

A Thesis Submitted for the Degree of PhD at the University of Warwick

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/174030>

Copyright and reuse:

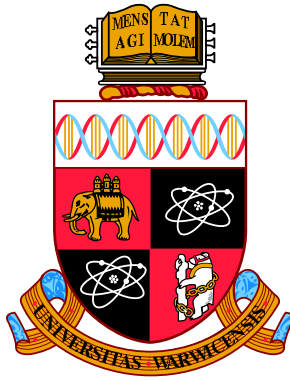
This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it.

Our policy information is available from the repository home page.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk



RIGIDITY PROPERTIES OF GRAPHS ASSOCIATED WITH PLANAR SURFACES

Marco Barberis

Thesis submitted for the degree of
Doctor of Philosophy in Mathematics

UNIVERSITY OF WARWICK

Mathematical Institute

February 2022

Contents

Acknowledgements	iii
Declarations	iv
Index of Notation	vi
Introduction	viii
1 Main Results and Background Material	1
1.1 Basic Objects and Conventions	1
1.1.1 Categories	1
1.1.2 Surfaces and Subsurfaces	2
1.1.3 Homotopy and Isotopy	3
1.1.4 Curves, Arcs and Regions	4
1.1.5 Mapping Class Groups	7
1.1.6 Graphs	8
1.1.7 Symmetric Groups	9
1.2 Realisation of Curves and Regions	10
1.2.1 Transversality and General Position	10
1.2.2 Geodesic Realisation and Representatives of Regions	10
1.2.3 Alexander’s Method	13
1.3 Graphs of Curves	15
1.3.1 The Classical Curve Graph	15
1.3.2 Subgraphs of the Curve Graph	16
1.4 Graphs of Discs and Regions	19
1.4.1 Graphs of k -Separating Discs	19
1.4.2 Graphs of Regions	20
1.5 Kneser Graphs	24
1.5.1 Standard Kneser Graphs	24
1.5.2 Extended Kneser Graphs	24
1.6 Rigid Subgraphs and Exhaustions	26
2 Rigidity of Graphs of Regions	28
2.1 Outline of the Chapter	28
2.2 Topological Properties and Alternating Hexagons	32
2.3 Connectedness of the Graphs	37

2.4	Projections to Extended Kneser Graphs	45
2.5	Nonstandard Regular Hexagons	50
2.6	Irregular Alternating Hexagons	67
2.7	Disc Graphs and Curve Graphs	79
2.8	Rigidity of Graphs of Discs	85
2.9	Rigidity of Graphs of Regions	94
3	Exhaustions by Finite Rigid Sets	111
3.1	Outline of the Chapter	111
3.2	Graph-Theoretical Machinery	113
3.3	One-Third Dehn Twists	116
3.4	Exhaustion for $\mathcal{C}_{ss}(S_{0,7})$	119
3.5	Exhaustion for $\mathcal{C}_{ss}(S_{0,8})$	136
3.6	Examples of non-Exhaustable Graphs	161
	Bibliography	165

Acknowledgements

I would like to thank the Engineering and Physical Sciences Research Council for funding my doctoral scholarship, and to the Mathematics Institute of Warwick for hosting me during this doctorate.

My supervisors, Brian Bowditch and Saul Schleimer, deserve special thanks, as they have been both leading and inspirational figures throughout this path. Brian, for having been invaluable in suggesting new possible results, and in guiding me towards them, while also being a tireless reader of my work, always providing the most useful suggestions to make it better. Saul, for having always been there to listen to my ideas and the invaluable help in brainstorming, and also for constantly having contagious passion and enthusiasms for every kind of mathematics. Moreover, thanks to both for having always supported me, and for the encouragement that let me surpass any hardship I encountered on my way.

I am thankful to all the people at the Mathematics Institute of Warwick, alongside all the people I met at conferences and seminars during these years, for the most profitable exchanges of ideas.

I would also like to thank my family, who has always believed in me, and all my friends, both those with whom I shared part of my path here in Warwick, and those who have always been at my side, even from afar.

Words cannot tell how thankful I am to you, Simona, who everyday push me to be a better person, and at the same time give me the strength to achieve any goal.

Declarations

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy.

I hereby declare that the work of this thesis is my own, except where explicitly indicated in the text.

No part of the thesis has been submitted by me for any other degree.

Abstract

We prove two kinds of results on rigidity of graphs arising from punctured spheres.

First, we prove that the graphs whose vertices are discs with a suitably bounded number of marked points, and whose edges are given by disjointness, are rigid, that is, every graph automorphism is topologically induced by an extended mapping class. We will also extend this rigidity result to subgraphs of the curve graph, with similar bounds on marked points enclosed by each curve, obtaining a generalisation of Bowditch's rigidity theorem for the strongly separating curve graph. Moreover, we will provide a complete topological classification of the rigid graphs of regions, which are graphs of isotopy classes of subsurfaces, sharpening a theorem of McLeay. Thus, our work verifies another case of Ivanov's Metaconjecture, which states that sufficiently rich objects naturally associated with surfaces have the extended mapping class group as the group of their automorphisms.

The second group of results concerns the existence of exhaustions by finite rigid subgraphs, that is, such that every embedding into the ambient graph is induced by a global graph automorphism. We will study the case of the strongly separating curve graphs of both the seven-holed sphere and the eight-holed sphere, that is the graph whose vertices having at least three punctures in each complementary component, with edges given by disjointness. These are inspired by related work of Aramayona-Leininger on exhaustions of the regular curve graph. It follows that our graphs have the co-Hopfian property, that is every self-embedding of the entire graph is actually an automorphism.

Index of Notation

- $K^*(m, k)$ Extended Kneser graph, page 24
- $M(A, B)$, page 48
- $P(A, B)$, page 48
- $\Delta(A, B)$ Common neighbouring disc, page 43
- $l(A, B)$, page 49
- $p(A, B)$, page 48
- $s(A, B)$, page 49
- $t(B, A)$, page 49
- $\mathcal{A}(\Sigma_m)$ $Mod^\pm(\Sigma_m)$ -equivariant collection of regions, page 20
- $\mathcal{A}(S, p)$ Loop graph, page 159
- $\mathcal{A}_D, \mathcal{A}_{D^c}, \mathcal{A}_{\partial C}$ Components of $\lambda(\{D, D^c\})$, page 92
- $\mathcal{C}(S)$ Curve graph, page 15
- $\mathcal{C}_k(\Sigma_m)$ k -separating curve graph, page 17
- $\mathcal{C}_{(k)}(\Sigma_m)$ Strict k -separating curve graph, page 17
- $\mathcal{C}_{ss}(S)$ Strongly separating curve graph, page 16
- $\mathcal{D}_{(k)}(\Sigma_m)$ Graph of k -punctured discs, page 19
- $\mathcal{D}_k(\Sigma_m)$ k -separating disc graph, page 19
- $\mathcal{GPC}_k(\Sigma_m)$ Graph of pairs of complementary discs, page 72
- $\mathcal{G}_\mathcal{A}(\Sigma_m)$ Graph of regions, page 20
- \mathcal{H} Heptagon in $\mathcal{C}_{ss}(\Sigma_7)$, page 117
- $\mathcal{H}(\Sigma_m)$ Set of pairs surrounding O , page 77
- $\mathcal{H}(\Sigma_m)$ Set of surrounding pairs, page 77
- $\mathcal{MPJ}_\mathcal{A}(\Sigma_m)$ Set of maximal perfect joins, page 92

$\mathcal{N}(Y, r)$ Regular neighbourhood, page 11
 \mathcal{O} Octagon with diagonals in $\mathcal{C}_{ss}(\Sigma_8)$, page 134
 \mathcal{O}^* , page 138
 $\mathcal{PC}_k(\Sigma_m)$ Set of pairs of complementary discs, page 71
 $\mathcal{R}(S)$ Set of regions on a surface, page 6
 $\mathcal{S}(X)$ Symmetric group, page 8
 \mathcal{S}_m Symmetric group over a set of m elements, page 8
 \mathcal{S}_∞ Symmetric group over a countable set, page 8
 $\mathcal{V}(\Gamma)$ Vertices of a graph, page 7
 $\text{Mod}(S)$ Mapping class group, page 6
 $\text{Mod}^\pm(S)$ Extended mapping class group, page 6
 $\nu(R)$ Complexity of a region, page 21
 $\nu_{\mathcal{A}}$ Complexity of a graph of region, page 21
 Π Marked points on a surface, page 2
 Σ_p p -punctured sphere, page 2
 $\xi(S)$ Complexity of a surface, page 2
 $i(A, B)$ Geometric intersection number, page 5
 $K(m, k)$ Kneser graph, page 24
 N_O , page 78
 $S_{g,p}^b$ Finite-type surface, page 2
 $V * W$ Graph join, page 91
 W_j^ε , page 142
 Z_j^ε , page 122

Introduction

Since 1981, when first introduced by Harvey in [Harv1], the curve complex of a surface has been a major object of interest for low-dimensional topologists and geometric group theorists. The curve complex on a surface has the (isotopy classes of) curves on that surface as vertices, and a k -dimensional simplex for each set of $k - 1$ pairwise disjoint curves. In many applications, including this thesis, the study of this complex is equivalent to the study of its 1-skeleton, known as the curve graph. The key property of this object comes from the fact that, although apparently encoding a meager amount of information, it proves to be a complete invariant for the topological type of the surface (apart from some sporadic cases), hence revealing to contain much more data than apparent. This is the reason why the curve complex, and many of its variants, find notable application throughout many aspects of low-dimensional topology and geometric group theory.

Among the many applications of the curve complex we recall some of the most notable ones. Harer, in [Hare1] and [Hare2], studied the homological type of curve complexes, and used that information to compute the cohomology of mapping class groups. The mapping class group is another fundamental object of study related to surfaces. It consists of the group of self-homeomorphisms of the surface, mod out by isotopy: this ensures that this group does not contain the wasteful information provided by the group of homeomorphisms, which is too big, while for finite-type surface is indeed finally presented. At the same time, the mapping class group retain all the relevant information about the surface, since it is a complete invariant for its topological type (apart from some sporadic cases). Since the mapping class group naturally acts on the curve complex, the two objects are deeply linked.

Masur and Minsky, in their pioneering works [MM1] and [MM2], continued using the curve complex in order to study the mapping class group, this time from a metrical viewpoint. This machinery was eventually generalised and led to the notion of Hierarchically Hyperbolic Spaces, as defined in [BHS] by Behrstock, Hagen, and Sisto. The curve complex has also been used in the study of 3-manifolds and Klenian groups: it is, for instance, a key element in the proof of the Ending Lamination Theorem, as first proven by Brock, Canary, and Minsky in [BCM].

In time, many subcomplexes of the curve complex have been studied, as well as many other complexes naturally associated to surfaces. Among the vast array of results related to such complexes, in this thesis we will focus on aspects of rigidity for planar surfaces. Our results can be divided into two main areas. The first group of results concerns combinatorial rigidity of three classes of graphs, that is, the fact that every automorphism of the graph is induced topologically by a mapping class. This property

expresses the fact that, apart from sporadic cases, the curve complex contains all the relevant information regarding the surface, and is a complete invariant for its topological type. The first rigidity result we prove is for the k -separating curve graph, that is, the subgraph of the curve graph induced by the vertices which represent the curves that have at least k marked points in either of their complementary component; see Definition 1.3.6.

Theorem A. *Let $k \geq 2$ and $m \geq 2k + 1$. Let $S_{0,m}$ be the m -punctured sphere. Then the k -separating curve graph $\mathcal{C}_k(S_{0,m})$ is rigid.*

Among the possible variation on the curve complex and graph, is it possible to use (isotopy classes of) subsurfaces as vertices, with adjacency still induced by disjointness. Indeed, the k -separating disc graph is defined as the graph with discs with a number of marked points between k and $m - k$ as vertices, with an edge joining vertices represented by disjoint discs; see Definition 1.4.1. We prove the rigidity of this graph, which is the following result.

Theorem B. *Let $k \geq 2$ and $m \geq 2k + 1$. Let $S_{0,m}$ be the m -punctured sphere. Then the k -separating disc graph $\mathcal{D}_{(k)}(S_{0,m})$ is rigid.*

The last result of this group is the most general and constitutes a complete classification of rigid graphs of region on punctured spheres. In the definition of a graph of regions vertices are a collection of isotopy classes of subsurfaces, which is stable under the action on the mapping class group, so that a natural action is defined. Once again, an edge is present between two regions if they are disjoint. For all the relevant definitions we refer to §1.4.2. We give a complete classification of all the rigid graphs of regions on planar surfaces. The requirements for rigidity are a bound on the minimal complexity of any region in the collection, where the complexity of a region is the minimum number of marked points that a disc containing it must have. Moreover, the graph of regions must not have pathological types of vertices, called vertices with holes and cork pairs: these vertices have the property of generating automorphisms of the graph which only exchange two vertices, while fixing all the other, any such automorphisms cannot be induced topologically. The aforementioned classification is the following.

Theorem C. *Let $m \geq 5$. Let $S_{0,m}$ be the m -punctured sphere. Let $\mathcal{A} \subseteq \mathcal{R}(S_{0,m})$ be a $\text{Mod}^\pm(S_{0,m})$ -invariant collection of regions, and let $\mathcal{G}_{\mathcal{A}}(S_{0,m})$ be the associated graph of regions. Then the graph is rigid if and only if $m \geq 2\nu_{\mathcal{A}} + 1$, and the graph has no vertices with holes and no cork pairs.*

The second group of results concerns the existence of exhaustions of the strongly separating curve graph on the 7-holed sphere $S_{0,7}$ and the 8-holed sphere $S_{0,8}$ by finite rigid sets, as in Definition 1.6.1, with trivial pointwise stabilisers. Strongly separating curve graph is the original name for the 3-separating curve graph, that is the subgraph of the curve graph induced by vertices representing curves with at least 3 marked points in each of their complementary components. In this context a rigid subgraph is such that every embedding into the ambient graph is induced by a global automorphism. The first result we prove is the following.

Theorem D. *Let $S = S_{0,7}$ or $S = S_{0,8}$. Then the strongly separating curve graph $\mathcal{C}_{ss}(S)$*

admits an exhaustion by finite rigid sets

$$\bigcup_{i \in \mathbb{N}} X_i = \mathcal{C}_{ss}(S).$$

Moreover, every subgraph X_i has trivial pointwise stabiliser.

As a corollary of the previous result we obtain that the aforementioned graphs satisfy the co-Hopfian property.

Theorem E. *Let $S = S_{0,7}$ or $S = S_{0,8}$. Then, for every injective graph homomorphism $i : \mathcal{C}_{ss}(S) \hookrightarrow \mathcal{C}_{ss}(S)$, there exists an extended mapping class $f \in \text{Mod}^\pm(S)$ such that i is induced by f .*

As for a historical note, the first rigidity result for the curve graph was proven by Ivanov in [Iv1], whose proof was completed by Korkmaz in [Kor]. The result has also been proven by Luo in [Lu] using different techniques. In the years that followed many rigidity results for various subgraphs of the curve graph have been found. Among these, Irmak in [Ir] proved rigidity of the graph of nonseparating curves, while Brendle and Margalit in [BM1] proved the rigidity of the graph of separating curves, extending a partial result due to Kida in [Ki]; lastly, the rigidity of the strongly separating curve graph was proven by Bowditch in [B2].

There are also rigidity results regarding graph whose vertices are still objects related to surfaces, but not necessarily curves. In this vein, for example, McCarthy and Papadopoulos proved the rigidity of the truncated graph of domains in [MP], Irmak and McCarty proved rigidity of the arc graph in [IrM], a result which has also independently been proven by Disarlo in [Di]; similarly the arc and curve complex has been proven to be rigid by Korkmaz and Papadopoulos in [KP].

All the results we have mentioned so far concern finite-type surfaces, but many results have also been proven for infinite-type surfaces. Among others, Hernández Hernández, Morales, and Valdez have proven that the curve graph is rigid in [HMV], while the rigidity of both the loop graph and the arc graph has been proven by Schaffer-Cohen in [S-C].

The previously mentioned results also have deep consequences for the computation of automorphisms of other objects naturally associated to surfaces. Just to name a few, the rigidity of the curve graph was already been used by Ivanov in [Iv1] to prove rigidity of the Teichmüller space, and Bowditch in [B3] proved the quasi isometric rigidity of the Weil-Petersson space reducing to the computation of automorphisms of the strongly separating curve graph.

The abundance of such positive rigidity results has led Ivanov to state the following metaconjecture in [Iv2].

Conjecture (Ivanov’s Metaconjecture). *Every object naturally associated to a surface S and having a sufficiently rich structure (that is apart from some low-complexity cases) has the mapping class group $\text{Mod}^\pm(S)$ as its group of automorphisms. Moreover, this can be proven by a reduction to the theorem about the automorphisms of the curve graph.*

With the previous metaconjecture in mind, many authors developed results applicable to wide classes of objects related to surfaces. For example, rigidity for graph of

regions has been dealt with by Brendle and Margalit for closed surfaces in [BM2] and by McLeay for punctured spheres in [Mc1] and for general punctured surfaces in [Mc2]. The vertices of such graphs are isotopy classes of subsurfaces, and edges between them are given by disjointness. All the aforementioned results bring as a corollary many interesting rigidity results for normal subgroups of the mapping class groups, proving that, under mild hypotheses, every group automorphism and abstract commensurator is induced topologically. However, these results are not sharp and do not provide a complete classification of the rigid graphs. Our Theorem *D* will instead be a complete classification of the rigid graphs of region on punctured spheres. The proof of the aforementioned result will pivot around the reduction to the computation of automorphisms of graphs of discs or curves, hence, ultimately, to the application of Ivanov's Theorem, as suggested by the metaconjecture.

In the latter part of the thesis we will instead focus on rigid subgraphs of the strongly separating curve graphs for the 7-holed sphere and the 8-holed sphere. Such graphs have for vertices isotopy classes of curves which do not bound a pair of pants on either side, and edges between two of them if they are disjoint. A rigid subgraph is a subgraph such that each of its embedding into the ambient graph is induced by an automorphism of the latter graph. The existence of finite rigid subgraphs is a problem which is interesting per se. Just to name a few results the problem was solved for the curve graph by Aramayona and Leininger in [AL1], whereas Hernández Hernández, Leininger, and Maungchang solved it for the pants graph in [HLM], and Shinkle did for the arc graph in [Shi]. For the strongly separating curve graph of the 7-holed sphere a finite rigid subgraph is provided in [B2].

Alongside the existence of finite rigid subgraphs, another interesting problem is the existence of exhaustions of the graph by finite rigid sets. This was proven, for the curve graph, once again by Aramayona and Leininger in [AL2], and, through a different method, by Hernández Hernández in [He]. In this thesis we show the existence of an exhaustion by finite rigid sets for the strongly separating curve graph of the 7-holed and the 8-holed sphere (Theorem *E*).

The existence of such exhaustions brings other rigidity results as a consequence. Indeed, Aramayona and Leininger in [AL2] gave a proof of the co-Hopfian property of the curve graph, which was already proven by Shackleton in [Sha], that is every injection of the graph into itself is induced by a graph automorphism. From the existence of an exhaustion by finite rigid sets we will deduce that the analogous result holds for the strongly separating curve graph of the 7-holed and the 8-holed sphere (Theorem *F*). Moreover, we will use this fact to present some examples of graphs that, although they are rigid and metrically nice, do not admit any exhaustion by rigid sets with trivial pointwise stabilisers.

The thesis will be structured as follows. In Chapter 1 we will fix notations and conventions, define the main objects and state the main results of the thesis, while providing the necessary background material.

In Chapter 2 we will deal with the rigidity result regarding graphs of regions. In order to achieve such a result we will first investigate the rigidity of graphs of discs and curves. Moreover, we will prove the connectedness of such graphs.

Lastly, in Chapter 3 we will construct exhaustions by finite rigid sets for the strongly separating curve graphs for the 7-holed sphere and the 8-holed sphere. In order to do so

we will introduce a general graph-theoretical machinery to produce such exhaustions, and deduce the co-Hopfian property from them. Furthermore, we give a couple of examples of graphs naturally associated with surfaces, which do not admit any such exhaustion.

Chapter 1

Main Results and Background Material

The definition of the main objects of study, and the technical statements of the main theorem of the thesis are contained in Sections 1.3, 1.4, and 1.6.

Section 1.5 contains the definition of the Kneser graph: an uncommon combinatorial tool which plays a big role in our arguments.

Section 1.1 contains general definitions of objects which are probably well known by the expert reader, and has the role of fixing conventions, but might be skipped at will without compromising the comprehension of the rest of the thesis.

Section 1.2 contains some results about realisation of isotopy classes of curves which are intuitive and constitute some unspoken assumption throughout many arguments. The main purpose of the section is to give formal proofs of some of these assumptions, due to a lack in the literature of the specific formulation needed, but can easily be skipped at first reading and just used as a reference, since we will discuss these subtle issues when they appear in our future arguments.

1.1 Basic Objects and Conventions

1.1.1 Categories

Throughout this thesis, unless otherwise stated, we will work inside the smooth (DIFF) category. This means that the objects will be smooth manifolds, and the morphisms between them will be smooth maps. This will let us make use of notion such as transversality and general position. Whenever we perform any sort of surgery the initially obtained object will just be piecewise-smooth, and will always tacitly assume that some kind of smoothing of the angles is performed, so to remain inside the smooth category.

However, in order to allow for more flexibility, the figures in this thesis are to be thought of in the topological (TOP) category.

The reader should feel free to employ the category they feel most comfortable with. Indeed, in this thesis we will always work with manifolds of dimension at most 2, and in those cases it is a well known fact that the smooth, piecewise-linear (PL), and topological

categories are equivalent (although not in a strictly categorical sense, as some contain more morphisms than other).

The existence of triangulations for topological surfaces and the equivalence between homeomorphisms and PL homeomorphisms (known as *Hauptvermutung* for surfaces) have been first proven by Radó in [R]. For a more modern exposition and proof in English we refer to [Mo, §8 Theorem 3] and [Mo, §8 Theorem 5], respectively.

The existence of unique smooth structures on PL surfaces has been proven by Thom in [Tho]; for a modern proof see [Thu, Theorem 3.10.8, and Theorem 3.10.9]. Lastly, existence and uniqueness of PL structures for smooth surfaces has been first proven by Whitehead in [Wh]; for a modern proof we refer to [Thu, Theorem 3.10.2].

1.1.2 Surfaces and Subsurfaces

For convenience we want to work with surfaces with marked points instead of punctures. This is the reason why we need to define our objects with the language of *topological pairs*, whose definition we will now recall, in accordance with [Sp, p.17].

Definition 1.1.1. A *topological pair* is pair (X, A) where X is a topological space, and $A \subset X$.

As for the maps between pairs we will use the following definition, which is stronger than the usual notion of morphism of pairs. Indeed, usually only Property 1 is required for morphisms of pairs, see again [Sp, p.17], and Property 2 will be the additional property we will need. The reason for this choice will be clear when the definition gets specialised to the case of surfaces.

Definition 1.1.2. Let $(X, A), (Y, B)$ be two topological pairs. We define a *strict map of pairs* $f : (X, A) \longrightarrow (Y, B)$ to be a continuous map $f : X \longrightarrow Y$ such that the following hold:

1. $f(A) \subseteq B$;
2. $f^{-1}(B) \subseteq A$.

We can now define what is the technical meaning the word *surface* will have for us.

Definition 1.1.3. Throughout this thesis a *surface* will be a topological pair (S, Π) where S is 2-dimensional, compact, connected, oriented smooth manifold (possibly with boundary), and $\Pi \subset S$ is a subset such that $\Pi \cap \partial S = \emptyset$. The points in Π will be called the *marked points* of S .

