformation classes of algebraically finite type. *Danske Vid. Selsk. Mat.-Fys. Medd.*, 21(2):89, 1944.

[Osi16] D. Osin. Acylindrically hyperbolic groups. *Trans. Amer. Math. Soc.*, 368(2):851–888, 2016.

[Sch95] Richard Evan Schwartz. The quasi-isometry classification of rank one lattices. *Inst. Hautes Études Sci. Publ. Math.*, 82:133–168 (1996), 1995.

[Ser85] Caroline Series. The modular surface and continued fractions. *J. London Math. Soc. (2)*, 31(1):69–80, 1985.

[Tho92] Carsten Thomassen. The Jordan-Schönflies theorem and the classification of surfaces. *Amer. Math. Monthly*, 99(2):116–130, 1992.

[Thu97] William P. Thurston. *Three-dimensional geometry and topology. Vol. 1*, volume 35 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1997. Edited by Silvio Levy.

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