Unless otherwise stated, we will assume that any surface (S, Π) is of *finite type*, that is such that its fundamental group $\pi_1(S \setminus \Pi)$ is finitely generated.

Maps between surfaces will be strict maps of pairs. The conditions we have imposed in the definition of strict map of pairs express the fact that marked points have to be mapped to marked points, and only marked points can be mapped to marked points.

In the following we will more often than not abuse notation and write S instead of the pair (S, Π) to indicate a surface. Similarly we will write $f : S \longrightarrow S'$ to denote a map between surfaces although such map is, technically speaking, a strict map of pairs.

The complete classification of finite-type surfaces was originally proven by Möbius (although his proof was only for triangulated surface): for a modern statement see [FM, Theorem. 1.1], and moreover see [Hi, §9, Theorem 1.9] for a proof (in the smooth category). In the wake of this result we have the following.

Definition 1.1.4. A finite-type surface (S, Π) is said to be of type $S_{g,p}^b$ if it has genus g , p marked points (that is $|\Pi| = p$) and b boundary components.

We define the complexity of a surface S of type $S_{g,p}^b$ as $\xi(S) = 3g + p + b - 3$.

While writing $S_{g,p}^b$ to denote the type of a surface, we will omit the superscript b when there are no boundary components.

A surface of type $S_{0,p}^b$ will be denoted with Σ_p^b (we will still omit b when it is zero). For the most part of the thesis we will only work on these surfaces. The surface Σ_p will be called *p-punctured sphere*. A surface of type Σ_p^1 will be called a *p-punctured disc* (or simply a disc if it has no marked points). The names for these surfaces are purely conventional, as we will work with marked points instead of punctures, but are chosen in order to keep the statements in the thesis closer to the existing literature. Any surface of type Σ_p^b with $p + b = 3$ will be loosely called a *pair of pants*, although we will make care to remark the specific type in every case we will encounter one. Lastly, a surface of type $\Sigma_{0,p}^2$ will be called a *p-punctured annulus* (or simply an annulus if it contains no marked points).

In terms of pairs, (R, P) is a *subsurface* of (S, Π) if $R \subseteq S$ is a 2-dimensional compact submanifold of S (possibly not connected), and $P = \Pi \cap R$. We will not require any condition about boundaries: indeed, it will be common for the subsurfaces we will consider to have their boundary lying in the interior of the ambient surface. We will abuse notation and write $R \subseteq S$ to indicate a subsurface, without explicit mention of the marked points. We remark that, under our definition, subsurfaces are automatically closed subsets of the ambient surface.

In line with our goal to work with closed objects, given a subsurface $R \subseteq S$, we will refer to the closure of $S \setminus R$ as the *complement* of R , and will indicate that by R^c (slightly abusing notation, as it technically is not the set-theoretical complement). In particular, given a punctured disc $D \subseteq S$ we will call D^c the *complementary disc* to D .

1.1.3 Homotopy and Isotopy

We will now recall some convention about homotopies and isotopies. These maps may be thought of as morphisms in any of the categories discussed in §1.1.1 without any change.

We now first give a general definition of homotopy and isotopy, see also [Hat, p.3] and [FM, p.33]. We give it for differential manifolds (hence in the DIFF category), although it is naturally defined with the morphisms in any of the other categories that might be employed.

Definition 1.1.5. Let X, Y be differential manifolds. Let $f, g : X \rightarrow Y$ be two smooth maps. A *homotopy* between f and g is a map $F : X \times [0, 1] \rightarrow Y$ such that, for every $x \in X$, $F(x, 0) = f(x)$ and $F(x, 1) = g(x)$. Moreover, if for every $t \in [0, 1]$ the map $F|_{X \times \{t\}}$ is injective, then F is said to be an *isotopy* between f and g .

All of the homotopies we will consider will be between pairs (M, P) and (N, Q) , hence they are strict map of pairs

$$(M \times [0, 1], P \times [0, 1]) \longrightarrow (N, Q).$$

In particular, for maps into surfaces, homotopies at any time map marked points, and only marked points, to marked points.

When dealing with surfaces with boundary, and unless otherwise stated, we will always assume that every homotopy is relative to the boundary: that is, it pointwise fixes the boundary at any time.

Moreover, let S be a surface and M be a manifold. We recall that we say that two maps $f, g : M \longrightarrow S$ are *ambient homotopic* (resp. *isotopic*) if there exists an automorphism $\varphi : S \longrightarrow S$, homotopic (resp. isotopic) to the identity, and such that $f = \varphi \circ g$.

Since we will only be dealing with manifolds of dimension at most two we may freely promote homotopies to isotopies. Indeed, two curves on a surface are homotopic to each other if and only if they are isotopic, see [FM, Proposition 1.10] for a proof. Similarly, two homeomorphisms (relative to the boundary) of a compact surface are homotopic if and only if they are isotopic, see [FM, Theorem 1.12] and the related discussion. Both of these results are originally due to Reinhold Baer. These results will cover all the cases of interest to us.

Moreover, we recall that any homeomorphism of a compact surface S is isotopic to a diffeomorphism of S . This was first proven independently by James Munkres, Stephen Smale and John Whitehead; see [FM, Theorem 1.13] for further reference.

1.1.4 Curves, Arcs and Regions

In order to clarify our conventions we will now give some definition about curves and arcs which we will follow throughout the thesis. Once again, the following definition are valid for each of the three categories presented in §1.1.1. We first start with *parametrised curves*.

Definition 1.1.6. Let (S, Π) be a surface. A closed (oriented) *parametrised curve* on (S, Π) is a strict map of pairs $\gamma : (S^1, \emptyset) \longrightarrow (S, \Pi)$, where (S^1, \emptyset) is the pair composed of the unit complex circle with no marked points. We equip S^1 with the standard counterclockwise orientation (given by its embedding into \mathbb{C}).

A parametrised curve is said to be *essential* if it not isotopic to a curve contained in a neighbourhood of either a point or a boundary component. It is said to be *peripheral* otherwise.

A parametrised curve is said to be *simple* if it is an embedding.

A trivial loop is obviously non essential. We remark that a loop around a marked point, although nontrivial from an homotopical perspective, is peripheral under our definition.

Since in the previous definition we are using strict map of pairs, this means that curves avoid marked points. Likewise, this means that a homotopy between two curves

consists in a strict map of pairs

$$(S^1 \times [0, 1], \emptyset) \longrightarrow (S, \Pi),$$

that is it avoids the marked points at all times.

In what follows we want the term *curve* to indicate not the single parametrised curve (hence the apparently odd name) but its isotopy class. Thus we make the following definition.

Definition 1.1.7. Let (S, Π) be a surface. An oriented *curve* on S is the isotopy class of an oriented essential simple closed parametrised curve.

A *curve* on S is a class of oriented curves up to the orientation (that is forgetting about the orientation).

We remark that, when a surface contains curves, there it contains countably many.

We remind that, in view of [FM, Proposition 1.10], in the previous definition homotopy classes could have been considered instead of isotopy classes, without any change of note. Moreover, two parametrised curves are isotopic if and only if they are ambient isotopic, see [FM, Proposition 1.11]. For a proof (in the smooth case) of the underlying theorem, known as Extension of Isotopy Theorem, we refer to [Hi, §8 Theorem 1.3].

We will abuse notation and always use lowercase Greek letters to identify both (parametrised) curves and their isotopy classes. Moreover, we will mostly be loose in distinguishing between a curve and one of its representatives. While this is mostly done in order not to make both the notation and the arguments too cumbersome, the fact that it is possible to pick particularly nice representatives (see § 1.2) further justifies this looseness.

We will now start defining arcs.

Definition 1.1.8. Let (S, Π) be a surface. Let $X \subseteq \{0, 1\}$. A *parametrised arc* on S is a strict embedding of pairs $\gamma : ([0, 1], X) \longrightarrow (S, \Pi)$.

An *oriented arc* on S is the isotopy class of a parametrised arc. Unless otherwise stated, we will require such isotopy to be relative to the endpoints, that is relative to $\partial[0, 1] = \{0, 1\}$.

An *arc* on S is a class of oriented arcs up to the orientation.

The set X appearing in the definition of parametrised arcs has the following meaning. For convenience reasons, we want to allow arcs to potentially have marked points as endpoints, but we never want the interior of an arc to contain any marked point. Three cases can now occur: first, if X is empty the arc is disjoint from marked points. Second, if X contains a single point then exactly one of the endpoints of the arc is a marked point. Last, if X contains two points the arc is from a marked point to another. All of these cases will be of relevance in various parts of the thesis.

Moreover, we notice that a parametrised arc a comes with a natural orientation. We will call the point $a(0)$ the *first endpoint*, while the point $a(1)$ will be called the *second endpoint*. Isotopies of parametrised arcs preserve this orientation, so oriented arcs will have a natural orientation.

We will use lowercase Latin characters for both parametrised arcs and arcs. Moreover, we will often abuse notation and confuse arcs with their representatives.

Definition 1.1.8 is intentionally vague. We will always impose some restrictions on the allowed isotopies between parametrised arcs, but the requirements will be defined in each case, according to the specific needs. One general rule will, however, be that an isotopy F between two parametrised arcs is a strict map of pairs

$$([0, 1] \times [0, 1], X \times [0, 1]) \longrightarrow (S, \Pi).$$

Similarly to what happens in the case of curves, homotopy classes could have been considered instead of isotopies. In all the cases of interest to us this does not produce any chance of note.

We will now give a definition of the intersection number between curves and arcs. In what follows the square brackets will indicate an isotopy class of maps.

Definition 1.1.9. Let S be a surface. Let $A = [X \rightarrow S]$ be a curve or an arc, and let $B = [Y \rightarrow S]$ be a curve or an arc. The *geometric intersection number* $i(A, B)$ between A and B is the number

$$i(A, B) = \min \{|f(X) \cap g(Y)| \text{ for } f \in A, g \in B\}.$$

In the previous definition, for the case of arcs, the isotopy classes are to be intended subjected to all the specific restrictions we will give in each case.

From now on, given two parametrised curves α, β we will often abuse notation and confuse the map with its image in the surface. This means we will write $\alpha, \beta \subset S$, we write $\alpha \cap \beta \subset S$ to mean the intersection of the images, and $|\alpha \cap \beta|$ for its cardinality. We will similarly abuse notation confusing parametrised arcs with their images.

We will say that a curve or an arc is *disjoint* from another curve or arc if their intersection number is 0, that is if they admit disjoint representatives. Moreover, we will refer to a (finite) collection of disjoint curves as a *multicurve*. We recall that the maximum size of a multicurve on a surface S is given by its complexity $\xi(S)$; see [FM, p.237].

We will say that a parametrised curve or arc is in *minimal position* with another curve or arc if the cardinality of their intersection is equal to the intersection number between their classes. In the case of two parametrised curves, the following proposition, known as *bigon criterion*, gives a necessary and sufficient condition for them to be in minimal position. For the formal definition of transversality we refer to [Hi, §3.2 p.74], also see § 1.2.1 for further discussion.

Definition 1.1.10. Let S be a surface. Let $\alpha, \beta \subset S$ be two transverse (see § 1.2.1 for definition) parametrised curves. We say that a connected component $D \subseteq S \setminus (\alpha \cup \beta)$ is a (punctured) *bigon* (or that α and β *form a bigon*) if D is a (punctured) disc and both $\partial D \cap \alpha$ and $\partial D \cap \beta$ are (parametrised) arcs.

If the component D is a disc not containing any marked point we will call it an *empty bigon*.

Proposition 1.1.11. *Let S be a surface. Let $\alpha, \beta \subset S$ be two transverse parametrised curves. Then α and β are in minimal position if and only if they do not form any empty bigon.*

The previous proposition is [FM, Proposition 1.7], and we refer to that for its proof.

We will now introduce the concept of *regions*, which will largely be used in what follows. For this definition, we follow [Mc2, § 1].

Definition 1.1.12. Let (S, Π) be a surface. A *region* is an isotopy class of connected a proper subsurface $S' \subsetneq S$ (possibly with boundary and marked points) such that every component of $\partial S'$ is an essential curve.

We will denote the set of regions on a surface S as $\mathcal{R}(S)$.

It is immediate to notice that any two representatives of a region must be homeomorphic, so the topological type of a region is well defined.

Given a region R , the boundary components of any two representatives are isotopic in pairs, hence their isotopy class is a well defined multicurve, which we will call the *boundary* of the region R , and will indicate by ∂R .

If R is homeomorphic to an annulus with no marked points then the boundary of R is homotopic to a single curve γ , hence the entire region R is homotopic to γ . We will refer to such a region as an *annular region* homotopic to γ , and will sometimes abuse language and confuse the annular region with the homotopy class of its boundary, which is a curve, although such curve is technically not a region (as regions are isotopy classes of subsurfaces only). This also proves that if the complexity $\xi(S) \geq 0$ and $S \neq \Sigma_3$, that is S is not a pair of pants, then the set of regions $\mathcal{R}(S)$ is infinite, as it contains at least one element for each curve.

Once again, two regions are said to be disjoint if they admit disjoint representatives.

We will often abuse notation and be loose in distinguishing region from their representatives. The potential issues coming from this will be addressed in § 1.2.2.

1.1.5 Mapping Class Groups

For the sake of clarity here we recall the definition of the (extended) mapping class group of a surface. For a more detailed discussion see, for instance, [FM, §2.1].

Definition 1.1.13. Let (S, Π) be a surface with marked points $\Pi \subset S$ and boundary ∂S . The *mapping class group* of S , which we will denote by $\text{Mod}(S)$, is the group of orientation preserving self-homeomorphism of the pair (S, Π) fixing ∂S pointwise, modulo isotopy relative to the boundary.

The *extended mapping class group* of S , which we will denote by $\text{Mod}^\pm(S)$, is the group of all self-homeomorphism of the pair (S, Π) fixing ∂S pointwise, modulo isotopy relative to ∂S .

Elements of the (extended) mapping class group will be called *(extended) mapping classes*.

We recall that the isotopies here considered are strict maps of pairs

$$F : (S \times [0, 1], \Pi \times [0, 1]) \longrightarrow (S, \Pi)$$

relative to ∂S .

Following common conventions, in our previous definition we allow for the homeomorphisms to permute the marked points, but they have to fix the boundary pointwise.

This means that marked points and boundary components play significant different roles when it comes to the definition of the mapping class group, whereas they are essentially the same when it comes to the study of curves. For example, given two surfaces $S_{g,p}^b$ and $S_{g,p'}^{b'}$ with $p + b = p' + b'$, their (extended) mapping class groups are different if $p \neq p'$.

We remark that in our definitions we might have considered diffeomorphisms instead of homeomorphisms, and homotopy instead of isotopy. The group obtained would have been naturally isomorphic to the ones we defined. Indeed, this is another instance in which all the possible choices for the category to work with are essentially equivalent. See [FM, §2.1] for further reference.

We recall that $\text{Mod}(S) < \text{Mod}^\pm(S)$ and the mapping class group has index 2 inside the extended mapping class group, the quotient being generated by any orientation-reversing map.

Given a curve $\gamma \subset S$ the (right) *Dehn-twist* around γ , as defined in [FM, §3.1.1], will be denoted by $T_\gamma \in \text{Mod}(S)$. We recall that a Dehn twist is a self-homeomorphism supported in an annular neighbourhood of a curve, which can be thought of as keeping one of the boundary components fixed and rotating the other by a full twist.

Similarly, let $\gamma \subset S$ be a curve such that the closure of one component of $S \setminus \gamma$ is of type $S_{0,2}^1$. Then the (right) *half Dehn-twists* around γ , as in [FM, §9.1.3], is a well defined mapping class and will be denoted by H_γ .

Lastly, we state the following very well-known result, originally proven (for closed surfaces) by Dehn in 1938; for a modern version of the papers in English see [De]. For a modern proof, see [FM, Corollary 4.15].

Theorem 1.1.14. *Let S be a finite-type surface. Then, both the mapping class group $\text{Mod}(S)$ and the extended mapping class group $\text{Mod}^\pm(S)$ are finitely generated.*

1.1.6 Graphs

Throughout the thesis we will make use of multiple (unoriented) graphs and subgraphs. For us the definition of a graph will be the following.

Definition 1.1.15. A *graph* Γ is a pair of sets $(\mathcal{V}(\Gamma), \mathcal{E}(\Gamma))$, where

$$\mathcal{E}(\Gamma) \subseteq \mathcal{V}(\Gamma) \times \mathcal{V}(\Gamma) / \sim,$$

where $(a, b) \sim (a, b)$ and $(a, b) \sim (b, a)$. The elements of the set $\mathcal{V}(\Gamma)$ are called *vertices* of the graph, while the elements of $\mathcal{E}(\Gamma)$ are known as *edges*.

We will not distinguish graphs from their geometric realisation: for us graphs will always be 1-dimensional simplicial complexes. Moreover, we endow all the graphs with the natural metric coming from giving length 1 to every edge, and indicate such metric with d . It follows that graphs for us will be geodesic metric spaces. In all our arguments only distances between two vertices will appear: we say that two vertices v, w are k -distant to mean that $d(v, w) = k$. In particular we remark that a pair of 1-distant vertices is exactly a pair of adjacent (that is joined by an edge) vertices, so we will use both terminologies in our arguments.

We will now recall what our definition of *graph homomorphisms* is.

Definition 1.1.16. Let Γ, Γ' be graphs. A *graph homomorphism* between Γ and Γ' is a map $\bar{\Phi} : \mathcal{V}(\Gamma) \rightarrow \mathcal{V}(\Gamma')$ such that for any two vertices $v, w \in \mathcal{V}(\Gamma)$ which are joined by an edge then $\bar{\Phi}(v)$ is adjacent to $\bar{\Phi}(w)$, too.

We remark that we can extend the map $\bar{\Phi}$, as in the previous definition, to a continuous map $\Phi : \Gamma \rightarrow \Gamma'$ defined on the entire graph by mapping every edge isometrically to the edge between the images of the two endpoints.

What follows in the definition of *graph isomorphism*.

Definition 1.1.17. A graph *isomorphism* is a bijective (on the entire graphs, not just between the sets of vertices) morphism or, equivalently, an isometry. We will denote the group of automorphisms of Γ with $\text{Aut}(\Gamma)$.

Equivalently, a graph isomorphism is a bijective map $\bar{\Phi} : \mathcal{V}(\Gamma) \rightarrow \mathcal{V}(\Gamma')$ such that $v, v' \in \mathcal{V}(\Gamma)$ are joined by an edge if and only if $\bar{\Phi}(v), \bar{\Phi}(v') \in \mathcal{V}(\Gamma')$ are adjacent.

We will now give the definition of *equivariance*, which is a property we will often use. The following definition is only given in terms of actions on graphs, which will be the only case of interest for us.

Definition 1.1.18. Let Γ, Γ' be two graphs. Let G be a group acting by automorphisms on both Γ and Γ' . Let $\rho_\Gamma : G \rightarrow \text{Aut}(\Gamma)$ be the group homomorphism induced by the action $G \curvearrowright \Gamma$, and let $\rho_{\Gamma'} : G \rightarrow \text{Aut}(\Gamma')$ be the one induced by the action $G \curvearrowright \Gamma'$. A group homomorphism $\varphi : \text{Aut}(\Gamma) \rightarrow \text{Aut}(\Gamma')$ is said to be *G-equivariant* if the following diagram

$$\begin{array}{ccc} G & \xrightarrow{\text{Id}} & G \\ \rho_\Gamma \downarrow & & \downarrow \rho_{\Gamma'} \\ \text{Aut}(\Gamma) & \xrightarrow{\varphi} & \text{Aut}(\Gamma') \end{array}$$

commutes.

1.1.7 Symmetric Groups

Let X be a set. We will denote the group of permutations on X , also known as *symmetric group* of X , with $\mathcal{S}(X)$. Given a subset $Y \subset X$ we will naturally identify $\mathcal{S}(Y)$ as a subgroup of $\mathcal{S}(X)$.

Given a set X we will use $|X| \in \mathbb{N} \cup \{\infty\}$ to denote its cardinality. The previous notation does not distinguish between different infinite cardinalities: however, apart from this subsection we will always just be interested in distinguishing finite sets from infinite one, so we will not need to be more precise than that.

It is a straightforward and well-known fact that the isomorphism class of a symmetric group $\mathcal{S}(X)$ only depends on the cardinality of the set X . If $|X| = m \in \mathbb{N}$ we will denote the isomorphism class of $\mathcal{S}(\{1, \dots, m\})$ as \mathcal{S}_m . We will moreover denote $\mathcal{S}(\mathbb{N})$ as \mathcal{S}_∞ : let us remark once again that this group is isomorphic to the symmetric group of any countable set. We observe that the group \mathcal{S}_∞ is uncountable, hence it cannot be contained in any finitely generated group. In particular, a subgroup of the (extended) mapping class group of a finite-type surface cannot be isomorphic to it.

1.2 Realisation of Curves and Regions

1.2.1 Transversality and General Position

Throughout the current section we will work in the smooth category. Similar definitions and results are also valid for the PL category, but a richer structure than the purely topological one is needed.

We begin by recalling that two parametrised curves are said to be *transverse* if, at every point in which they intersect, (the pushforward of) their tangent spaces are not zero and together span the entire tangent of the surface. For instance, in a plane two non-parallel lines are transverse, whereas two tangent circumferences are not, as the tangent spaces at the intersection point coincide, so they span a 1-dimensional subspace of the plane. For a more precise and technical definition we refer to [Hi, §3.2 p.74]. We recall that two transverse parametrised curves intersect in a finite number of points, see [Hi, §3 Theorem 3.3] for reference.

Given a finite collection of parametrised curves as above, we call a *triple point* a point belonging to the intersection of (at least) three different curves. This is a pathological behaviour we want to avoid. This is indeed possible up to homotopy. Actually, the following slightly stronger statement holds true.

We will now give definition of *general position*

Definition 1.2.1. Let $\gamma_1, \dots, \gamma_n$ be a finite collection of distinct parametrised curves. We say those curves are in *general position* if they are pairwise transverse and there are no *triple points*.

Given that transversality is a dense condition, and given a collection of parametrised curves $\gamma_1, \dots, \gamma_n$, there exists a collection of curves $\gamma'_1, \dots, \gamma'_n$ which are pairwise transverse to each other and such that γ_i is isotopic to γ'_i . For a reference for this result see [Kos, §2 Corollary 2.5].

Proposition 1.2.2. Let $\gamma_1, \dots, \gamma_n$ be a collection of parametrised curves in general position. Let γ be a parametrised curve. Then there exists a parametrised curve γ' , isotopic to γ , such that the collection $\gamma_1, \dots, \gamma_n, \gamma'$ is in general position.

Moreover, if the parametrised curves in the collection $\gamma_1, \dots, \gamma_n$ are not pairwise in minimal position, all the bigons between them can be removed without interfering with transversality or creating triple points, hence resulting in a collection of curves in general and minimal position.

From now on, unless otherwise stated, when dealing with any collection of curves $[\gamma_1], \dots, [\gamma_n]$ we will abuse notation and identify it with a collection of representatives $\gamma_1, \dots, \gamma_n$ in general and minimal position. The legitimacy of such an assumption follows from the previous discussion.

1.2.2 Geodesic Realisation and Representatives of Regions

Let (S, Π) be a surface with marked points $\Pi \subset S$, such that for the Euler characteristic (see [Hat, p.146]) we have $\chi(S \setminus \Pi) < 0$ (this will cover all the cases of interest for us). Let μ be a finite-area complete metric on $S \setminus \Pi$ with constant curvature -1 , which exists

due to [FM, Theorem 1.2]. We will abuse language and call such a metric a *hyperbolic* metric on S . When dealing with surfaces in the remainder of this section we will always assume they come equipped with a hyperbolic metric. Since all the surfaces we will be interested in throughout the thesis have negative Euler characteristic, this is not restrictive.

Any (essential) curve γ on S admits a unique representative that is a closed geodesic (where the term must be interpreted in the context of Riemannian geometry, hence as a locally length-minimising curve) for the metric μ , see [FM, Proposition 1.3]. Two distinct geodesics are transverse, thanks to the uniqueness property of [Do, §4.4 Proposition 5]. Moreover, two distinct simple closed geodesics are in minimal position, see [FM, Corollary 1.9]. It follows that, given a collection of distinct curves $[\gamma_1], \dots, [\gamma_n]$ their geodesic representatives can be chosen to obtain a collection of representatives which are transverse and in minimal position.

Such a collection, though, is not necessarily free from triple points: for this reason we will generally not use geodesics as representatives. We will however rely on geodesic representatives of curves in two cases: to represent punctured discs, and in order to construct representatives of regions through the procedure we will later discuss.

Let \mathcal{D} be a collection of regions on S which are all isotopy classes of punctured discs. For any disc $D \in \mathcal{D}$ let γ_D be the geodesic representative of ∂D . We define a representative W_D as the closure of the component of $S \setminus \gamma_D$ isotopic to D . When dealing with graphs of discs (see §1.4.1) we will always adopt this set of representatives, and consistently abuse notation identifying a disc D with its representative W_D . Given this conventional choice, intersections and inclusion between discs are well defined. In particular, our choice of representative has the nice property that for any disc D we have $D \cap D^c = \gamma_D$, where γ_D is once again the geodesic representative of ∂D . We will also abuse notation here, writing $D \cap D^c = \partial D$. Moreover any two distinct discs A, B are disjoint if and only if $W_A \cap W_B$ is either empty (that is $B \subsetneq A^c$) or it is a single curve (which occurs when $B = A^c$); otherwise the intersection $A \cap B$ is a union of discs, possibly punctured.

In the case of collections of regions which are not all discs, we need a bit more work in order to define representatives satisfying similarly nice properties. We introduce the necessary technical background for this construction with the next two results.

First, we give a version of the Collar Lemma in terms of lifts of curves, as it will be useful later. This version follows directly from the hyperbolic geometry involved in the proof of the usual Collar Lemma itself. We refer to [Hal] for a proof in the context of Riemann surfaces (which completes the original proof provided by Linda Keen in [Ke]) which directly adapts to our statement. Otherwise see [FM, Lemma 13.6] and their following remark for another proof. Said proof can be followed step by step, only replacing the hyperbolic geometry identities with their analogues in the universal cover, to get the statement we are interested in.

Before we present the result, we recall some notation, which will also be employed elsewhere in the thesis. Let (X, d) be a metric space and $Y \subset X$ be a subset. Let $r \in \mathbb{R}_{\geq 0}$. We will denote the (open) r -regular neighbourhood of Y as

$$\mathcal{N}(Y, r) := \{x \in X \text{ s.t. } d(x, Y) < r\}.$$

The closed neighbourhood will be denoted with $\overline{\mathcal{N}(Y, r)}$.

Moreover, given a parametrised curve $\gamma : S^1 \rightarrow S$ we will follow [FM, §1.2.1] and we will call *lift* of γ a smooth map $\tilde{\gamma} : \mathbb{R} \rightarrow \mathbb{H}^2$ such that the following diagram

$$\begin{array}{ccc} \mathbb{R} & \xrightarrow{\tilde{\gamma}} & \mathbb{H}^2 \\ \pi_S \downarrow & & \downarrow \pi_{S^1} \\ S^1 & \xrightarrow{\gamma} & S \end{array}$$

commutes, where the two descending arrow are the universal covering maps. We recall that a curve admit infinitely many lifts, and that the lift of a geodesic on S is a geodesic ray in \mathbb{H}^2 .

We are now ready to state the version of the Collar Lemma we will use.

Theorem 1.2.3 (Collar Lemma). *Let S be a surface endowed with a hyperbolic metric μ . Then there exists a function $K : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ with the following property. Let α, β be simple closed (local) geodesics on S (possibly the same). Let $\tilde{\alpha}$ (resp. $\tilde{\beta}$) be a lift of α (resp. β) to the universal cover \mathbb{H}^2 . Then*

$$\overline{\mathcal{N}(\tilde{\alpha}, K(L(\alpha)))} \cap \overline{\mathcal{N}(\tilde{\beta}, K(L(\beta)))} = \emptyset \quad \text{if and only if} \quad \tilde{\alpha} \cap \tilde{\beta} = \emptyset.$$

From the Collar Lemma we deduce the following corollary, which will be the key for our construction. It is merely a technically convenient result but, since we have not been able to find any reference for it in the literature, we will give a proof.

Corollary 1.2.4. *Let (S, Π) be a finite-type surface such that $\chi(S \setminus \Pi) < 0$, endowed with a hyperbolic metric μ . Let α, β be simple closed geodesics on S . Let $l \leq K(L(\alpha))$ and $h \leq K(L(\beta))$, where K is the function as in Theorem 1.2.3. Let α' (resp. β') be a component of $\partial\mathcal{N}(\alpha, l)$ (resp. $\partial\mathcal{N}(\beta, h)$). Then the two parametrised curves α' and β' are in general and minimal position.*

Proof. Let $\tilde{\alpha}'$ (resp. $\tilde{\beta}'$) be a lift of α' (resp. β') to the universal cover \mathbb{H}^2 . Let $\tilde{\alpha}$ (resp. $\tilde{\beta}$) be the lifts of α (resp. β) such that $\tilde{\alpha}' \subseteq \partial\mathcal{N}(\tilde{\alpha}, l)$ (resp. $\tilde{\beta}' \subseteq \partial\mathcal{N}(\tilde{\beta}, h)$). In the disc model the lift $\tilde{\alpha}'$ (resp. $\tilde{\beta}'$) is an arc of an Euclidean circumference between the same endpoints of $\tilde{\alpha}$ (resp. $\tilde{\beta}$), although that circumference is not orthogonal to the boundary circle. From the Collar Lemma (Theorem 1.2.3) it follows that two such lifts intersect if and only their endpoints are linked in the boundary circumference $\partial\mathbb{H}^2$, and in that case they intersect exactly once.

First, we will prove that α' and β' are in minimal position. From the Collar Lemma (Theorem 1.2.3) it follows that $\tilde{\alpha}' \cap \tilde{\beta}' \neq \emptyset$ if and only if $\tilde{\alpha} \cap \tilde{\beta} \neq \emptyset$. It follows that there exists a bijective correspondence, which respects the action of $\pi_1(S \setminus \Pi)$ on \mathbb{H}^2 via deck transformations, between $\pi^{-1}(\alpha') \cap \pi^{-1}(\beta')$ and $\pi^{-1}(\alpha) \cap \pi^{-1}(\beta)$ where the map $\pi : \mathbb{H}^2 \rightarrow S$ is the universal covering map. The aforementioned correspondence preserves $\pi_1(S)$ -orbits. Since those orbits are in bijective correspondence with intersection points for the parametrised curves on the surface it follows that $|\alpha' \cap \beta'| = |\alpha \cap \beta|$. Since geodesic are always in minimal position it follows that $|\alpha \cap \beta| = i(\alpha, \beta)$. Lastly, since

the curve α' (resp. β') is isotopic to α (resp. β) it follows that $i(\alpha', \beta') = i(\alpha, \beta)$. From this we conclude that $|\alpha' \cap \beta'| = i(\alpha', \beta')$, that is α' and β' are in minimal position.

Lastly, we claim that α' and β' are in general position, that is they intersect transversely. If not, since the covering map is conformal, there would exist two lifts $\tilde{\alpha}', \tilde{\beta}'$ which are tangent in a point. Since such lifts are arcs of Euclidean circumferences with linked endpoints (as they intersect) it is impossible for them to be tangent. This is a contradiction, and hence our claim is proven. \square

We can now define nice representatives of regions in the following way. If R is an annular region homotopic to a simple closed geodesic $\gamma \subset S$ let

$$W_R = \overline{\mathcal{N}\left(\gamma, \frac{K(L(\alpha))}{2}\right)}.$$

If R is a nonannular region let $\gamma^1, \dots, \gamma^n$ be the simple closed geodesics representing the multicurve ∂R . Let X be the connected component of $S \setminus (\gamma^1 \cup \dots \cup \gamma^n)$ isotopic to R . Then we define a representative for R by

$$W_R = X \setminus \left(\bigcup_1^n \mathcal{N}(\gamma^i, K(L(\gamma^i))) \right).$$

The previous set of representatives is defined in a way such that two distinct regions R, Q are disjoint if and only if we have $W_R \cap W_Q = \emptyset$. Moreover, it follows directly from Corollary 1.2.4 that the boundaries of two such representatives are in general and minimal position. Indeed, such collection is the union of two multicurves so there exist no three curves pairwise intersecting, hence there are no triple points.

Given two region R, Q their intersection, defined as the isotopy class of $W_R \cap W_Q$, is now well-defined as a union of disjoint isotopy classes of subsurfaces (not necessarily regions, as they might have nonessential boundary components). We will denote this intersection as $R \cap Q$. Moreover, the notion of inclusion between regions is now well defined. This was the main reason for the choice of representatives we have introduced.

In what follows we will constantly abuse notation and, when we refer to a region R , we will actually refer to the representative W_R , as constructed before.

1.2.3 Alexander's Method

Another way to obtain a good definition of the intersection of two regions is the use of the so-called *Alexander Method*, which we are now going to discuss, as it will also be needed later in the thesis. We start with the following lemma, that is [FM, Lemma 2.9], which we refer to for the proof.

Lemma 1.2.5. *Let S be a surface. Let $\gamma_1, \dots, \gamma_n$ be a collection of essential simple closed parametrised curves such that the following hold:*

1. *The parametrised curves γ_i are pairwise in general and minimal position;*
2. *The parametrised curves γ_i are pairwise non-isotopic;*
3. *There are no triangles, that is, for every three distinct parametrised curves $\gamma_i, \gamma_j, \gamma_k$ at least two are pairwise disjoint.*

Let $\gamma'_1, \dots, \gamma'_n$ be another collection of simple closed parametrised curves with the same properties, and such that for every i the curves γ_i and γ'_i are isotopic. Then there exists an ambient isotopy of S which simultaneously maps every parametrised curve γ_i to γ'_i .

Given two regions R, Q any two sets of representatives for $\partial R \cup \partial Q$ which are in general and minimal position automatically satisfy the hypotheses of the previous lemma. It follows that, given any two such collection, there exists a global isotopy mapping one to the other. It follows that the intersection between two regions can be defined this way, and it is well defined up to a global isotopy of the surface.

The statement of Alexander's Method which we will need later is the following. This is a weaker version of [FM, Proposition 2.8], and we refer to that for the proof.

Proposition 1.2.6 (Alexander's Method). *Let S be a surface. Let $\gamma_1, \dots, \gamma_n$ be a collection of essential oriented simple closed parametrised curves such that the following hold:*

1. *The parametrised curves γ_i are pairwise in general and minimal position;*
2. *The parametrised curves γ_i are pairwise non-isotopic;*
3. *There are no triangles, that is for every three distinct parametrised curves $\gamma_i, \gamma_j, \gamma_k$ at least two are pairwise disjoint.*
4. *The parametrised curves γ_i fill S , that is every component of $S \setminus (\gamma_1 \cup \dots \cup \gamma_n)$ is either a disc or a once-punctured disc.*

Let $f \in \text{Homeo}^+(S, \partial S)$ be an orientation-preserving self-homeomorphism of S relative to the boundary. Suppose that for every $i = 1, \dots, n$ the parametrised curve $f(\gamma_i)$ is isotopic to γ_i and it has the same orientation. Then f is isotopic to the identity Id_S .

1.3 Graphs of Curves

1.3.1 The Classical Curve Graph

We recall the definition of the *curve graph* of a surface, as initially introduced by Harvey in [Harv1].

Definition 1.3.1. Let S be a surface (not necessarily of finite type). The *curve graph* $\mathcal{C}(S)$ of S is the abstract graph defined as follows:

Vertices There is one vertex for every curve (that is for every isotopy class of simple closed essential curves) on S ;

Edges There is an edge between two vertices corresponding to two curves if they are disjoint (that is they admit disjoint representatives).

We will almost always abuse notation and indicate the vertices of the curve graph as curves, that is with lowercase Greek letters.

We notice that if the complexity is $\xi(S) \leq 0$ and S is not the torus $S_{1,0}$, every parametrised curve on S is nonessential, so the curve graph has no vertices. If S is either of type $S_{1,0}$ or $\xi(S) = 1$ every two curves on S intersect so, given our definition, the curve graph is discrete, hence not very interesting. A different definition for the graph can be given in these cases, only requiring two curves to have the minimal possible intersection for them to be connected by an edge, producing a Farey graph; see [FM, §4.1.1] for further details. However, this refinement of the definition will play no role in this thesis, as we will only focus on surfaces of complexity $\xi(S) \geq 2$. In these cases the curve graph is actually connected, as was already known to Harvey, see [Harv1, Proposition 2]. For another proof, there given in the closed case but adaptable to the general one, also see [FM, Theorem 4.4].

Let α, β be two parametrised curves on S . Moreover, let $f, g \in \text{Homeo}(S)$ be two isotopic self-homeomorphisms. If α and β are isotopic it follows that $f(\alpha) = f(\beta)$ so $f([\alpha])$ is a well-defined curve. Moreover $f([\alpha])$ is isotopic to $g([\alpha])$, so, given a mapping class $[f] \in \text{Mod}^\pm(S)$, its action on the vertices of the curve graphs is well defined. If α and β are disjoint then $f(\alpha)$ is disjoint from $f(\beta)$. It follows that a group action $\text{Mod}^\pm(S) \curvearrowright \mathcal{C}(S)$ by graph automorphisms is well-defined. Given an extended mapping class $f \in \text{Mod}^\pm(S)$ and a vertex of the graph, that is a curve $\gamma \in \mathcal{V}(\mathcal{C}(S))$, we will loosely denote the action of the mapping class on the curve as $f(\gamma)$.

Apart from some sporadic cases, the aforementioned action is faithful and induces all the automorphisms of the graph. This is the content of the following celebrated and important rigidity result, known as Ivanov's Theorem, as the first proof is due to Nikolai Ivanov in [Iv1] in the case of genus at least two, while the missing cases were proven by Korkmaz in [Kor]. An independent proof covering all the cases has been provided by Luo in [Lu].

Theorem 1.3.2 (Ivanov's Theorem). *Let S be a surface of type $S_{g,p}^0$ with complexity $\xi(S) \geq 2$ and different from $S_{1,2}$ or $S_{2,0}$. Then the curve graph $\mathcal{C}(S)$ is rigid, that is the group homomorphism*

$$\rho : \text{Mod}^\pm(S) \longrightarrow \text{Aut}(\mathcal{C}(S))$$

induced by the natural action is an isomorphism.

1.3.2 Subgraphs of the Curve Graph

In this subsection we will introduce some important subgraphs of the curve graph that will play a major role in this thesis. In order to define them we will first introduce some properties of curves. The following definition introduces the first two such properties and the first important subgraph, the *strongly separating curve graph* introduced by Bowditch in [B2].

We recall that by subgraph of a graph Γ induced by a set of vertices V we mean the graph having V as the set of vertices and having an edge between two of them if and only those vertices are connected by an edge in Γ .

Definition 1.3.3. Let S be a surface. A curve $\gamma \subset S$ is said to be *separating* if the set $S \setminus \gamma$ is disconnected.

A separating curve $\gamma \subset S$ is said to be *strongly separating* if no closure of any component of $S \setminus \gamma$ is a pair of pants, that is of type $S_{0,p}^b$ with $p + b = 3$.

The *strongly separating curve graph* $\mathcal{C}_{ss}(S)$ is the subgraph of the curve graph induced by vertices representing strongly separating curves.

In other terms, the strongly separating curve graph is the graph whose vertices are strongly separating curves, and two curves are joined by an edge if and only if they are disjoint.

Since surface self-homeomorphisms preserve the set of separating curves, and the topological types of the complementary components, and vertices and edges are defined up to isotopy, the extended mapping class group naturally acts by automorphisms on the strongly separating curve graph. Apart from some sporadic cases, the strongly separating curve graph is rigid. This result, which is [B2, Theorem 1.1], was originally proven by Bowditch in order to prove quasi isometric rigidity of the Weil-Petersson space, see [B3].

Theorem 1.3.4. Let S be a surface of type $S_{g,p}^0$ with $g + p \geq 7$. Then the strongly separating curve graph $\mathcal{C}_{ss}(S)$ is rigid, that is the group homomorphism

$$\rho : \text{Mod}^\pm(S) \longrightarrow \text{Aut}(\mathcal{C}_{ss}(S))$$

induced by the natural action is an isomorphism.

Restricted to the case of punctured spheres, our Theorem A is a generalisation of the previous rigidity result, although it says nothing about the case of non-planar surfaces.

From now on we will exclusively focus our attention to the case of punctured spheres. We recall that we will denote a p -punctured sphere, that is a surface of type $S_{0,p}^0$, as Σ_p . On punctured spheres every curve is separating (this is the Jordan Curve Theorem, first stated by Camille Jordan in [J]: for a modern proof, using algebraic topology, see [Hat, Proposition 2B.1.b]) and we can easily classify curves according to the number of marked points on each side. Practically speaking, in order to make the notation less cumbersome, we will classify them in terms of the number of marked points on the side which contain less. This is the content of the following definition. We recall that, given

a curve, the closure of their complementary components are regions well defined up to isotopy, see §1.2, in particular their topological type is well defined.

Definition 1.3.5. Let $m \in \mathbb{N}$ with $m \geq 4$. Let Σ_m be the m -punctured sphere. Let $\gamma \subset \Sigma_m$ be a curve. The closure of the two connected components of $\Sigma_m \setminus \gamma$ are two punctured discs, of types $S_{0,l}^1$ and $S_{0,h}^1$, with $2 \leq l \leq h$. We say that γ is a *l-separating curve*.

A 1-separating curve would be nonessential, hence the hypothesis $l \geq 2$ is not restrictive. Moreover, we observe that two isotopic curves have the same topological type.

The definition of the following two subgraphs of the curve graph comes naturally.

Definition 1.3.6. Let Σ_m be a m -punctured sphere and let $k \geq 2$. We define the *k-separating curve graph* as the subgraph of the curve graph induced by h -separating curves with $k \leq h$: we will denote it with $\mathcal{C}_k(\Sigma_m)$.

The subgraph of $\mathcal{C}_k(\Sigma_m)$ induced by exactly k -separating curves only will be called *strict k-separating curve graph* and denoted by $\mathcal{C}_{(k)}(\Sigma_m)$.

The graphs we have just defined are only interesting if $m \geq 2k+1$. Indeed, if $m < 2k$ there exists no k -separating curve on Σ_m , hence the graph is empty. When $m = 2k$ two k -separating curves on Σ_m are disjoint if and only if they cobound an annulus, hence they are isotopic: it follows that the graph is discrete. When $m \geq 2k+1$ the (strict) k -separating curve graphs are actually connected: we will give multiple proofs of this fact in § 2.3, both following from the rigidity results and a *a priori* ones.

The simplest case we can have is for $k = 2$ and $m \geq 5$: the graph $\mathcal{C}_2(\Sigma_m)$ corresponds to the usual (separating) curve graph $\mathcal{C}(\Sigma_m)$. When $k = 3$ and $m \geq 7$ the graph $\mathcal{C}_3(\Sigma_m)$ is the strongly separating curve graph $\mathcal{C}_{ss}(\Sigma_m)$ described by Bowditch in [B2]. Moreover, we observe that when $m = 2k+1$ every vertex of the k -separating curve graph corresponds to a k -separating curve, so the graph $\mathcal{C}_k(\Sigma_{2k+1})$ coincides with the (strict) k -separating curve graph $\mathcal{C}_{(k)}(\Sigma_{2k+1})$.

We observe that there exist the following two natural inclusions of graphs:

$$\begin{aligned}\mathcal{C}_k(\Sigma_m) &\hookrightarrow \mathcal{C}_{k+1}(\Sigma_m); \\ \mathcal{C}_{(k)}(\Sigma_m) &\hookrightarrow \mathcal{C}_k(\Sigma_m).\end{aligned}$$

For every k , every self-homeomorphism of the ambient surface preserves the set of k -separating curves, and its action on the set of those curves is transitive, see [FM, §1.3]. Since vertices and edges are defined up to isotopy, the extended mapping class group naturally acts on the (strict) k -separating curve graph. Indeed, in all but some sporadic cases those graphs are rigid. This is the content of our first main theorem (which covers the case of the non-strictly separating curve graphs, the other following from Theorem D).

Theorem A. *Let $k \geq 2$ and $m \geq 2k+1$. Let Σ_m be the m -punctured sphere. Then the k -separating curve graph $\mathcal{C}_k(\Sigma_m)$ is rigid, that is the group homomorphism*

$$\rho : \text{Mod}^\pm(\Sigma_m) \longrightarrow \text{Aut}(\mathcal{C}_k(\Sigma_m))$$

induced by the natural action is an isomorphism.

We have already remarked that the previous theorem is an extension of [B2, Theorem 1.1] for planar surfaces. Moreover, this result is a sharpening of [Mc1, Theorem 1.5] (which is [Mc2, Theorem 2] in the planar case), which only holds for $m \geq 3k - 1$.

Lastly, let us note that the bound on the complexity of the surface in term of k we have used in our result is sharp. Indeed, if $m \leq 2k$ we have noticed that $\mathcal{C}_k(\Sigma_m)$ is either empty or countable and discrete. In the first case the automorphism group is empty, in the other case it is the infinite permutation group \mathcal{S}_∞ . In neither case such a group can be isomorphic to the extended mapping class group of the surface, which is nontrivial but finally generated, thanks to the observation made in §1.1.7.

1.4 Graphs of Discs and Regions

1.4.1 Graphs of k -Separating Discs

During most of Chapter 2 we will not deal with the graph of k -separating curve, but with a closely related graph whose vertices will not be curves but rather discs. This is mostly due to technical reasons: indeed, in our arguments we would have to repeatedly mention the discs bounded by a curve, and the fact that there is a possible ambiguity in the choice of the side would have made the notation, and hence the arguments, much more cumbersome. In order to avoid this issue we will work with graphs whose vertices are discs, since their boundary curves are always well-defined, so our notation will be much less convoluted. Moreover, these graphs are an interesting first example of graphs of regions, which we will introduce in Definition 1.4.3, so they are worth considering on their own.

These are the reason for the following definitions.

Definition 1.4.1. Let $k, m \in \mathbb{N}$. Let Σ_m be the sphere with m -marked points. We define the k -separating disc graph $\mathcal{D}_k(\Sigma_m)$ as the abstract graph given by the following.

Vertices There is vertex for each isotopy class of p -punctured discs $D \subseteq \Sigma_m$, with $k \leq p \leq m - k$;

Edges There is an edge between two vertices if the corresponding classes of discs are disjoint (that is they admit disjoint representatives).

Analogously to the k -separating curve graph the aforementioned graph is only interesting for $m \geq 2k + 1$. Indeed, if $m < 2k$ the graph is empty, and if $m = 2k$ it has infinitely many connected components, each consisting of a pair of complementary discs $\{D, D^c\}$ only.

Similarly to the definition of the strict k -separating curve graph we have given in Definition 1.3.6 we will now define a graph whose elements are k -punctured discs only, which will come up in some of our later arguments.

Definition 1.4.2. Let $k, m \in \mathbb{N}$. Let Σ_m be the sphere with m -marked points. We define the graph of k -punctured discs $\mathcal{D}_{(k)}(\Sigma_m)$ as the full subgraph of $\mathcal{D}_k(\Sigma_m)$ induced by vertices corresponding to discs containing exactly k marked points.

It is immediate to observe that, when $m \geq 2k + 1$, two k -punctured discs in Σ_m are disjoint if and only if their boundary curves are. Indeed, this proves that the graph $\mathcal{D}_{(k)}(\Sigma_m)$ is naturally isomorphic to the strict k -separating curve graph $\mathcal{C}_{(k)}(\Sigma_m)$, when $m \geq 2k + 1$.

For any p , any self-homeomorphism of the ambient surface preserves the set of p -punctured discs, and its action on such set is transitive. Moreover, since vertices and edges are defined up to isotopy, it follows that the extended mapping class group naturally acts on the k -separating disc graph by graph automorphisms. It turns out that, apart from some sporadic cases, this graph is actually rigid. This is the content of the following result, to the proof of which the most part of Chapter 2 will be dedicated.

Theorem B. *Let $k \geq 2$ and $m \geq 2k + 1$. Let Σ_m be the m -punctured sphere. Then the k -separating disc graph $\mathcal{D}_k(\Sigma_m)$ is rigid, that is the group homomorphism*

$$\rho : \text{Mod}^\pm(\Sigma_m) \longrightarrow \text{Aut}(\mathcal{D}_k(\Sigma_m))$$

induced by the natural action is an isomorphism.

A weaker version of the previous theorem was already known, as it follows from [Mc1, Theorem 3] (or [Mc2, Theorem 2] in the planar case) for $m \geq 3k - 1$.

Once again, we remark that the bound in the previous result is sharp. Indeed, when $m < 2k$ the automorphism group is empty, and when $m = 2k$ the graph is a countable disjoint union of pairs of vertices joined by an edge, hence the automorphism group contains an infinite permutation group \mathcal{S}_∞ as a subgroup. Since this group is uncountable, as noted in §1.1.7, the automorphism group cannot be isomorphic to any mapping class group.

1.4.2 Graphs of Regions

We will now introduce the last class of graphs naturally arising from surfaces, that of graphs of regions, which we will encounter later in this thesis. Such class of graphs encompass all the previously defined ones, hence will be the object of our most general rigidity result. This kind of graphs has first been introduced by McCarthy and Papadopoulos in [MP], and was popularised, in the context of closed surfaces, by Brendle and Margalit in [BM1]; we will follow the definition for punctured spheres given by McLeay in [Mc1, §1]. For the definition of regions, we refer to Definition 1.1.12.

Definition 1.4.3. Let $\mathcal{A} \subseteq \mathcal{R}(\Sigma_m)$ be a $\text{Mod}^\pm(\Sigma_m)$ -invariant set of regions. We define the *graph of regions* subordinate to \mathcal{A} , written $\mathcal{G}_\mathcal{A}(\Sigma_m)$, as the abstract graph defined as follows.

Vertices There is a vertex for every region $R \in \mathcal{A}$;

Edges There is an edge between two vertices if the corresponding regions are disjoint (that is they admit disjoint representatives).

The k -separating disc graph introduced in Definition 1.4.1 is a first example of a graph of regions. Moreover, the curve graph, and hence its subgraphs, can be thought as a graph of regions, by taking as set \mathcal{A} the set of annuli with essential boundary. The graph of regions $\mathcal{G}_\mathcal{A}(\Sigma_m)$ is naturally isomorphic to the curve graph, and this isomorphism is $\text{Mod}^\pm(\Sigma_m)$ -equivariant.

Given that we have required the set of regions \mathcal{A} to be $\text{Mod}^\pm(\Sigma_m)$ -invariant the mapping class group naturally acts by automorphisms on the graph of regions $\mathcal{G}_\mathcal{A}(\Sigma_m)$. We will produce a complete classification of the graphs of regions on planar surfaces which are rigid (Theorem D). The two ingredients we are currently missing in order to give the precise statement are the following.

First, we will need a way to measure the complexity of the set of regions \mathcal{A} , such that we can impose a lower bound on the complexity of the surface, in order to avoid empty or disconnected graphs. This measure of complexity, which will be a natural number $\nu_\mathcal{A}$,

will play a role analogous to the one the number k played for the graphs of k -separating curves or discs.

Second, since we are now allowing the considered regions to be not only discs or annuli, some new potentially pathological configurations can arise, which might prevent the graph from being rigid. In particular, these configurations allow for the existence of automorphisms, called *exchange automorphisms*, which permute two vertices while fixing all the other, and such automorphisms cannot be induced by mapping classes. We will give a complete topological description of the pathologies which generate such obstructions to rigidity.

The definition of the complexity of a set of regions, which is the same definition as in [Mc1, §1], is the following.

Definition 1.4.4. Let $m \geq 5$. Let $R \in \mathcal{R}(\Sigma_m)$ be a region of the m -punctured sphere Σ_m . We define the *complexity* $\nu(R)$ to be the minimum natural number such that there exists a $\nu(R)$ -punctured disc $D \subseteq \Sigma_m$ containing R .

Given a set of regions $\mathcal{A} \subseteq \mathcal{R}(\Sigma_m)$ we define its *complexity* $\nu_{\mathcal{A}}$ to be the minimum of $\nu(R)$ with R varying among every region $R \in \mathcal{A}$.

For examples of the complexity of region see Figures 1.1 and 1.2.

We immediately remark that, since boundary components of regions are essential by definition, it follows that for every collection of regions \mathcal{A} it always holds $\nu_{\mathcal{A}} \geq 2$.

We notice that the complexity of a h -punctured disc is h , and the complexity of an annulus homotopic to a h -separating curve is exactly h . It follows that the complexity of the k -separating curve (resp. disc) graphs, thought as graphs of regions, is exactly k , as expected.

We begin with the definition of *vertices with holes*. Before that, we introduce some notation: given a region $R \in \mathcal{R}(\Sigma_m)$, a *complementary disc* of R is the isotopy class of the closure of one of the connected components of $\Sigma_m \setminus R$ (which is well-defined region, as it is a disc with at least two marked points).

Definition 1.4.5 (Vertices with holes). Let $\mathcal{A} \subseteq \mathcal{R}(\Sigma_m)$ be a $\text{Mod}^{\pm}(\Sigma_m)$ -invariant set of regions. A vertex $R \in \mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is said to be a *vertex with a hole* if one of the complementary discs of R does not contain any subsurface representing a region of \mathcal{A} .

We notice that, by definition, an annular region without marked points never has holes, since any complementary discs contains a subsurface which is isotopic to it.

For an example of a region with holes let us consider the set $\mathcal{A} \subseteq \mathcal{R}(\Sigma_5)$ given by once-punctured annuli with essential boundary (that is surface of type $S_{0,1}^2$). Any of those regions has a hole, since if one of its complementary discs, which have two punctures each, contained a once-punctured annulus, this would have one boundary component isotopic to a marked point. For a picture of these regions see Figure 1.1.

The second type of pathological vertices is that of *cork pairs*.

Definition 1.4.6 (Cork pairs). Let $\mathcal{A} \subseteq \mathcal{R}(\Sigma_m)$ be a $\text{Mod}^{\pm}(\Sigma_m)$ -invariant set of regions. A pair of vertices (P, D) , with $P, D \in \mathcal{A}$ are said to form a *cork pair* if P is represented by an annular region with a complementary disc D and no proper, nonperipheral (that is not isotopic to ∂D) subsurface of D represents an element of \mathcal{A} .

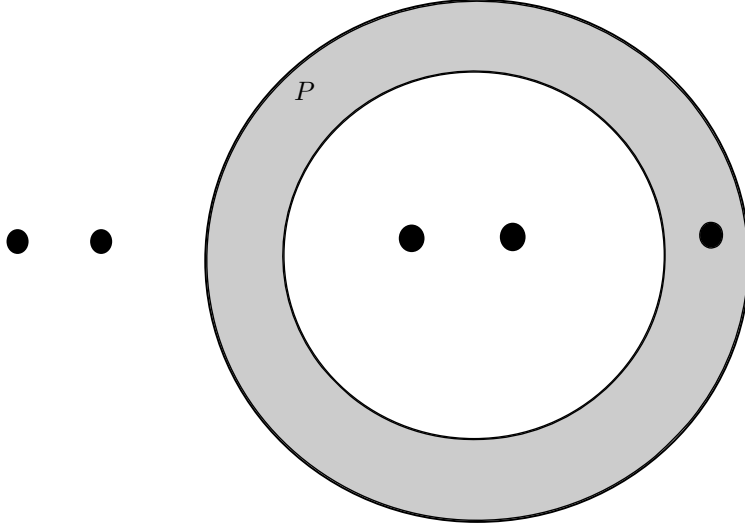


Figure 1.1: The once-punctured annulus P is shaded in gray. Any of the two complementary disc is a hole for it. Note that the complexity is $\nu(P) = 3$.

An example of a graph of regions admitting cork pairs is the following. Let $\mathcal{A} \subseteq \mathcal{R}(\Sigma_5)$ be the set of annular regions and twice-punctured discs. Then any annular region and the twice-punctured disc it bounds form a cork pair. Indeed, any essential curve in a 2-punctured disc is isotopic to the boundary. This can be seen in Figure 1.2 (which also explains the name of the object, as the disc is the “cork” of the bottle, while the annular region is the “neck”).

We can now state our rigidity result for graphs of planar regions, which is a complete classification.

Theorem C. *Let $m \geq 5$. Let Σ_m be the m -punctured sphere. Let $\mathcal{A} \subseteq \mathcal{R}(\Sigma_m)$ be a $\text{Mod}^\pm(\Sigma_m)$ -invariant collection of regions, and let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be the associated graph of regions. Then the graph is rigid, i.e. the natural homomorphism*

$$\rho : \text{Mod}^\pm(\Sigma_m) \longrightarrow \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$$

is an isomorphism, if and only if $m \geq 2\nu_{\mathcal{A}} + 1$, and the graph has no vertices with holes and no cork pairs.

The previous theorem is a sharp version of [Mc1, Theorem 3] (that is [Mc2, Theorem 2] in the planar case), which proves the same rigidity result with the weaker bound $m \geq 3\nu_{\mathcal{A}} - 1$.

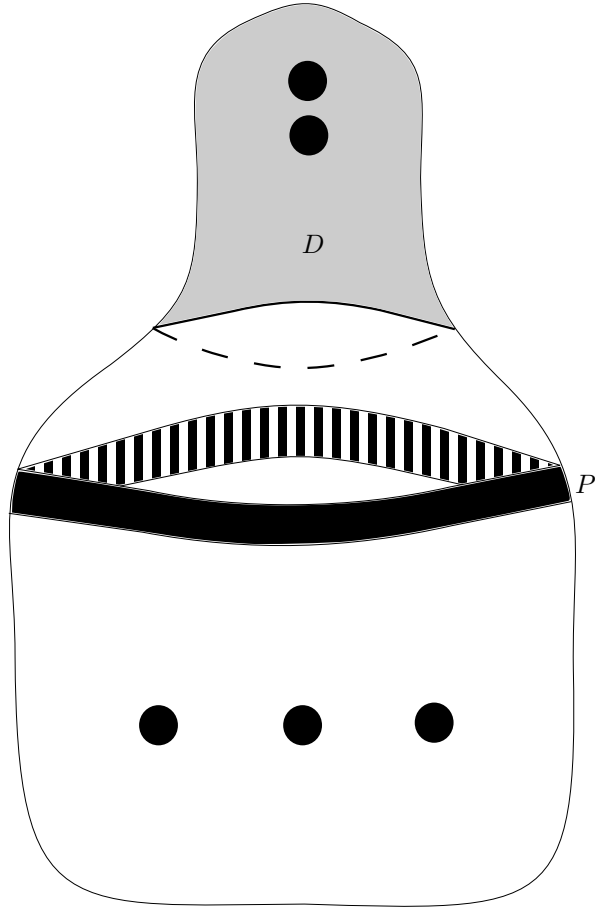


Figure 1.2: The 5-holed sphere is represented as the doubling of the “bottle” in figure. The regions (P, D) form a cork pair. The annulus P is the “neck of the bottle”, while the twice-punctured disc D is the “cork”. Note that the complexity is $\nu(P) = \nu(D) = 2$.

1.5 Kneser Graphs

1.5.1 Standard Kneser Graphs

Throughout many of the arguments of this thesis, we will need to keep track of the combinatorics of marked points, and relate it to configurations arising from graphs of curves or discs. This is the reason for us to introduce the *Kneser graphs*, whose definition is originally due to Kneser in [Kn]. For a more accessible English reference for the definition, see [SU, §3.2].

Definition 1.5.1. Let $m, k \in \mathbb{N}$, $m \geq k$. Let $\Pi = \{1, \dots, m\}$. The *Kneser graph* $K(m, k)$ is the abstract graph defined as follows.

Vertices There is one vertex for every subset $V \subseteq \Pi$ of cardinality exactly k ;

Edges There is an edge between two vertices if the corresponding sets are disjoint.

Since permutations preserve cardinalities and disjointness, we notice that the symmetric group \mathcal{S}_m naturally acts on $K(m, k)$ by graph automorphisms.

The reason we are interested to Kneser graphs is the following. Let Σ_m be the m -punctured sphere, and let the marked points be labelled by $\Pi = \{1, \dots, m\}$. Let $\gamma \subset S$ be a k -separating curve. If $m \geq 2k + 1$ then there exists a unique complementary disc for γ , which contains exactly k marked points. Let $\pi(\gamma) \subseteq \Pi$ be the set of those marked points: this is a subset of Π of cardinality k , hence a vertex of the Kneser graph $K(m, k)$. Moreover, if γ, γ' are two distinct disjoint k -separating curves then their k -punctured complementary discs are disjoint (otherwise the two curves would cobound an annulus, hence they would be isotopic), so in particular the vertices $\pi(\gamma), \pi(\gamma') \in K(m, k)$ are disjoint, hence adjacent. It follow that it is defined a graph homomorphism:

$$\pi : \mathcal{C}_{(k)}(\Sigma_m) \longrightarrow K(m, k)$$

from the strict k -separating curve graph to the Kneser graph $K(m, k)$.

This morphism will be a key component of the arguments in Chapter 3.

1.5.2 Extended Kneser Graphs

For the arguments throughout Chapter 2 we will need the following variation on Kneser graphs, which is original.

Definition 1.5.2. Let $m, k \in \mathbb{N}$, $m \geq k$. Let $\Pi = \{1, \dots, m\}$. The *extended Kneser graph* $K^*(m, k)$ is the abstract graph defined as follows.

Vertices There is one vertex for every subset $V \subseteq \Pi$ such that $k \leq |V| \leq m - k$;

Edges There is an edge between two vertices if the corresponding sets are disjoint.

A vertex given by a subset of cardinality h will be called a h -vertex.

Since permutations preserve cardinalities and disjointness, we notice that the symmetric group \mathcal{S}_m naturally acts on $K^*(m, k)$ by graph automorphisms.

Let Σ_m be the m -punctured sphere, and let its marked points be labelled by $\Pi = \{1, \dots, m\}$. Let $D \subset S$ be a h -punctured disc with $k \leq h \leq m-k$. Let $\pi(D) = D \cap \Pi \subseteq \Pi$ be the set of marked points contained in D : this is a h -vertex of the extended Kneser graph $K^*(m, k)$. Moreover, given two disjoint punctured discs D, D' the sets $\pi(D), \pi(D')$ are disjoint, hence adjacent as vertices of the extended Kneser graph. It follows that a graph homomorphism:

$$\pi : \mathcal{D}_k(\Sigma_m) \longrightarrow K^*(m, k)$$

from the k -separating disc graph to the extended Kneser graph $K^*(m, k)$ is well-defined. This morphism maps h -punctured discs to h -vertices of the extended Kneser graph. Moreover, we have $\pi(D^c) = \pi(D)^c$, where D^c is the closure of the complementary of the disc D , while $\pi(D)^c$ is the set of the marked points not contained in D .

1.6 Rigid Subgraphs and Exhaustions

So far, we have only discussed a notion of “global” (also called combinatorial or simplicial) rigidity for graphs, that is the natural action of the mapping class group inducing isomorphisms with the group of automorphisms. We will now introduce a “local” concept of rigidity for subgraphs, which is in the sense of [AL2]. This will be the main topic of Chapter 3.

Definition 1.6.1. Let Γ be a graph. A subgraph $X \subseteq \Gamma$ is said to be *rigid* if for every graph embedding $i : X \hookrightarrow \Gamma$ there exists a graph automorphism $\varphi \in \text{Aut}(\Gamma)$ such that $i = \varphi|_X$. A graph which is rigid, as a subgraph of itself, is said to be *co-Hopfian*.

We also recall the definition of exhaustions of graphs.

Definition 1.6.2. Let Γ be a graph. A sequence of subgraphs $X_i \subseteq \Gamma$ is said to be an *exhaustion* if $X_i \subseteq X_{i+1}$ and

$$\bigcup_{i \in \mathbb{N}} X_i = \Gamma.$$

If every set X_i is finite (resp. rigid) we say that the exhaustion is by finite (resp. rigid) sets.

Aramayona and Leininger proved the existence not only of a finite rigid subgraph of the curve graph ([AL1, Theorem 1.1]), which is an interesting problem per se, but also the existence of an exhaustion by finite rigid sets ([AL2, Theorem 1.1]).

We will focus on the case of strongly separating curve graph of the 7- and 8-holes sphere. A finite rigid set for $\mathcal{C}_{ss}(\Sigma_7)$ has been discovered by Bowditch in ([B2, Proposition 3.1]). In the same paper, Bowditch also produced a finite rigid set of the strict graph of 3-separating curve graph of Σ_8 , that is, $\mathcal{C}_{(3)}(\Sigma_8)$ ([B2, Lemma 7.2]), which we will adapt to find a finite rigid set for $\mathcal{C}_{ss}(\Sigma_8)$.

Our main result in this field, whose proof will occupy most of Chapter 3, will be existence of an exhaustion by finite rigid sets for the strongly separating curve graph of the 7- and 8-holes sphere.

Theorem D. *Let $S = \Sigma_7$ or $S = \Sigma_8$. Then the strongly separating curve graph $\mathcal{C}_{ss}(S)$ admits an exhaustion by finite rigid sets*

$$\bigcup_{i \in \mathbb{N}} X_i = \mathcal{C}_{ss}(S).$$

Moreover, every subgraph X_i has trivial pointwise stabiliser.

From the previous theorem our last result, which is a co-Hopfian property for the strongly separating curve graphs of the 7- and 8-holes sphere will follow.

The analogous co-Hopfian property for the standard curve graph has first been proven by Shackleton ([Sha, Theorem 1]), and also deduced from the existence of exhaustion by rigid sets by Aramayona and Leininger ([AL2, Corollary 1.2]).

Theorem E. *Let $S = \Sigma_7$ or $S = \Sigma_8$. Then, for every injective graph self-embedding $i : \mathcal{C}_{ss}(S) \hookrightarrow \mathcal{C}_{ss}(S)$, there exists an extended mapping class $f \in \text{Mod}^\pm(S)$ such that $i = f$, that is the self embedding i coincides with the map induced on the curve graph by the mapping class f .*

In the cases of interest to us this is a stronger statement than [B2, Theorem 1.1] (that is Theorem A), which was the combinatorial rigidity.

This co-Hopfian property is a direct consequence of the existence of exhaustions by rigid sets with trivial pointwise stabilisers. We will discuss an example in which this property fails, hence the graph does not admit any such exhaustion by rigid sets, in §3.6.

Chapter 2

Rigidity of Graphs of Regions

2.1 Outline of the Chapter

In this chapter we will prove Theorems A, B, and D.

For the most part we will focus on k -separating disc graphs. As remarked in §1.4.1, the choice to study this graph instead of the perhaps more natural graph of curves comes from the fact that the boundary of a disc is uniquely defined, whereas the “side” of a curve is not. Not having to deal with this issue will make our arguments more streamlined, although they can be adapted to the case of the graph of curves without any deep modification.

The general strategy employed for the proof of Theorem B will resemble the arguments used in [B2] and [Mc1], although the technical details will mostly differ. Indeed, the proof of rigidity of the graph $\mathcal{D}_k(\Sigma_m)$ will be an inductive argument on k . For the base case, that is, $k = 2$, we will employ the fact that the automorphism group of this graph is isomorphic to the automorphism group of the 2-separating curve graph (Proposition 2.7.7), that is, the standard curve graph. Given this fact, the base case follows directly from Ivanov’s Theorem (Theorem 1.3.2). We remark that this behaviour is consistent with Ivanov’s Metaconjecture (see page x), as our argument effectively reduces the computation of automorphisms of the k -separating disc graph to the computation of automorphisms of the curve graph.

The inductive step of the argument, that is, the existence of a $\text{Mod}^\pm(\Sigma_m)$ -equivariant injection from $\text{Aut}(\mathcal{D}_k(\Sigma_m))$ to $\text{Aut}(\mathcal{D}_{k-1}(\Sigma_m))$ (Lemma 2.8.12) will be the biggest technical challenge of the proof. Indeed, in order to prove this property we will actually prove that we can reconstruct the entire structure of the graph $\mathcal{D}_{k-1}(\Sigma_m)$ from the purely combinatorial structure of its subgraph $\mathcal{D}_k(\Sigma_m)$. In terms of topology we need to describe $(k-1)$ - and $(m-k+1)$ -punctured discs and their disjointness in terms of discs in $\mathcal{D}_k(\Sigma_m)$ and properties expressible in terms of disjointness only. As we will see in §2.2 the topological type of discs, and complementarity are recognisable (we will more formally define the term later) from the graph of discs, hence reconstructing $(k-1)$ -punctured discs will suffice. Analogously to [B2], the $(k-1)$ -discs will be represented by *surrounding pairs*, whose topological definition is the following.

Definition 2.1.1. Let $k \geq 2$, and $m \geq 2k + 1$. Let A, B be two k -punctured discs on

Σ_m . We say that A, B form a *surrounding pair* $\{A, B\}$ if their intersection $A \cap B$ is a $(k-1)$ -punctured disc.

We recall that discs are regions, and the intersection of regions is well defined up to ambient isotopy, as discussed in §1.2. Moreover, we remark that surrounding pairs are unordered.

Two problems arise when we try to employ the previous definition to represent $(k-1)$ -punctured discs. The first one, which is the easiest one to solve, is the fact that multiple distinct surrounding pairs may have the same intersection, that is, they represent the same $(k-1)$ -punctured disc. It follows that the set of surrounding pairs is not fit for representing $(k-1)$ -punctured discs, but luckily a suitable quotient is (Proposition 2.8.4). The relation which encodes the fact that two surrounding pairs have the same intersection is generated by the equivalence of couples of surrounding pairs $\{A, B\} \sim \{B, C\}$ when the triple $\{A, B, C\}$ is a *surrounding triple*, as in the following definition.

Definition 2.1.2. Let $k \geq 2$, and $m \geq 2k+1$. Let A, B, C be three k -punctured discs on Σ_m . We say that A, B, C form a *surrounding triple* $\{A, B, C\}$ if pairwise they form surrounding pairs and their triple intersection $A \cap B \cap C$ is a $(k-1)$ -punctured disc.

The second issue which we will need to take care of for our argumentative strategy, that is to use surrounding pairs to represent $(k-1)$ -punctured discs, to work will prove more problematic and will require the vast majority of the chapter to be addresses. This issue is the fact that the definition of surrounding pairs and triple we have given so far is based on topological information. Such information is, *a priori*, not recorded in the graph of k -separating discs, which records disjointness but apparently no information regarding any nonempty intersection between discs. Luckily for us, it turns out that such information can actually be extracted from the purely combinatorial structure of the graph. This means that it is possible to characterise surrounding pairs and triple just in terms of combinatorial properties of vertices of the graph. Ultimately, this proves that those pairs and triples are preserved not only by mapping classes but by any graph automorphism (Corollary 2.2.10). This result will be the key piece to reconstruct the graph $\mathcal{D}_{k-1}(\Sigma_m)$ and its automorphisms in terms of graph-theoretical properties of $\mathcal{D}_k(\Sigma_m)$. As a final remark, once the rigidity of the graph is proven, the preservation of surrounding pairs and triples must be true *a posteriori* so, as in many similar rigidity proofs, the main struggle is to extract topological information from the combinatorics of the graph only. The practical way in which we will achieve this goal will be to prove that surrounding pairs and triples appear in a class of hexagons, the *standard alternating hexagons* (Definition 2.2.7), which is preserved under graph automorphisms. This is similar to the approach developed in [B2, §3]: in that case surrounding pairs and triples were characterised in terms of heptagons. Unfortunately, it is impossible to utilise heptagons in the general case, and constructions similar to Bowditch's would involve an increasing number of vertices as k grows, so any universal proof using them would probably be impossible. Luckily, hexagons are suitable for the recognition of surrounding pairs and triples, independently of k and m . Unfortunately, differently from heptagons in [B2, Proposition 3.1], hexagons are no longer rigid as subgraphs (see Definition 1.6.1), that is, they are not unique up to the action of the mapping class group. We will still be

able to prove that the class of standard alternating hexagons can be recognised through purely combinatorial properties (Theorem 2.2.8), and this will suffice for our needs, but this argument will be quite long and technical, and constitutes the core of our proof.

The proof of rigidity for Theorem D will essentially resemble the one of [Mc1, Theorem 3], although with some modifications, as we will use graphs of discs instead of graphs of curves. However, to prove the other direction of the theorem, the sharpness of the result, two original arguments will be given.

In order to reduce the computation of automorphisms of a generic graph of regions $\mathcal{G}_A(\Sigma_m)$ to the ones of the graph of discs $\mathcal{D}_{\nu_A}(\Sigma_m)$ we will study particular subgraphs, called *maximal perfect joins* (see Definition 2.9.19). When the complexity ν_A (Definition 1.4.4) is low enough, and no vertices with holes or cork pairs are present, each of these joins will represent exactly the subgraphs of regions that are contained in either a disc D or its complementary D^c . This will produce a bijection between the set of maximal perfect joins and the set of pairs of complementary discs. Since maximal perfect joins are defined in purely combinatorial terms, they are preserved under graph automorphisms, hence any automorphism of the graph of regions will induce a bijection of the set of pairs of complementary discs. It will turn out that this map actually induces an automorphism of the graph of ν_A -punctured discs (Proposition 2.9.27). This will let us reduce to the computation of the automorphisms of the graph of discs, hence Theorem B can be employed to prove Theorem D .

As for the sharpness of the result, we will first prove that a certain type of graph automorphisms are incompatible with rigidity. These automorphisms are the so-called *exchange automorphisms* (Definition 2.9.2), and consist of automorphisms that permute two vertices, while fixing all the others. We will show directly (Proposition 2.9.3) that, if the complexity is low enough, such automorphisms can never be induced by mapping classes. Afterward, we will characterise this combinatorial obstruction to rigidity in topological terms. Indeed, we will prove that the presence of either vertices with holes (Definition 1.4.5) or cork pairs (Definition 1.4.6) implies the existence of exchange automorphisms. Theorem D proves, *a posteriori*, that the existence of such vertices is not only a sufficient condition for the existence of exchange automorphisms, but it is also necessary. Lastly, we will prove that if the complexity of the graph is too high, with respect to the complexity of the surface, then either the graph is disconnected, or it admits a vertex with holes or it is one of two graphs for which we can directly compute the automorphism group. In all these cases the graph turns out not to be rigid, hence proving the sharpness of Theorem D .

This chapter will be divided as follows. We start with Section 2.3 where we will discuss various proof of the connectedness of graphs of discs, curves and regions. Some of these proofs will be proof *a posteriori*, and for those we will assume the rigidity results, or some intermediate result in their proofs, while others will be completely independent *a priori* proofs.

In Section 2.2 we will prove that some topological properties of discs are recognisable from the combinatorics of the graph $\mathcal{D}_k(\Sigma_m)$. Moreover, we will introduce the notion of alternating hexagons we will work with throughout the following sections.

In Section 2.4 we will study the projection of alternating hexagons in the graph of discs $\mathcal{D}_k(\Sigma_m)$ to the extended Kneser graph $K^*(m, k)$, hence studying the combinatorics

of the distribution of the marked points contained by the vertices.

In Sections 2.5 and 2.6 we will prove the preservation of standard alternating hexagons by graph automorphisms (Theorem 2.2.8). Each section will cover one of the two main different cases, differentiated according to the projection of the hexagons to the Kneser graph: the arguments are similar in spirit, but the technical details differ substantially between them.

In Section 2.7 we will relate the k -separating disc graphs, hence its automorphism group, to the k -separating curve graphs, proving that Theorems A and B are equivalent. This fact will also be needed to reduce the base case of the inductive argument to Ivanov's Theorem.

In Section 2.8 we will formalise the inductive argument we sketched earlier and prove rigidity for k -separating disc graphs, that is, Theorem *B*.

Lastly, in Section 2.9 we will expand on the argument outlined above and prove the characterisation of rigid graphs of regions, that is, Theorem *D*.

2.2 Topological Properties and Alternating Hexagons

In what follows we will from time to time change our notation, and write $\mathcal{D}_k(\Sigma_{k+n+1})$ with $n \geq k$ instead of $\mathcal{D}_k(\Sigma_m)$ with $m \geq 2k+1$: the quantities m and $k+n+1$ will alternatively play the same role. This will simplify the notation in many statements and arguments.

Our main technical result, Corollary 2.2.10, will require that $k \geq 3$ (we will point out why in more detail after the statement of Theorem 2.2.8), but this will be enough for our needs, as Lemma 2.8.12 would still require $k \geq 3$. This is the reason why, unless otherwise stated, from now on we will always assume $k \geq 3$. However, many intermediate statements will still hold for $k \geq 2$, hence we will state them with this weaker bound, in order for their proofs to appear more natural.

The goal of this section is to introduce the hexagons we want to study in the graph of k -separating hexagons, in order to recognise surrounding pairs and triples. In the first part of this section we will show how some topological properties of discs can be detected just by looking at the combinatorial structure of the graph.

We say that a certain property, or a set of vertices, is *recognisable* in the graph to mean that it is preserved under graph automorphisms.

We will start by proving that inclusion between two discs up to isotopy (which is well defined, see §1.2.2) is a recognisable property in the graph. We can actually characterise it in the following way.

Proposition 2.2.1. *Let $k \geq 2$ and $n \geq k$. Let A, B be two distinct vertices of the graph $\mathcal{D}_k(\Sigma_{k+n+1})$. Then $A \subseteq B$ if and only if every vertex of $\mathcal{D}_k(\Sigma_m)$ adjacent to B is adjacent to A . In particular inclusion between discs is a recognisable property.*

We remark that the above condition can be expressed in more graph-theoretical terms as $\text{link}(B) \subseteq \text{link}(A)$.

Proof of Proposition 2.2.1. If $A \subseteq B$ it is immediate to observe that every disc disjoint from B must be disjoint from A , too.

For the converse let A, B be vertices such that $\text{link}(B) \subseteq \text{link}(A)$. From this it follows that B^c is adjacent to A , since it is adjacent to B . Hence, we have that $A \subseteq (B^c)^c = B$, which was our claim. \square

Now that we have proven that inclusion between discs is recognisable, we can go a step further and prove that the topological type of a disk, that is, the number of marked points in it, is preserved under graph automorphisms. Since every graph automorphism induced by the extended mapping class group preserves the topological types of discs, the topological types being recognisable is a necessary condition for the graph to be rigid.

In order to describe the topological types we will use maximal chains of nested discs, which we will now define.

Definition 2.2.2. Let $k \geq 2$ and $n \geq k$. A *chain* of discs is a collection of discs $\{D_0, \dots, D_j\} \subset \mathcal{D}_k(\Sigma_{k+n+1})$ such that $D_0 \subsetneq D_1 \subsetneq \dots \subsetneq D_j$.

A chain is said to be *maximal* if it is not properly contained in any other chain.

If two discs A, B contain the same number of marked points and $A \subseteq B$ then $B \setminus A$ is an annulus, hence A and B are isotopic. It follows that the number of marked points contained by discs in a chain must be strictly increasing. Hence it is immediate to observe that there cannot exist infinite chains.

Since inclusion is recognisable in the graph of discs due to Proposition 2.2.1, chains of discs are recognisable too.

Given a disc in $\mathcal{D}_k(\Sigma_{k+n+1})$, there exists a maximal chain containing it. The following proposition gives us the characterisation of topological types of discs in terms of maximal chains and, alongside the previous observations, proves that topological types are recognisable.

Proposition 2.2.3. *Let $k \geq 2$ and $n \geq k$. Let $D_0 \subsetneq \cdots \subsetneq D_j$ be a maximal chain of discs in $\mathcal{D}_k(\Sigma_{k+n+1})$. Then $j = n - k + 1$ and D_i is a $(k + i)$ -punctured disc.*

Let D be a vertex of $\mathcal{D}_k(\Sigma_{k+n+1})$. Then D is a $(k + i)$ -punctured disc if and only if there exists a maximal chain $D_0 \subsetneq \cdots \subsetneq D_{n-k+1}$ such that $D = D_i$. In particular, the topological type of discs is recognisable in the graph.

Proof. We proceed by induction on i . First, we claim that the disc D_0 is k -punctured. If not, there would exist a k -punctured disc D properly contained in D_0 , and the chain $D \subsetneq D_0 \subsetneq \cdots \subsetneq D_j$ would be strictly longer than the chain we started with, contradicting maximality. Our claim is proven.

If the disc D_i is $(k + i)$ -punctured then the subsequent disc D_{i+1} must contain at least $k + i + 1$ marked points, otherwise we would have $D_i = D_{i+1}$. We claim that the disc D_{i+1} contains exactly $k + i + 1$ marked points. If not, the disc D_{i+1} would contain at least two marked points not in D_i , let us denote them with p, q . Then there would exist a disc $D \subseteq D_{i+1}$ containing p but not q . From this it would follow that $D_i \subsetneq D \subsetneq D_{i+1}$, hence the chain would not be maximal. We have reached a contradiction, proving the inductive step of the argument, hence the proposition.

In order to prove the existence of the disc D the following construction can be employed. Let $a \subset D_{i+1} \setminus D_i$ be a simple arc with one endpoint on D_i and p as the other one, not containing q . Let $D \subseteq D_{i+1}$ be a regular neighbourhood of $D_i \cup a$ small enough to be a punctured disc not containing q . \square

The following proposition, whose proof is immediate, alongside the previous result, also proves that complementary discs are recognisable in the graph.

Proposition 2.2.4. *Let $k \geq 2$ and $n \geq k$. Let h be such that $k \leq h \leq m - h$. Let A be a h -punctured disc in $\mathcal{D}_k(\Sigma_m)$. Then there exists a unique $(m - h)$ -punctured disc B which is disjoint from A . This disc is the complementary disc of A . In particular, (unordered) pairs of complementary discs are recognisable in the graph of discs.*

We are now ready to give the definition of the objects we will be focusing on for the most part of the chapter: *alternating hexagons*. First, we begin with a graph-theoretical definition.

Definition 2.2.5. Let Γ be a graph. Let $h \geq 3$. Let C_h be the 2-regular connected graph with h vertices. A h -cycle in Γ is a map $f : C_h \rightarrow \Gamma$.

A cycle is said to be *embedded* if it is injective.

A *hexagon* in Γ is an embedded 6-cycle.

We can now define alternating hexagons in the graph of k -separating discs.

Definition 2.2.6. Let $k \geq 2$ and $n \geq k$. Let $H = (D_1, \dots, D_6)$ be a 6-cycle in the graph $\mathcal{D}_k(\Sigma_{k+n+1})$. We say that such a cycle is *alternating* if all the vertices are either k -punctured or n -punctured discs, and a vertex D_i is k -punctured if and only the adjacent vertices $D_{i\pm 1}$ are n -punctured, where subscripts are to be considered modulo 6.

An *alternating hexagon* is an alternating embedded 6-cycle.

Thanks to Proposition 2.2.3 and Proposition 2.2.1 every property in the previous definition is recognisable in the graph, so the definition is purely graph-theoretical.

An alternating hexagon may not be isometrically embedded, that is, it can admit diagonals in $\mathcal{D}_k(\Sigma_{k+n+1})$.

We will now present a class of alternating hexagons, which constitute the simplest possible example, but will play a pivotal role in our argument by letting us recognise surrounding pairs and triples. We will call these hexagons *standard*. A *standard alternating hexagon* can be defined as a hexagon which coincide with the one in Figure 2.1, up to action of a mapping class. For clarity, in that picture, as in all the others to follow, we have only drawn the boundaries of the discs in the hexagons. We do, however, also give a formal definition of such hexagons, which, although cumbersome, does not rely on pictures.

Definition 2.2.7. Let $k \geq 2$, $n \geq k$. Let T be the 1-skeleton of a standard 2-simplex, with vertices v_1, v_2, v_3 . Let now $ST \cong S_{0,0}^0$ be the suspension of T , endowed with the metric induced by every 2-simplex being an equilateral Euclidean triangle of edglength 1. We will consider the triangle T as naturally embedded into \mathbb{R}^2 , hence the suspension as a naturally (but not isometrically) embedded subset $ST \subset \mathbb{R}^3$. We will call $T \subseteq ST$ the equator and the two collapsed points in the suspension north and south poles. Let l_1, l_5, l_3 be the edges from the north pole to vertices v_1, v_2, v_3 respectively; and let l_4, l_2, l_6 be the edges from the south pole to v_1, v_2, v_3 respectively. Let now D_i be the $\frac{1}{3}$ -regular closed neighbourhood of l_i (see § 1.2.2). Let us now mark $k-1$ points in the ball of radius $\frac{1}{6}$ around the north pole, mark $n-1$ points in the ball of radius $\frac{1}{6}$ around the south pole, and mark each vertex of T . The surface we obtain is homeomorphic to Σ_{k+n+1} . Under the identification given by the aforementioned homeomorphism the discs (D_1, \dots, D_6) form a alternating hexagon H in $\mathcal{D}_k(\Sigma_{k+n+1})$.

An alternating hexagon in $\mathcal{D}_k(\Sigma_{k+n+1})$ is said to be *standard* if it belongs to the orbit of H under the action of $\text{Mod}^\pm(\Sigma_{k+n+1})$.

Let us remark that standard alternating hexagon all look the same, that is, as in Figure 2.1, up to the action of the extended mapping class group.

The following theorem, concerning the recognition of standard alternating hexagons, will be the main technical result of the chapter, and the key to the proof of Theorem B. Its proof will be split in two different cases and dealt with in Sections 2.5 and 2.6.

Theorem 2.2.8. *Let $k \geq 3$ and $m \geq 2k+1$. Then, the set of standard alternating hexagons in $\mathcal{D}_k(\Sigma_m)$ is preserved by graph automorphisms.*

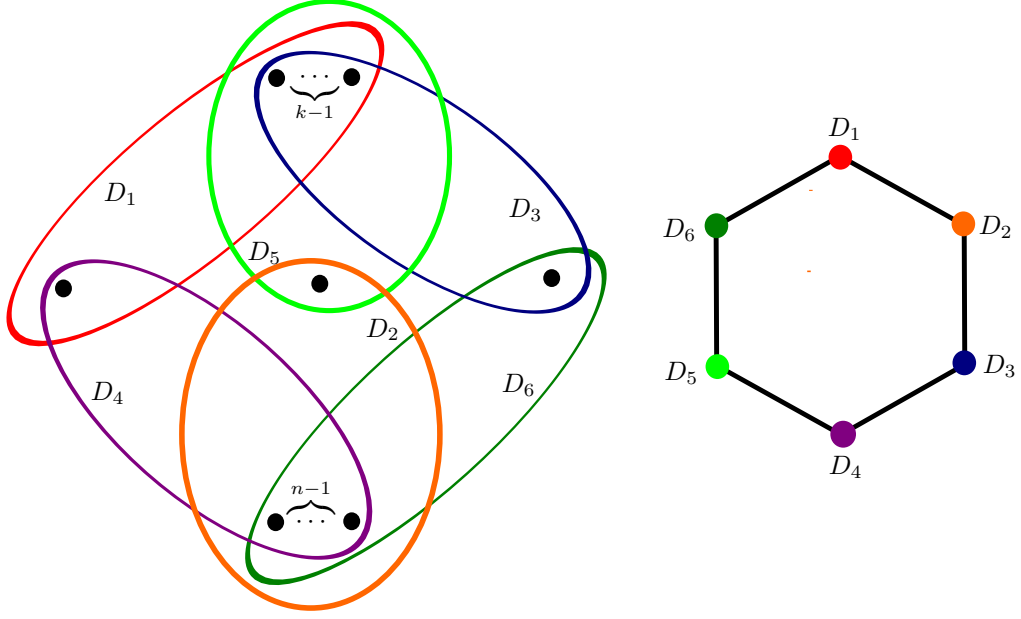


Figure 2.1: Standard alternating hexagon.

The need for the hypothesis $k \geq 3$ in the previous statement comes from one single intermediate result in the proof, namely Proposition 2.6.4. Indeed, this result fails for $k = 2$ (we will give a counterexample when we prove it in §2.6). It is very likely that the proof can be adapted to the case $k = 2$, but this is completely unnecessary to prove rigidity as Lemma 2.8.12 would still require that $k \geq 3$ anyway.

In order to be able to use the previous result, we now want to prove that, using standard alternating hexagons, we can recognise surrounding pairs and triples in the graph of discs.

Proposition 2.2.9. *Let $k \geq 2$ and $m \geq 2k + 1$. Let A, B, C be distinct k -punctured discs in Σ_m . Then A and B form a surrounding pair if and only if they are at distance 2 in a standard alternating hexagon in $\mathcal{D}_k(\Sigma_m)$*

Moreover, A, B, C form a surrounding triple if and only if they are three pairwise non-adjacent vertices in a standard alternating hexagon.

Proof. The fact that k -punctured discs at distance 2 in a standard alternating hexagon form surrounding pairs and triples follows immediately from Definition 2.2.7, and can be visualised in Figure 2.1.

For the converse let $\{A, B, C\}$ be a surrounding triple. The argument for pairs can be deduced from this by only considering the constructions related to A and B . Let $O = A \cap B \cap C$, which is a $(k - 1)$ -punctured disc. Let P_A (resp. P_B, P_C) be $A \cap O^c$ (resp. $B \cap O^c, C \cap O^c$), which is a once-punctured disc only containing a single marked point we will denote with p_A (resp. p_B, p_C). In P_A (resp. P_B, P_C) there exists a unique arc a (resp. b, c) with p_A (resp. p_B, p_C) as the first endpoints and the second one to ∂O , up to isotopy which keeps the second endpoint on ∂O at all times.

Since the action of the homeomorphism group is transitive among discs of the same topological type (see [FM, p.37]), the disc O can be identified with the region $O' =$

$D_1 \cap D_3 \cap D_5$ in Definition 2.2.7. Moreover, up to the action of a homeomorphism which fixes O' pointwise, the arc a, b, c can be simultaneously identified with disjoint arcs inside D_1, D_3 and D_5 , with one endpoint on $\partial(O')$ and with v_1, v_3 or v_2 , respectively, as the other endpoint. Under this identification, the disc A (resp. B, C) is isotopic to a small regular neighbourhood of $O' \cup a$ (resp. $O' \cup b, O' \cup c$) which contains k -marked points, which is isotopic to D_1 (resp. D_3, D_5). It follows that, up to the action of a mapping class, we have $A = D_1, B = D_3$, and $C = D_5$. Hence A, B, C are three vertices pairwise at distance 2 in a standard hexagon. \square

Theorem 2.2.8 and Proposition 2.2.9 immediately prove the following corollary, which will be the key result that, in Section 5, will let us reconstruct the $(k-1)$ -separating disc graph from the k -separating one.

Corollary 2.2.10. *Let $k \geq 3$ and $m \geq 2k + 1$. Then, the set of surrounding pairs and the set of surrounding triples in $\mathcal{D}_k(\Sigma_m)$ are preserved by graph automorphisms.*

2.3 Connectedness of the Graphs

In this section we will discuss and prove the following connectedness results for the graphs of discs $\mathcal{D}_k(\Sigma_m)$, graph of curves $\mathcal{C}_k(\Sigma_m)$ and more generally graphs of regions $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$. This is not only an interesting argument *per se*, but also a way to familiarise with the type of arguments and the techniques that will be heavily employed in the following sections. Beacuse we are willingly anticipating this material, we will use some results that will only be proven by the end of the section. However, the connectedness results we achieve in this section will never be employed to prove those results. We begin with the result for graphs of discs.

Theorem 2.3.1. *Let $k \geq 2$ and $m \geq 2k + 1$. Then the k -separating disc graph $\mathcal{D}_k(\Sigma_m)$ is connected.*

Since every disc which is a vertex of $\mathcal{D}_k(\Sigma_m)$ is disjoint from a k -punctured discs, the previous result is indeed equivalent to proving that the strict graph of k -punctured discs $\mathcal{D}_{(k)}(\Sigma_m)$ (see Definition 1.4.2) is connected.

The result for graphs of curves is the following.

Theorem 2.3.2. *Let $k \geq 2$ and $m \geq 2k + 1$. Then the k -separating curve graph $\mathcal{C}_k(\Sigma_m)$ is connected.*

Since every curve which is a vertex of $\mathcal{C}_k(\Sigma_m)$ is disjoint from a k -separating curve, the previous result is indeed equivalent to proving that the strict k -separating curve graph $\mathcal{C}_{(k)}(\Sigma_m)$ (see Definition 1.3.6) is connected. We recall that, when $m \geq 2k + 1$, such graph is naturally isomorphic to the graph $\mathcal{D}_{(k)}(\Sigma_m)$. It follows that Theorem 2.3.2 is equivalent to Theorem 2.3.1.

Lastly, the result for graphs of regions is the following. For the definition of the complexity of a graph of region we refer to Definition 1.4.4, whereas the definition of vertices with holes is given in Definition 1.4.5.

Theorem 2.3.3. *Let $A \subseteq \mathcal{R}(\Sigma_m)$ such that $m \geq 2\nu_{\mathcal{A}} + 1$ and such that the associated graph of region $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ does not contain vertices with holes. Then $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is connected.*

When there are no vertices with holes, it follows from Lemma 2.9.12 that the set of minimal complexity elements of $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ contains either the set of $\nu_{\mathcal{A}}$ -punctured discs or the set of annular surfaces representing $\nu_{\mathcal{A}}$ -separating curves (possibly both). It follows that the graph induced by the minimal complexity elements contains an isomorphic copy of either the graph of k -punctured discs $\mathcal{D}_{(\nu_{\mathcal{A}})}(\Sigma_m)$ (see Definition 1.4.2 for definition) or the strict k -separating curve graph $\mathcal{C}_{(\nu_{\mathcal{A}})}(\Sigma_m)$ (see Definition 1.3.6 for its definition). Every region of \mathcal{A} is adjacent to both a $\nu_{\mathcal{A}}$ -punctured discs and an annular surfaces representing $\nu_{\mathcal{A}}$ -separating curve. Hence, connectedness of $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is then equivalent to the connectedness of one of the two previous graphs (which are isomorphic to each other). It follows that the three previously stated results are equivalent.

It can be noticed that, in the previous theorems, the hypotheses needed for connect-edness recall the ones required for rigidity. Indeed, connectedness of the graphs can be proven as a direct consequence of rigidity, as Proposition 2.3.5 shows in the general case of graphs of regions. Moreover, we observe that, thanks to Proposition 2.9.1 and

Lemma 2.9.11 for the graphs of regions without vertices with holes and corks pairs connectedness is equivalent to rigidity. In the case of graphs with vertices with holes or cork pairs, however, this equivalence is no longer true: indeed, in those cases the obstruction to rigidity does not come from the existence of multiple connected components, but from the existence of exchange automorphisms.

In order to prove that rigidity implies connectedness, we give the following general graph-theoretical result.

Proposition 2.3.4. *Let Γ be a graph such that the following hold:*

1. *There exists a $\text{Aut}(\Gamma)$ -orbit with at least three elements;*
2. *Every connected component of Γ has trivial pointwise stabiliser for the action of $\text{Aut}(\Gamma)$.*

Then the graph Γ is connected.

Proof. Let us argue by contradiction and assume that Γ is not connected. Let $C_1, C_2 \subseteq \Gamma$ be two different connected components. Let $v \in \mathcal{V}(\Gamma)$ be a vertex whose $\text{Aut}(\Gamma)$ -orbit has at least three elements. It follows that either such orbit intersects one component at least twice, or there exists a third connected component C_3 .

We will first deal with the case of the $\text{Aut}(\Gamma)$ -orbit of v intersecting one component at least twice. Without loss of generality we can assume that $v \in C_1$, and that there exists $\varphi \in \text{Aut}(\Gamma)$ such that $\varphi(v) \in C_1$ and $\varphi(v) \neq v$, since the orbit of v intersects C_1 at least one distinct vertex. In particular we have $\varphi \neq \text{Id}_\Gamma$. Since graph automorphisms preserve connected components it follows that $\varphi(C_1) = C_1$. We can now define a graph automorphism of Γ as

$$\psi(w) = \begin{cases} \varphi(w) & \text{if } w \in C_1 \\ w & \text{otherwise.} \end{cases}$$

This automorphism ψ fixes the component C_2 pointwise, hence $\psi = \text{Id}_\Gamma$, which is a contradiction.

We can now move to the case in which the graph Γ admits three connected components C_1, C_2, C_3 , all intersecting the $\text{Aut}(\Gamma)$ -orbit of v . Without loss of generality, we can assume that $v \in C_1$ and that $\varphi \in \text{Aut}(\Gamma)$ is a graph automorphism such that we have $\varphi(v) \in C_2$. In particular we have $\varphi \neq \text{Id}_\Gamma$. Since graph automorphisms preserve connected components it follows that $\varphi(C_1) = C_2$. We can now define a graph automorphism of Γ as

$$\psi(w) = \begin{cases} \varphi(w) & \text{if } w \in C_1 \\ \varphi^{-1}(w) & \text{if } w \in C_2 \\ w & \text{otherwise.} \end{cases}$$

This automorphism ψ fixes the component C_3 pointwise, hence $\psi = \text{Id}_\Gamma$, which is a contradiction. The proposition is proven. \square

We will now prove that, for graphs of regions, rigidity implies connectedness.

Proposition 2.3.5. *Let $\mathcal{A} \subseteq \mathcal{R}(\Sigma_m)$ be such that the associated graph of regions $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is rigid, that is $\text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m)) = \text{Mod}^\pm(\Sigma_m)$. Then the graph $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is connected.*

Proof. Our goal is to apply Proposition 2.3.4. Condition 1 is satisfied, since the orbit of every vertex is infinite.

Since $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is rigid, Corollary 2.9.15 proves that the graph does not contain either any vertex with holes or any cork pair, and that $m \geq 2\nu_{\mathcal{A}} + 1$.

Every connected component of $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ contains an element of minimal complexity. Indeed, every vertex of $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ admits at least a complementary disc, and since it has no holes, every such disc contains at least $\nu_{\mathcal{A}}$ marked points. Thanks to Lemma 2.9.12, it contains either a $\nu_{\mathcal{A}}$ -punctured disc or an annular surfaces representing $\nu_{\mathcal{A}}$ -separating curves, which is disjoint from the original vertex. We will only focus on the case where the set of minimal complexity elements of $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ contains the set of $\nu_{\mathcal{A}}$ -punctured discs, the other being extremely similar.

Let now C be a connected component of $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$: we want to compute its pointwise stabiliser. Let $f \in \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$ such that $f|_C = \text{Id}_C$: since the graph is rigid then $f \in \text{Mod}^{\pm}(\Sigma_m)$ is an extended mapping class. Let $A \in C$ be a $\nu_{\mathcal{A}}$ punctured disc. It not only holds that $f(A) = A$ but that for every $\nu_{\mathcal{A}}$ -punctured disc $B \subseteq A^c$ we have $f(B) = B$. It follows that the homeomorphism f is supported in A , up to isotopy. Similarly, let D be a ν -punctured discs disjoint from A , which exists as $m \geq 2\nu_{\mathcal{A}} + 1$: from a similar argument it follows that f is supported in D , up to isotopy. We can deduce that f is supported in $A \cap D = \emptyset$, up to isotopy, hence isotopic to the identity. We have proven that the pointwise stabiliser of every connected component of $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is trivial. We have checked Condition 2, so we can apply Proposition 2.3.4, proving that $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is connected. \square

The proof of connectedness which we have just given is completely *a posteriori*: indeed, the connectedness of the graphs has never been used throughout the entire proof of rigidity.

Another way in which connectedness can be proven is by using the tools introduced in our proof of rigidity, in particular using an inductive argument revolving around the inclusion $\mathcal{D}_k(\Sigma_m) \subseteq \mathcal{D}_{k-1}(\Sigma_m)$, Lemma 2.8.11 and the definition of edges in Definition 2.8.10. Indeed, it is not hard to prove that if two vertices representing h -punctured discs with $k \leq h \leq m - h$ are joined by a path in $\mathcal{D}_{k-1}(\Sigma_m)$ then they are also disjoint by a path avoiding vertices representing either $(k - 1)$ or $(m - k + 1)$ -punctured discs. Hence, the path is contained in $\mathcal{D}_k(\Sigma_m) \subset \mathcal{D}_{k-1}(\Sigma_m)$, hence connectedness of $\mathcal{D}_{k-1}(\Sigma_m)$ implies that also $\mathcal{D}_k(\Sigma_m)$ is connected. The base case is then given by connectedness of the graph $\mathcal{D}_2(\Sigma_m)$, which is equivalent to the graph $\mathcal{C}_2(\Sigma_m)$ being connected. The last mentioned graph is the usual curve graph, and its connectedness has been well-known since its definition: see [Harv1, Proposition 2].

The proofs of connectedness we have presented so far, although perfectly valid, have the issue of essentially being *a posteriori*, and relying on the most part, if not at all, of the work needed to prove rigidity. It turns out that it is possible to give *a priori* direct proofs of the connectedness of the graphs. The first proof of this kind we give is the following, which has the upside of being quite short, although not particularly constructive, as it uses an explicit set of generators for the mapping class group, finding which is a far from trivial result. Such an approach will make use the following lemma, due to Andrew Putman, the easy proof of which is left to the reader, see [Pu, Lemma 2.1].

Lemma 2.3.6. *Let Γ be a graph. Let G be a group acting on Γ by graph automorphisms. Let H be a set of generators for the group G such that $H^{-1} = H$. Suppose there exists a vertex $v \in \mathcal{V}(\Gamma)$ such that the following hold:*

1. *The orbit $G \cdot v$ intersects every connected component of Γ ;*
2. *For every $h \in H$ there is a path in Γ between v and $h(v)$.*

Then the graph Γ is connected.

We will now prove Theorem 2.3.1 through the use of an explicit set of generators for the mapping class group.

Before we begin the proof we recall that, given a set X , its *disjoint union* with itself is the set $X \sqcup X = (X \times \{0\}) \cup (X \times \{1\})$.

Proof of Theorem 2.3.1 via a generating set of $\text{Mod}(\Sigma_m)$. Throughout the entire proof we will identify the m -punctured sphere as the doubling of a regular m -gon, similarly to what we will consistently do in §3.4 and §3.5.

Let \mathcal{P}_m be a regular polygon with m -sides in \mathbb{R}^2 , with vertices $\{v_1, \dots, v_m\}$ cyclically ordered. Let l_i be the side between v_i and v_{i+1} .

Let $S = \mathcal{P}_m \sqcup \mathcal{P}_m / \sim$ be the doubling of the m -gon, where $(x, 0) \sim (y, 1)$ if and only if $x = y \in \partial \mathcal{P}_m$. The marked surface $(S, \{v_1, \dots, v_m\})$ is a m -holed sphere.

Let D be the k -punctured disc isotopic to a small regular closed neighbourhood of $l_1 \cup \dots \cup l_{k-1}$. Let γ_i be the 2-separating curve isotopic to the boundary of a small regular neighbourhood of l_i . Let H_i be the (right) half-Dehn twist around γ_i , see [FM, §9.1.3]. Let $H = \{H_i^{\pm 1}\}$: the discussion in [FM, § 4.4.4] proves that the collection H generates the entire mapping class group $\text{Mod}(\Sigma_m)$.

We will now apply Lemma 2.3.6 with $\Gamma = \mathcal{D}_{(k)}(\Sigma_m)$ (we have already observed how proving connectedness for this graph is enough), $G = \text{Mod}(\Sigma_m)$, H as defined above, and $v = D$.

Condition 1 follows immediately from the transitivity of the action, that is from fact that for every disc $A \in \mathcal{V}(\mathcal{D}_{(k)}(\Sigma_m))$, which contains exactly k marked points as D does, there exists a mapping class $f \in \text{Mod}(\Sigma_m)$ such that $f(D) = A$, hence A itself belongs to the orbit of D .

In order to prove Condition 2, we first observe that for every $j \notin \{k, m\}$ the curve γ_j is disjoint from ∂D , hence $H_j^{\pm 1}(D) = D$. The following constructions are pictured in Figure 2.2. Let B be the k -punctured disc isotopic to a small regular closed neighbourhood of $l_{k+2} \cup \dots \cup l_{2k+1}$: since $m \geq 2k+1$ this disc is disjoint from D . Similarly, let C be the k -punctured disc isotopic to a small regular closed neighbourhood of $l_{k+1} \cup \dots \cup l_{2k}$. The disc B is disjoint both from $H_k^{\pm 1}(D)$, which is contained in a small regular neighbourhood of $l_1 \cup \dots \cup l_k$, and B is also disjoint from D . It follows that D and $H_k^{\pm 1}(D)$ are joined by a path in the graph. Similarly, the disc C is disjoint both from $H_m^{\pm 1}(D)$ and from D , hence D and $H_m^{\pm 1}(D)$ are joined by a path in the graph. It follows that the condition is satisfied and hence the graph $\mathcal{D}_{(k)}(\Sigma_m)$ is connected. \square

Lastly, we would like to provide a proof of connectedness which is still *a priori*, and very constructive, as it is based on a series of surgery operation between discs. This is

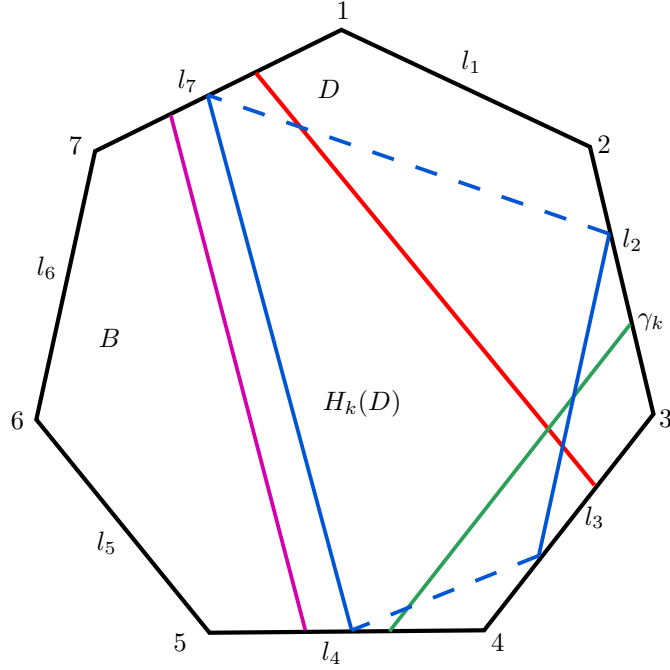


Figure 2.2: The figure is for $k = 3$ and $m = 7$. The red curve is ∂D , the green curve is γ_k , the blue curve is $\partial(H_k(D))$ (the dotted segments are to be intended on the “back” of the heptagon), and the purple curve is ∂B .

an inductive argument somewhat similar to usual proof of connectedness for the curve graph, see [FM, Theorem 4.4] for an example of such a proof.

Direct proof of Theorem 2.3.1. Once again we will only prove that the graph of k -punctured discs $\mathcal{D}_{(k)}(\Sigma_m)$ is connected, which is sufficient, as we have already remarked.

Let A, B two vertices of $\mathcal{D}_{(k)}(\Sigma_m)$, that is two k -punctured discs of Σ_m . Our argument will involve multiple inductions. We start by arguing by induction on $i(\partial A, \partial B)$. If such intersection number is zero then either $A \subseteq B$ or $A \subseteq B^c$. Since both discs contain exactly k marked points in the former case it follows that $A = B$, while in the latter it follows that A and B are disjoint, hence adjacent. In both cases A and B belong to the same connected component.

Let us now assume that $i(\partial A, \partial B) = 2$. We proceed by induction on $|(A \Delta B) \cap \Pi|$, where $A \Delta B = (A \setminus B) \cup (B \setminus A)$ is the symmetric difference, and Π the set of marked points on the surface. We remark that such cardinality is always even. Since $i(\partial A, \partial B) = 2$ then both $A \cap B^c$ and $B \cap A^c$ are connected subspaces, and in particular they are bigons. If it held $|(A \Delta B) \cap \Pi| = 0$ then the aforementioned components would be empty bigons, contradicting minimal position.

If $|(A \Delta B) \cap \Pi| = 2$ then the region $A \cup B$ is a disc containing exactly $k + 1$ marked points. It follows that the region $D = (A \cup B)^c$ is a disc containing $m - (k + 1) \geq k$ marked points, hence there exists a k -separating disc in D , which is then adjacent to both A and B , which then belong to the same connected component.

Let us now assume that $|(A \Delta B) \cap \Pi| > 2$. The following construction is pictured in Figure 2.3. Let $p \in A \cap B^c$ be a marked point and let $q \in B \cap A^c$ be a marked point. Let a

be an arc inside A with one endpoint on $\partial B \setminus A$ and the other on p . Similarly, let b be an arc inside B with one endpoint on $\partial B \setminus A$ and the other on q . Let $\overline{\mathcal{N}(a)}$ be the closure of a small regular neighbourhood of a , not intersecting either any marked point other than p . Similarly, let $\mathcal{N}(b)$ be a small open regular neighbourhood of b , not intersecting $B \cap A^c$ and not containing any marked point other than q . Let $D = (B \cup \overline{\mathcal{N}(a)}) \setminus \mathcal{N}(b)$. This is a k -punctured disc whose boundary intersects ∂A and ∂B at most twice. Moreover $|(A \Delta C) \cap \Pi| = |(A \Delta B) \cap \Pi| - 2$ and $|(B \Delta C) \cap \Pi| = 2$. We can now apply our inductive hypothesis to prove that A and B belong to the same connected component of C .

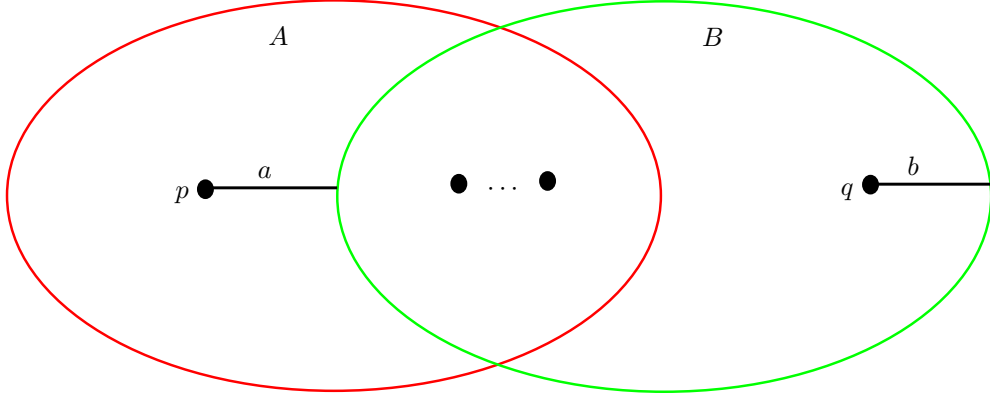


Figure 2.3: For clarity we are omitting some of the marked points.

We can now move to the case of discs A, B such that $i(\partial A, \partial B) > 2$. A picture of this case is provided in Figure 2.4. First, assume that the intersection $A \cap B$ contains a punctured bigon, that is a connected component whose boundary is composed of a single arc of ∂A and a single arc of ∂B , and which contains at least one marked point. Let D be such a bigon, and let $h \geq 1$ be the number of marked points contained in D . We proceed by induction on $h \geq 1$. Let p be a marked point in D and let q be a marked point in $A^c \cap B^c$. Let a (resp. b) be a simple arcs in D (resp. $A^c \cap B^c$) with one endpoint on $\partial B \setminus \partial A$ and the other on p (resp. q). Let $\mathcal{N}(a)$ be a small open regular neighbourhood of the arcs a , contained in A and non containing any marked point apart from p . Let $\overline{\mathcal{N}(b)}$ be a small closed regular neighbourhood of the arcs b , non containing any marked point apart from q , disjoint from A and whose intersection with B is connected. The region $C = (B \setminus \mathcal{N}(a)) \cup \overline{\mathcal{N}(b)}$ is a k -punctured disc. Moreover if $h = 1$ then $i(\partial A, \partial C) = i(\partial A, \partial B) - 2$, otherwise $i(\partial A, \partial C) = i(\partial A, \partial B)$, and C and A form a bigon containing $h - 1 \geq 1$ marked points. In both cases our inductive hypotheses prove that A and C belong to the same connected component. Moreover $i(\partial B, \partial C) = 2$, hence the previous case prove that B belongs to the same connected component of C , henceforth to the same component of A .

We are now left to the case where $A \cap B$ does not contain any punctured bigon. A picture of the following constructions is provided in Figure 2.5. Once again we argue by induction on the number of marked points contained in the symmetric intersection $|(A \Delta B) \cap \Pi|$. There always exists an outermost arc among the connected components of $\partial B \cap A$, that is an arc c such that the interior of one of the components of $A \setminus c$ does not

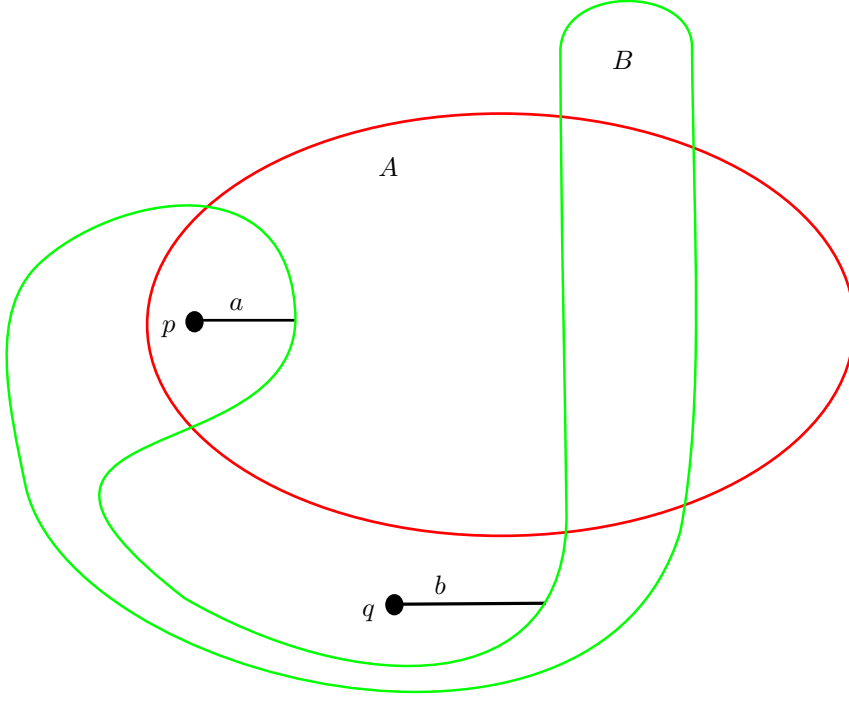


Figure 2.4: For clarity we are omitting some of the marked points.

intersect ∂B . Due to this fact, the absence of bigons in $A \cap B$ implies that $A \cap B^c$ contains at least one punctured bigon. In particular it follows that $A \cap B^c$ contains at least one marked point, hence we can assume that $|(A \Delta B) \cap \Pi| > 0$ without loss of generality. Let p be a marked point in $A \cap B^c$ and let q be a marked point in $B \cap A^c$. Let a (resp. b) be a simple arc in $A \cap B^c$ (resp. $B \cap A^c$) with one endpoint on $\partial B \setminus A$ and the other on p (resp. q). Let $\overline{\mathcal{N}(a)}$ be a small closed regular neighbourhood of the arcs a , non containing any marked point apart from p and whose intersection with B is connected.. Let $\mathcal{N}(b)$ be a small closed regular neighbourhood of the arcs b , non containing any marked point apart from p and disjoint from A . The region $C = (B \setminus \mathcal{N}(b)) \cup \overline{\mathcal{N}(a)}$ is a k -punctured disc. we have $i(\partial A, \partial C) \leq i(\partial A, \partial B)$ and $|(A \Delta C) \cap \Pi| = |(A \Delta B) \cap \Pi| - 2$, hence our inductive hypotheses prove that A and C belong to the same connected component. Moreover, $i(\partial B, \partial C) = 2$, hence a previous argument proves that B lies in the same connected component of C . In particular it follows that A and B belong to the same component, and the proof is complete. \square

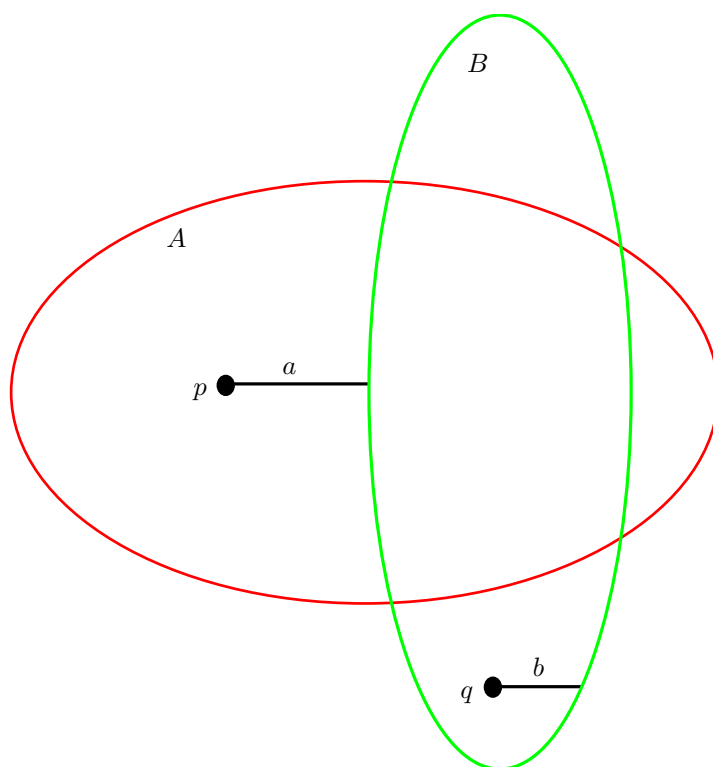


Figure 2.5: For clarity we are omitting some of the marked points.

2.4 Projections to Extended Kneser Graphs

In this section we will study the combinatorial configurations of marked points contained in the discs forming an alternating hexagon. In order to study this behaviour, we will make use of extended Kneser graphs $K^*(m, k)$ (Definition 1.5.2) and the natural map

$$\pi : \mathcal{D}_k(\Sigma_m) \longrightarrow K^*(m, k),$$

as described in §1.5.2. We recall that the projection of a disc $\pi(D)$ is the set of marked points inside the disc D , that is $\pi(D) = D \cap \Pi$, where Π denotes the set of marked points of Σ_m . Moreover, two projections are joined by an edge if they are disjoint, that is if the two discs do not share any marked point.

First, we observe the following very simple lemma, which will be useful later.

Lemma 2.4.1. *Let $k \geq 2$ and $m \geq 2k + 1$. Let h be such that $k \leq h \leq m - k$. Let $v \neq w$ be two distinct h -vertices of the extended Kneser graph $K^*(m, k)$ such that there exists a $(m - h - 1)$ -vertex z adjacent to both. Then $|v \cap w| = h - 1$.*

Proof. Since $v, w \subseteq z^c$ and $|z^c| = h + 1$ it follows that $|v \cup w| \leq h + 1$. Since $v \neq w$ it follows that $|v \cup w| > h$ so $|v \cup w| = h + 1$. The inclusion-exclusion principle yields

$$h + 1 = |v \cup w| = |v| + |w| - |v \cap w| = 2h - |v \cap w|$$

hence $|v \cap w| = h - 1$. □

Since we will be studying the projections of alternating hexagons in the graph of discs, we will now define the analogous alternating property in the context of extended Kneser graphs. Indeed, the following definition is analogous in spirit to Definition 2.2.6.

Definition 2.4.2. Let $k \geq 2$, $n \geq k$. Let $h \in \mathbb{N}$ even with $h \geq 4$. Let $H = (v_1, \dots, v_h)$ be a h -cycle in the extended Kneser graph $K^*(k + n + 1, k)$. Such h -cycle is said to be *alternating* if the vertices are either k -vertices or n -vertices, and the vertex v_i is a k -vertex if and only if its neighbours $v_{i \pm 1}$ are n -vertices, where subscripts are considered modulo h .

An *alternating hexagon* in the extended Kneser graph $K^*(k + n + 1, k)$ is an embedded alternating 6-cycle.

It follows immediately that an alternating h -cycle in $\mathcal{D}_k(\Sigma_m)$ projects to an alternating h -cycle in $K^*(m, k)$. However, even if the original cycle were embedded, the projection cycle may not be. In particular, the projection of an alternating hexagon in the graph of discs may not be a hexagon in the extended Kneser graph. For an example of such a non-embedded projection see Figure 2.6.

We will now study the structure of alternating hexagons in the extended Kneser graph $K^*(m, k)$, with $m \geq 2k + 1$. It will turn out that there exists a unique such hexagon, up to the action of the symmetric group and translation on the indices of vertices. That is the content of the following proposition, which reminds of [B2, Lemma 3.3].

Proposition 2.4.3. *Let $k \geq 2$ and $m \geq 2k + 1$. Up to the action of \mathcal{S}_m , and a translation on the indices of vertices, there exists a unique alternating hexagon in the extended Kneser graph $K^*(m, k)$.*

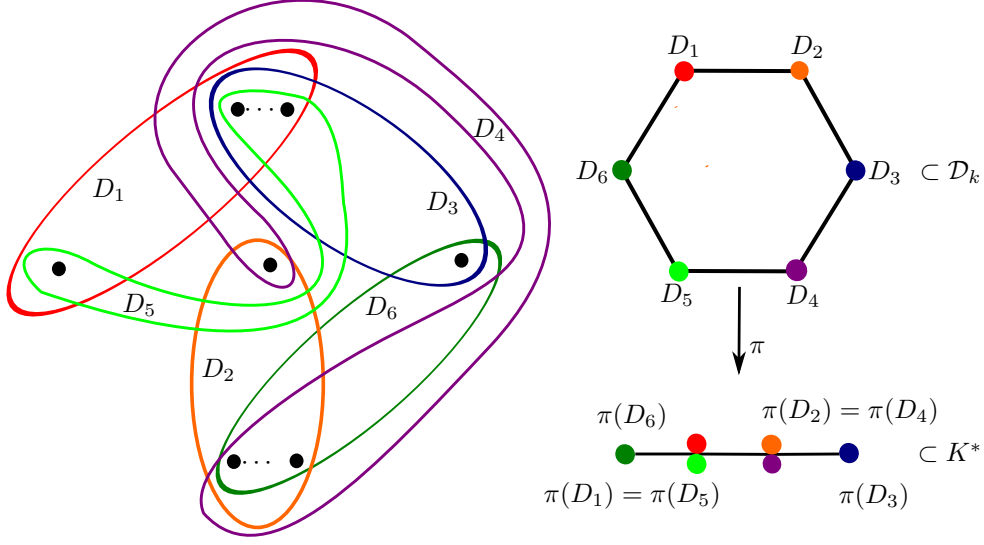


Figure 2.6: A hexagon with non-embedded projection.

The model for that unique hexagon is the following. Let us define $U = \{1, \dots, k-1\}$, $V = \{k, \dots, m-3\}$, and $r = \{m-2\}$, $p = \{m-1\}$, and $q = \{m\}$. The six vertices of the model, in order, are the following:

$$\begin{aligned}
 v_1 &= U \cup r; \\
 v_2 &= V \cup p; \\
 v_3 &= U \cup q; \\
 v_4 &= V \cup r; \\
 v_5 &= U \cup p; \\
 v_6 &= V \cup q.
 \end{aligned}$$

We can now move to the proof of the proposition.

Proof of Proposition 2.4.3. Before starting the proof we recall that, given two sets $Y \subseteq X$, we identify the symmetric group $\mathcal{S}(Y)$ as the subgroup of $\mathcal{S}(X)$ which pointwise fixes $X \setminus Y$. Moreover, $\mathcal{S}_m = \mathcal{S}(\Pi)$ where $\Pi = \{1, \dots, m\}$.

Given a vertex v we will write $\text{Stab}(v)$ to denote the setwise stabiliser of v , that is the subset of \mathcal{S}_m which preserves v setwise, although not necessarily pointwise. This is exactly the stabiliser for the action of \mathcal{S}_m on $K^*(m, k)$. We remark that, given a vertex v and a set $\mathcal{X} \subset \Pi$ such that either $\mathcal{X} \subseteq v$ or $v \cap \mathcal{X} = \emptyset$, we have $\mathcal{S}(\mathcal{X}) \subseteq \text{Stab}(v)$.

For the rest of the proof we will consider $n \geq k$ such that $m = k + n + 1$.

Let (v_1, \dots, v_6) be a hexagon. Up to a translation of the indices of vertices we can assume v_1 to be a k -vertex. Given that the action of \mathcal{S}_m is transitive on k -vertices we can assume that, up to the action of the symmetric group, we have

$$v_1 = U \cup r.$$

we have $v_2 \subseteq v_1^c = V \cup q \cup p$. Thanks to a previous observation we have that

$$\mathcal{S}(V \cup q \cup p) \subset \text{Stab}(v_1).$$

Henceforth, up to a permutation in $\mathcal{S}(V \cup q \cup p)$, the second vertex can be chosen to be

$$v_2 = V \cup p.$$

Similarly, the third vertex is a subset of k elements of $U \cup q \cup r$ which must be different from $v_1 = U \cup r$, so it must contain q . Up to a permutation in

$$\mathcal{S}(U \cup r) \subseteq \text{Stab}(v_1) \cup \text{Stab}(v_2)$$

we can choose it to be

$$v_3 = U \cup q.$$

Similarly the fourth vertex is a subset of n elements of $V \cup p \cup r$ which must be different from $v_2 = V \cup p$, so it must contain r . Up to a permutation in

$$\mathcal{S}(V \cup p) \subseteq \text{Stab}(v_1) \cup \text{Stab}(v_2) \cup \text{Stab}(v_3)$$

we can choose it to be

$$v_4 = V \cup r.$$

We claim that either q or p does not belong to v_5 . If, otherwise, both q and p were contained in v_5 then we would have $U \cup p \cup q \cup r \subseteq v_1 \cup v_5$, hence $v_6 \subseteq (v_1 \cup v_5)^c \subseteq V$, which is impossible since $|V| = n - 1$ but v_6 is a n -vertex. Since $v_5 \subseteq v_4^c = U \cup q \cup p$ it follows that either $v_5 = U \cup p$ or $v_5 = U \cup q$, but in the latter case we would have $v_5 = v_3$, contradicting the fact that the 6-cycle is embedded. It follows that

$$v_5 = U \cup p.$$

we have $v_6 \subseteq v_1^c \cap v_5^c = V \cup q$ and since $|V \cup q| = n$ and v_6 is a n -vertex it follows that

$$v_6 = V \cup q.$$

Uniqueness, up to the action of \mathcal{S}_m and a translation on indices, is proven. \square

The following corollary follows directly from the proof we have just given. It implies that an alternating hexagon in an extended Kneser graph is uniquely determined by four consecutive vertices, and will be useful later.

Corollary 2.4.4. *Let (v_1, v_2, v_3, v_4) be four vertices of $K^*(m, k)$ such that, for $i = 1, 2, 3$, the vertices v_i and v_{i+1} are joined by an edge. Then there exists at most one pair of vertices (v_5, v_6) such that (v_1, \dots, v_6) is an alternating hexagon.*

Proposition 2.4.3 gives us a complete combinatorial structure for the alternating hexagons which project to a hexagon in the extended Kneser graph. However, the previous proposition is unfortunately far from enough to have a complete grasp of the structure of alternating hexagons in the graph of discs.

Indeed, infinitely many discs project to the same vertex in the Kneser graph, so information about the projection is far from univocally determining the disc. This notwithstanding, the projection to the extended Kneser graph will prove to be an invaluable tool for the topological arguments to follow.

Unfortunately, as noted earlier, not every alternating hexagon in the graph of discs projects to an alternating hexagon in the extended Kneser graph. This leads to the following definition, which subdivides the set of hexagons in the graph of discs into two categories: *regular* and *irregular* ones.

Definition 2.4.5. An alternating hexagon H in $\mathcal{D}_k(\Sigma_m)$ is said to be *regular* if its image $\pi(H) = (\pi(D_1), \dots, \pi(D_6))$ under the map π in the extended Kneser graph $K^*(m, k)$ is an alternating hexagon. It is said to be *irregular* otherwise.

We recall that we keep on systematically identifying h -cycles with h -tuples of vertices with the property that immediately subsequent vertices are adjacent in the graph.

In the last part of this section we will start studying the possible configurations arising from the projection of irregular hexagons, in order to complete our understanding of the combinatorics of the marked points for hexagons in $\mathcal{D}_k(\Sigma_m)$. However, in the current section we will only take care of the combinatorial arguments, whereas the identification of the only possible projection for irregular graph requires topological tools to be completed, and will be taken care of in Section 2.6.

The following first result proves that the image of the projection of an alternating hexagon in the curve graph to the extended Kneser graph, when it fails to be a hexagon, is a tree. We will later study which trees can arise in this way.

Proposition 2.4.6. *Let $k \geq 2$, $m \geq 2k + 1$. Let H be an alternating hexagon in $\mathcal{D}_k(\Sigma_m)$ such that $\pi(H) = (v_1, \dots, v_6)$ is not embedded. Then the full subgraph induced by $\pi(H)$ does not contain any embedded h -cycle, for $h \geq 3$. In particular the projection $\pi(H)$ is a tree.*

Proof. The 6-cycle $\pi(H)$ cannot clearly contain any embedded h cycle for $h > 6$.

Moreover, $\pi(H)$ cannot contain any embedded 6-cycle, for such a cycle would be the entire $\pi(H)$, which would then be a hexagon, contradicting the hypothesis.

First, we claim that the subgraph induced by $\pi(H)$ cannot contain an embedded 3-cycle. We argue by contradiction, and, up to a translation on the indices of vertices, we assume that (v_1, \dots, v_3) is an embedded 3-cycle. It follows that those three vertices are pairwise disjoint, hence

$$m \geq |v_1 \cup v_2 \cup v_3| = |v_1| + |v_2| + |v_3| \geq 2k + (m - k - 1) = m + k - 1$$

which is a contradiction as $k \geq 2$. The second inequality follows from the fact that the vertices v_1, v_3 are either k -vertices or $(m - k - 1)$ -vertices, and the vertex v_2 is a $(m - k - 1)$ -vertex or a k -vertex, respectively, as the hexagon H was alternating and p -punctured discs are mapped to p -vertices by π . Since $m \geq 2k + 1$ the first case is the one which minimises the sum of cardinalities, hence the inequality always holds. Our claim is proven.

If $\pi(H)$ contained a 5-cycle then, without loss of generality, we would have that v_1 and v_5 would be adjacent. In particular, those vertices would form a 3-cycle alongside v_6 , which is impossible thanks to the previous claim.

We are left to show that the subgraph induced by $\pi(H)$ cannot contain an embedded 4-cycle. We argue by contradiction, and, up to a translation on the indices of vertices, we assume that (v_1, \dots, v_4) is an embedded 4-cycle. Moreover, this 4-cycle is alternating, as it is the projection of an alternating cycle. It follows that, up to another translation of the indices of vertices, we can assume that v_1 and v_3 are k -vertices. Both v_2 and v_4 , which are adjacent to both v_1 and v_3 must be included in the set $X = (v_1 \cup v_3)^c$. Since $v_1 \neq v_3$, the 4-cycle being embedded, it follows that $|v_1 \cup v_3| \geq k + 1$, hence $|X| \leq m - k - 1$. It follows that there may exist at most one $(m - k - 1)$ -vertex adjacent to both v_1 and v_3 , hence we would have $v_2 = v_4$, contradicting injectivity. The claim is proven and the proof is complete. \square

The only possible trees arising from projections of alternating hexagons to the extended Kneser graphs are exactly the following.

Corollary 2.4.7. *Let $k \geq 2$, $m \geq 2k + 1$. Let H be an alternating hexagon in $\mathcal{D}_k(\Sigma_m)$ such that $\pi(H) = (v_1, \dots, v_6)$ is not a hexagon. Then, up to a translation on the indices of vertices, one of the following happens:*

1. *We have $v_1 = v_3 = v_5 \neq v_2 = v_4 = v_6$;*
2. *We have $v_1 = v_3 = v_5$, $v_2 = v_4 \neq v_6$ and $v_1 \neq v_2, v_6$;*
3. *We have $v_1 = v_3 = v_5$ and $v_i \neq v_j$ otherwise.*
4. *We have $v_1 = v_5$, $v_2 = v_4$, and $v_i \neq v_j$ otherwise.*

Proof. The highlighted configurations are the only possible ones, once it is considered that two neighbouring vertices $v_i, v_{i\pm 1}$ must be different, as they come from projections of disjoint discs. \square

Case 1 in the previous corollary represents the case in which the entire hexagon is collapsed to a single edge, that is all the odd vertices of H project to the same vertex, and all the even vertices project to the same vertex.

Case 2 represents the case where the hexagon H projects to a couple of consecutive edges.

Case 3 describes the projection of the hexagon H to a “tripod”, that is three distinct vertices with a single edge in common. In this case all the three odd or even vertices of H project to the same vertex, while the other three vertices project to three distinct other vertices. For a picture of a hexagon with this behaviour we refer to Figure 2.18.

Case 4 can be thought as a “folding” of the hexagon H over a diagonal (passing through two opposite vertices), that is a projection onto three consecutive edges. This is the behaviour happening in Figure 2.15.

While, *a priori*, all the configurations in Corollary 2.4.7 are allowed by the combinatorics, it turns out (Corollary 2.6.10) that, when $k \geq 3$, the last one is the only possible configuration. When $k = 2$, Case 3 can occur. The obstructions to some types of projections occurring, though, must be looked for in the topology, not just in the combinatorics. We will deal with this in the first half of Section 2.6.

2.5 Nonstandard Regular Hexagons

The goal of this section and the following one is to prove that standard alternating hexagons, hence surrounding pairs and triples, are recognisable in the k -separating disc graph, that is, Theorem 2.2.8. In order to do this we will find a combinatorial property to tell apart nonstandard hexagons from the standard ones. This property is the content of the statement to come. We will prove the case of nonstandard regular hexagons, that is, such that they project to a hexagon in the extended Kneser graph (Definition 2.4.5), in this section, and will take care of irregular hexagons in the next one.

Proposition 2.5.1. *Let $k \geq 3$, $n \geq k$. An alternating hexagon H in $\mathcal{D}_k(\Sigma_{k+n+1})$ is standard if and only if for any choice of four consecutive vertices there exists a different alternating hexagon $H' \neq H$ which shares the four chosen vertices with H .*

Since Proposition 2.5.1 characterises standard hexagons in terms of a purely graph-theoretical property, it follows immediately that the set of standard hexagons is preserved by graph automorphisms, that is, it directly implies Theorem 2.2.8.

In the previous sections, we have mostly showed that some topological properties can be expressed in terms of combinatorial properties of the graph of discs, and we have extracted some combinatorial information thanks to the projection to extended Kneser graphs. In this and the following section we will make use of that information, but the arguments will become essentially topological. We start with a different topological characterisation of standard hexagons, which will be useful later.

Proposition 2.5.2. *Let $k \geq 2$, $n \geq k$. An alternating hexagon H in $\mathcal{D}_k(\Sigma_{k+n+1})$ is standard if and only if for every two 2-distant vertices $A, B \in H$ we have $i(\partial A, \partial B) = 2$.*

Before we can prove this proposition, we will need some technical facts and introduce some notation that will be useful throughout both this and the following section. We will start with a result concerning the uniqueness of common neighbours.

Lemma 2.5.3. *Let $k \geq 2$ and $m \geq 2k + 1$ and let h be such that $k \leq h \leq m - k - 1$. Let A, B be two distinct h -punctured discs in Σ_m . Then there exists at most one $(m - h - 1)$ -punctured disc Δ such that $d(A, \Delta) = d(B, \Delta) = 1$.*

Proof. Let A, B be discs as in the statement. Let $X = A^c \cap B^c$: any disc which is disjoint from both A and B must be contained in X . If $\pi(A) \neq \pi(B)$ it follows that X contains at most $m - h - 1$ marked points, hence it can contain at most one $(m - h - 1)$ -punctured disc.

Let us now suppose that $\pi(A) = \pi(B)$. If we had $A \subseteq B$ from the fact that they are both h -punctured discs it would follow they are isotopic, giving a contradiction. Let $\alpha = \partial A$ and let a be an innermost arc among the closure of components of $\alpha \setminus B$, that is, such that one component of $B^c \setminus a$ does not intersect A . Such an innermost arc always exists. Let Y be the closure of aforementioned component of $B^c \setminus a$. Since α is in minimal position with ∂B it follows that both Y and $B^c \setminus Y$ are not empty bigons, hence they must contain at least one marked point each. Since $\pi(A) = \pi(B)$ then $Y \not\subseteq A$. It follows that $B^c \setminus a$, hence $B^c \cap A^c$, has at least two connected components. As the marked points in B^c are exactly $m - h$ it follows that at every component of $B^c \cap A^c$ can contain at

most $m - h - 1$ marked points. If there existed more than one component of Y containing at least $m - h - 1$ marked points, then, since $k \geq 2$, the surface would contain at least $2(m - h - 1) + h \geq m + k - 1 > m$ marked points, which is a contradiction. It follows that there exists at most one connected component of X containing $m - h - 1$ marked points. Any $(m - h - 1)$ -punctured disc disjoint from both A and B must be contained in such component, if it exists, hence it must be isotopic to its closure. It follows that every two $(m - h - 1)$ -punctured disc disjoint from both A and B are isotopic to each other. The lemma is proven. \square

Definition 2.5.4. Under the hypotheses of the previous lemma we will refer to the common neighbour disc Δ as $\Delta(A, B)$.

An important step towards the proof of Proposition 2.5.2 is given by the following lemma.

Lemma 2.5.5. *Let $k \geq 2$, $n \geq k$. Let $H = (D_1, \dots, D_6)$ be an alternating hexagon of $\mathcal{D}_k(\Sigma_{k+n+1})$ such that for every two 2-distant vertices $A, B \in H$ we have $i(\partial A, \partial B) = 2$. Then the hexagon H is regular.*

Proof. We first claim that every two 2-distant vertices in H project to different vertices in the extended Kneser graph. We argue by contradiction: let A, B be different 2-distant discs such that $i(\partial A, \partial B) = 2$ and $\pi(A) = \pi(B)$. A picture of the following constructions is given in Figure 2.7. Since ∂A and ∂B only intersect twice it follows that the subspace $X = A \cap B^c$ is connected, and its boundary is composed of a single arc of ∂A and a single arc of ∂B . Since every marked point contained in A is contained in B too, it follows that X does not contain any marked point. Hence X is an empty bigon between ∂A and ∂B , contradicting minimal position.

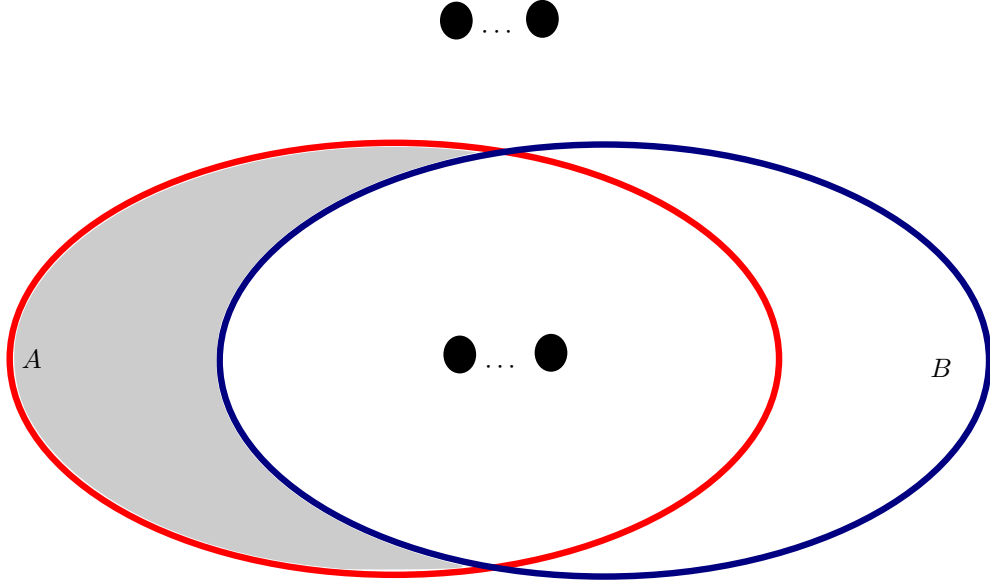


Figure 2.7: The gray region is X .

We are left to show that two opposite vertices in H , say D_1, D_4 , cannot be mapped to the same vertex in the extended Kneser graph $K^*(k + n + 1, k)$. If this happened we

would have a k -punctured disc and a n -punctured disc projecting to the same vertex, hence it should hold $n = k$. In this case, though, since both adjacent and 2-distant vertices must have different projections there would exist an embedded 3-cycle $(\pi(D_1) = \pi(D_4), \pi(D_2), \pi(D_3))$ inside $K^*(2k + 1, k)$. Since the three vertices $\pi(D_1), \pi(D_2)$, and $\pi(D_3)$ would be pairwise disjoint, we would then have

$$2k + 1 \geq |\pi(D_1) \cup \pi(D_2) \cup \pi(D_3)| = |\pi(D_1)| + |\pi(D_2)| + |\pi(D_3)| = 3k,$$

which is impossible since $k \geq 2$. \square

We are now ready to prove a crucial step for our characterisation of standard hexagons based on the intersection numbers, which will also be useful later.

Lemma 2.5.6. *Let A, B, C be three k -separating discs such that both the triple intersection and any pairwise intersection of their projections to the extended Kneser graphs has cardinality exactly $k - 1$. Then if $i(\partial A, \partial B) = 2$ the pair $\{A, B\}$ is a surrounding pair.*

Moreover, if the boundaries of any pair of discs intersect exactly twice the triple $\{A, B, C\}$ is a surrounding triple.

We remark that the condition on the intersection in the extended Kneser graph is automatically satisfied when the projections of the discs are the three k -vertices of a regular alternating hexagon, thanks to Proposition 2.4.3.

Proof of Lemma 2.5.6. First, we prove the statement for pairs. Let A, B be as in the statement, let $O = A \cap B$. Since $i(\partial A, \partial B) = 2$ it follows that ∂O is composed of a single arc of ∂B inside A and a single arc of ∂A inside B . In particular O is connected. It follows that O contains exactly the marked points in $\pi(A) \cap \pi(B)$, which are $k - 1$ by hypothesis. It follows that O is a $(k - 1)$ -punctured disc, hence $\{A, B\}$ is a surrounding pair.

We now claim that if C is a k -punctured disc whose boundary intersects both ∂A and ∂B exactly twice, then $\{A, B, C\}$ is a surrounding triple. A picture of the following constructions is provided in Figure 2.8.

From the previous argument we know that $\{A, B\}$, $\{A, C\}$, and $\{B, C\}$ are all surrounding pairs. By hypothesis We have

$$|\pi(A) \cap \pi(B) \cap \pi(C)| = k - 1 \geq 1,$$

as $k \geq 2$. It follows that $X = A \cap B \cap C$ contains $k - 1$ marked points, so in particular it is nonempty. Since X is the intersection of discs whose boundaries pairwise intersect twice each of its connected components must have only one boundary component, that is, X is a punctured disc. We claim this region X is connected, hence it is a $(k - 1)$ -punctured disc. If $A \cap B \subseteq C$ then $X = O$ and our claim follows immediately.

Let us now suppose $A \cap B \not\subseteq C$. Let Y be a connected component of X . At least one arc of ∂Y is contained in ∂C , hence there exists at least two intersection points between ∂C and $\partial(A \cap B)$. It follows that if X were not connected there would be at least four intersection between ∂C and $\partial(A \cap B)$. In particular, since we are assuming all the curves to be in minimal position without triple points, all the aforementioned intersections are

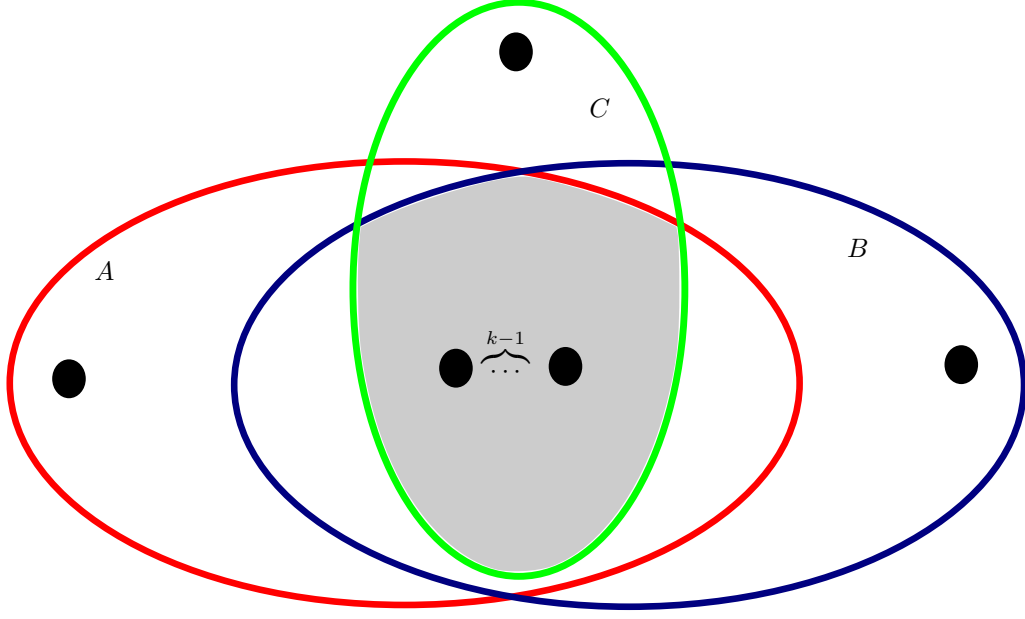


Figure 2.8: The gray region is X .

realised in $\partial(A \cap B) \subset (A \cup B)^\circ$, where $(A \cup B)^\circ$ denotes the interior of the subspace $A \cup B$. Since $C \not\subset A \cup B$, it follows that ∂C must intersect $\partial(A \cup B) \not\subset (A \cup B)^\circ$, hence this intersection has not been counted yet. It follows that ∂C intersects $\partial A \cup \partial B$ at least five times, which is impossible, since it intersects each of the two curves twice. This is contradiction, which proves our claim that $X = A \cap B \cap C$ is connected, hence it is a $(k-1)$ -punctured disc. It follows that $\{A, B, C\}$ is a surrounding triple. \square

We can now prove the characterisation of standard alternating hexagons via intersection numbers.

Proof of Proposition 2.5.2. The fact that the discs in a standard alternating hexagon have the required property follows directly from the definition and can easily be visualised in Figure 2.1.

For the converse let (A_1, \dots, A_6) be an alternating hexagon as in the statement. Thanks to Lemma 2.5.5 the projection of such hexagon to the extended Kneser graph is an alternating hexagon. Hence, Lemma 2.5.6 implies that $\{A_1, A_3, A_5\}$ is a surrounding triple. Thanks to Proposition 2.2.9 it follows that, up to the action of the mapping class group, the triple (A_1, A_3, A_5) can be identified with (D_1, D_3, D_5) in the standard alternating hexagon as defined in Definition 2.2.7. From Lemma 2.5.3 it follows that, under the previous identification, the common neighbour $A_{i+1} = \Delta(A_i, A_{i+2})$ must be equal to $\Delta(D_i, D_{i+1}) = D_{i+1}$, for $i = 1, 3, 5$. It follows that the hexagon is standard. \square

The alternating hexagon in Figure 2.9 provides an example of a hexagon which is regular, that is, it projects to a hexagon in the extended Kneser graph, but is not standard.

Let us note that the hexagon in Figure 2.9 shares the four consecutive vertices D_6, D_1, D_2, D_3 with a standard hexagon. This fact, alongside the transitivity property of the action of the mapping class group on a standard hexagon provided by the

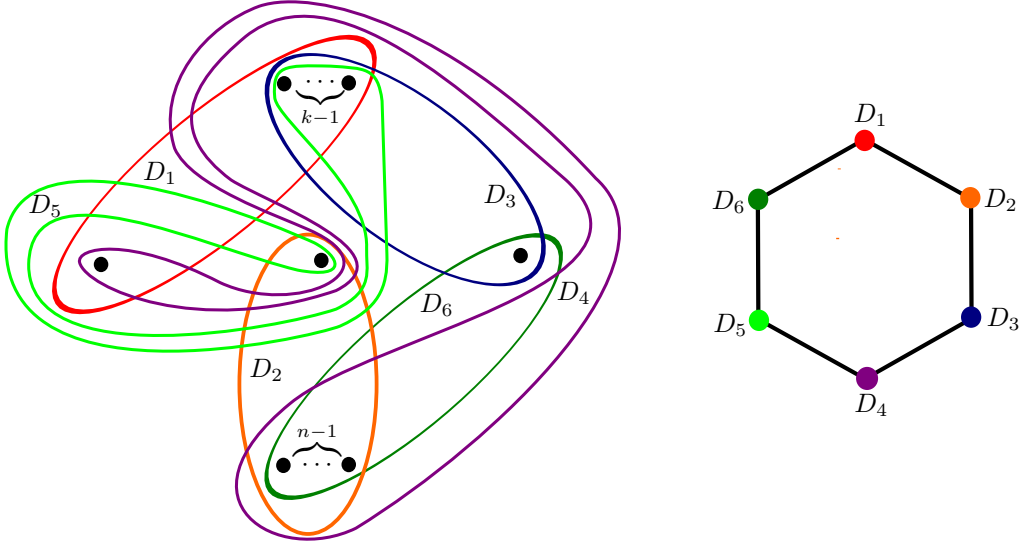


Figure 2.9: A nonstandard regular alternating hexagon.

next proposition, proves the “only if” implication of Proposition 2.5.1. Indeed, for every choice of four consecutive vertices in a standard alternating hexagon there exists a different alternating hexagon which shares them.

Proposition 2.5.7. *Let $k \geq 2$, $n \geq k$. Let H be a standard alternating hexagon in $\mathcal{D}_k(\Sigma_{k+n+1})$, and let $\text{Isom}^*(H)$ be the subgroup of simplicial automorphisms of H preserving the topological types of discs. Then there exists a subgroup $G < \text{Mod}^\pm(\Sigma_{k+n+1})$ acting on the hexagon H by graph automorphisms such that the induced group homomorphism $\Phi : G \longrightarrow \text{Isom}^*(H)$ is surjective.*

If $n > k$ there are two different topological types of discs in the hexagon, hence the group $\text{Isom}^*(H)$ preserves odd and even vertices and, thanks to Lemma 2.5.3, such action is completely determined by the action on one of these subset, on which it is transitive. It follows that is, this case We have $\text{Isom}^*(H) \cong \mathcal{S}_3$. When $n = k$, instead, the action is transitive on all the vertices of the hexagon, hence the group is the full dihedral group of the hexagon (which has order 12), that is, $\text{Isom}^*(H) \cong \mathcal{D}_6$.

Proof of Proposition 2.5.7. Throughout the entire proof will use the description of a standard hexagon given in Definition 2.2.7, considering the punctured sphere to embedded in \mathbb{R}^3 , as described there. A cartoon of the maps described in this proof is provided in Figure 2.10.

First, we introduce the “rotation of an angle $\frac{2\pi}{3}$ along the equator”, and study its action on the hexagon. More formally, let τ be a homeomorphism supported in a $\frac{1}{3}$ -neighbourhood of the equator, that is, a 3-punctured annulus which we will identify with

$$A = (T \times [-1, 1]) \setminus (\{v_1, v_2, v_3\} \times \{0\}).$$

In the coordinates given by such parametrisation we define the homeomorphism τ to rotate the latitude $T \times \{h\}$ of an angle $(1 - |h|)\frac{2\pi}{3}$. The homeomorphism τ acts on the standard hexagon by $D_i \mapsto D_{i-2}$.

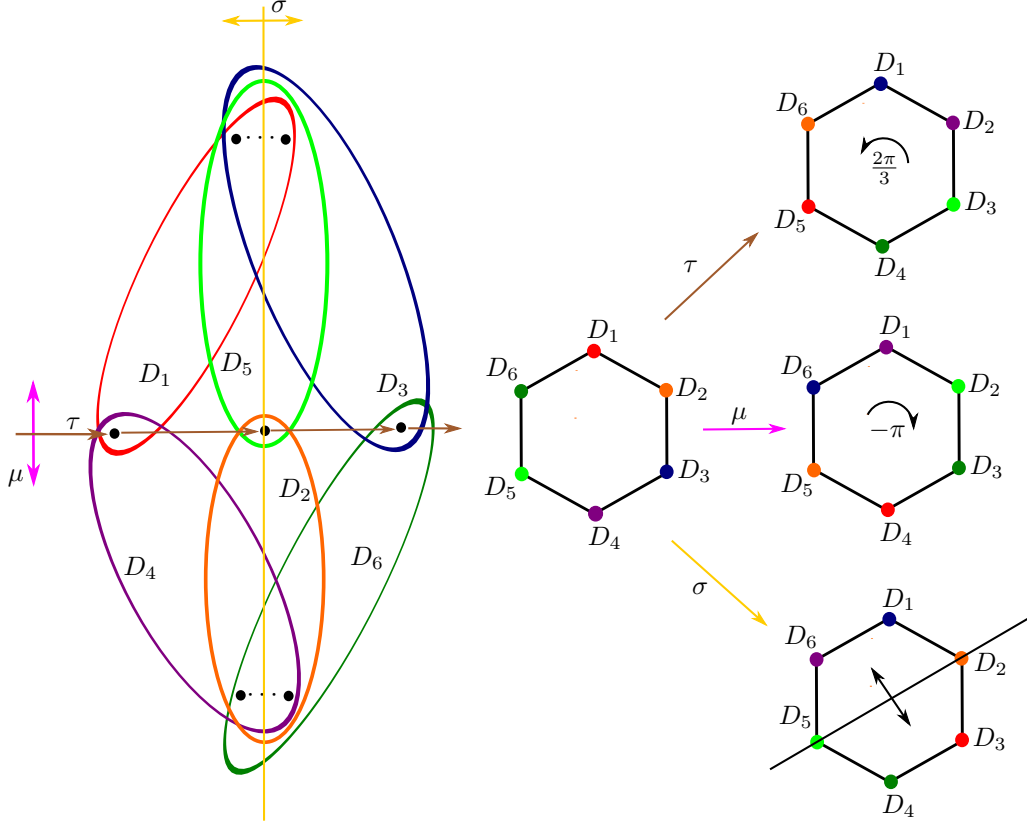


Figure 2.10

Second, we consider a “reflection along a vertical plane”, and study its action. More formally, let σ be the reflection along the plane (we assume that we are assuming Σ_{2k+1} to be naturally embedded in \mathbb{R}^3) containing l_2 and l_5 (we recall that we can assume the marked points at the poles to be arranged in a way such that this map is well-defined). This homeomorphism permutes two of the three marked points on the equator. On the standard hexagon σ acts by $D_i \mapsto D_{4-i}$.

Let $G = \langle \sigma, \tau \rangle < \text{Mod}^\pm(\Sigma_{k+n+1})$. Then G acts on H and let $\Phi : G \longrightarrow \text{Aut}^*(H)$ be the induced group homomorphism. We can write as a composition of cycles the two permutations $\Phi(\tau) = (1\ 5\ 3)(2\ 6\ 4)$ and $\Phi(\sigma) = (1\ 3)(4\ 6)$. The action of these permutations is transitive on the set of both odd and even vertices. From the observation made before the statement it follows that, if $n > k$, $\Phi(\tau)$ and $\Phi(\sigma)$ generate the entire group $\text{Isom}^*(H)$, that is, We have $\Phi(G) = \text{Isom}^*(H)$.

Lastly, when $n = k$ we can define a new homeomorphism of the sphere preserving the hexagon. This homeomorphism is the “reflection along the equator”, which was previously not well-defined because the number of marked points around the north pole and around the south pole differed. Formally, let μ be the reflection along the equator: this map is well-defined in this case since the number of marked points around the north pole is equal to the number of punctures around the south pole, and we can assume they are placed symmetrically with respect to this reflection. The homeomorphism μ acts on the standard hexagon by $D_i \mapsto D_{i+3}$.

In particular, the composition $\rho = \mu \circ \tau$ acts on the hexagon by $D_i \mapsto D_{i+1}$. Let $G = \langle \sigma, \rho \rangle < \text{Mod}^\pm(\Sigma_{2k+1})$. Then $\Phi(\rho)$ is a rotation of the hexagon by an angle of $\frac{\pi}{3}$ (if we think to the hexagon as embedded in the plane), where $\Phi(\sigma)$ is a reflection along a diagonal. It follows that $\Phi(\rho)$ and $\Phi(\sigma)$ generate the entire dihedral group \mathcal{D}_6 , that is, $\Phi(G) = \text{Isom}(H)$. \square

Before we enter the core of the argument for the “if” implication in Proposition 2.5.1, we will now introduce some notation and state some technical results that will be useful throughout both this section and the following one.

Lemma 2.5.8. *Let $k \geq 2$, $m \geq 2k + 1$, and let h be such that $k \leq h \leq m - k - 1$. Let A, B be two 2-distant h -punctured discs in $\mathcal{D}_k(\Sigma_m)$ such that $\pi(A) \neq \pi(B)$ and there exists a $(m - h - 1)$ -punctured disc D such that $D = \Delta(A, B)$, the common neighbour as in Definition 2.5.4. Then there exists a unique marked point in $\pi(A) \setminus \pi(B)$.*

We remark that the technical hypothesis on the two discs A, B in the previous lemma, which will be a recurrent condition in many other statements to follow, is always satisfied for a couple of discs which are 2-distant vertices of an alternating hexagon.

We will now give a couple of definitions, descending from the previous result, which will be useful later.

Definition 2.5.9. Under the hypotheses of the previous lemma, we will denote with $p(A, B)$ the unique marked point in $\pi(A) \setminus \pi(B)$.

Moreover, we define $M(A, B)$ to be the (closure of) the connected component of $A \cap B^c$ containing $p(A, B)$.

We observe that the objects we have just defined are not symmetrical: indeed, it always holds $p(A, B) \neq p(B, A)$ and $M(A, B) \neq M(B, A)$.

We will now prove the lemma.

Proof of Lemma 2.5.8. Since $\pi(A) \neq \pi(B)$ it follows that the set $\pi(A) \cup \pi(B)$ has cardinality at least $h + 1$, and it has cardinality at most $h + 1$ since it is a subset of the $(h + 1)$ -punctured disc D^c . It follows that the set $\pi(A) \cup \pi(B)$ has cardinality exactly $h + 1$ and an application of the Inclusion-Exclusion Principle yields that $\pi(A) \cap \pi(B)$ has cardinality exactly $h - 1$, so $|\pi(A) \setminus \pi(B)| = 1$. \square

We introduce the following.

Definition 2.5.10. Let $k \geq 2$, $m \geq 2k + 1$, and let h be such that $k \leq h \leq m - k - 1$. Let A be a h -punctured disc in Σ_m and B a $(m - h - 1)$ -punctured disc disjoint from A . The region $A^c \cap B^c$ is a pair of pants (to be precise surface of type $S_{0,1}^2$) which will be denoted with $P(A, B)$.

We remark that the previous condition is always satisfied for two adjacent discs in an alternating hexagon. Moreover, we notice that the previous definition is symmetrical, that is, $P(A, B) = P(B, A)$.

There are three different types of arcs in the previously introduced pair of pants, which we will classify as follows.

Definition 2.5.11. Let $P(A, B)$ be the pair of pants as in Definition 2.5.10. Then we define the following types of (parametrised) arcs in $P(A, B)$:

Mixed arcs are arcs with one endpoint on ∂A and one on ∂B .

TA -arcs are arcs with both endpoints of the same curve ∂A .

TB -arcs are arcs with both endpoints of the same curve ∂B .

All of these arcs will be considered up to isotopy which keeps each endpoint on the same boundary component (although does not necessarily fix it) at all times.

Since the subsurface $P(A, B)$ is a pair of pants the following lemma is immediate.

Lemma 2.5.12. *Under the aforementioned hypotheses, every two nontrivial parametrised TA -arcs (resp. TB -arcs, mixed arcs) are isotopic up to orientation.*

We can now observe the following lemma.

Lemma 2.5.13. *Let $k \geq 2$, $m \geq 2k + 1$, let h be such that $k \leq h \leq m - k - 1$. Let A, B be two 2-distant h -punctured discs in Σ_m such that $\pi(A) \neq \pi(B)$ and there exists a $(m - h - 1)$ -punctured disc D such that $D = \Delta(A, B)$. Then the region $M(A, B)$, as in Definition 2.5.9, is a once-punctured disc whose boundary is the union of a single arc of ∂A and a single arc of ∂B .*

Proof. First we observe that, given two parametrised disjoint TB -arcs, say a, a' , the set $B^c \setminus (a \cup a')$ is composed of three connected components, and only one is such that its boundary contains both arcs. It follows that the two arcs are isotopic if and only if said component contains no punctures.

Every arc of $(\partial A) \setminus B$ is a TB -arc in $P(B, D)$, so they are all isotopic to each other due to Lemma 2.5.12. From the previous observation it follows that every component of $A \cap B^c$ whose boundary contains multiple TB -arcs cannot contain any puncture in it. In particular it cannot contain $p(A, B)$, and it follows that such a component cannot be $M(A, B)$. \square

Given the previous result, we can now give the following definition.

Definition 2.5.14. Under the same hypotheses as in the previous lemma we will denote the closure of the component $\partial M(A, B) \cap \partial A$, which is an arc thanks to Lemma 2.5.13, with $s(A, B) \subseteq \partial A$.

Moreover, we will denote with $t(B, A) \subseteq \partial B$ the closure of the arc $\partial B \setminus M(A, B)$.

Once again, we stress the fact that the notation we have just introduced is not symmetrical.

Thanks to Lemma 2.5.13 we can now introduce a convenient simplification, and describe the component $M(A, B)$ as an arc, rather than as a subsurface.

Definition 2.5.15. Under the same hypotheses as in Lemma 2.5.13, let $l(A, B)$ be the unique arc in $M(A, B)$ with one endpoint on $\partial B \cap M(A, B)$ and $p(A, B)$ as the other one, up to isotopy which fixes $p(A, B)$ and keeps the other endpoint on $\partial B \cap M(A, B)$ at any time.

Once again, we remark that the previously introduced notation is not symmetrical.

Uniqueness of the arc $l(A, B)$, up to isotopy, follows immediately from the fact that $M(A, B)$ is a once-punctured disc, and that $\partial B \cap M(A, B)$ is connected, which is Lemma 2.5.13.

We observe that $M(A, B)$ is isotopic to a closed small regular neighbourhood of $l(A, B)$. Conversely, given two arcs in B^c with one endpoints on ∂B and $p(A, B)$ as the other one, two sufficiently small regular neighbourhoods are isotopic if and only if the two arcs are. Thanks to these observations from now on, instead of studying the isotopy class of the region $M(A, B)$, we will study that of the arc $l(A, B)$.

We can observe the following simple but crucial facts.

Lemma 2.5.16. *Let $k \geq 2$, $m \geq 2k + 1$, let h be such that $k \leq h \leq m - k - 1$. Let A, B be two 2-distant k -punctured discs in Σ_m such that $\pi(A) \neq \pi(B)$ and there exists a $(m - h - 1)$ -punctured disc D such that $D = \Delta(A, B)$. Then the region $M(A, B) \cup B$ is a punctured disc isotopic to $\Delta(A, B)^c$.*

Moreover, the curve $s(A, B) \cup t(B, A)$ is isotopic to $\partial\Delta(A, B)$. In particular, if ζ is a curve in Σ_m disjoint from ∂B it follows that

$$i(\partial\Delta(A, B), \zeta) = i(s(A, B), \zeta).$$

Proof. The region $X = M(A, B) \cup B$ contains $h + 1$ marked points and is contained in the $(h + 1)$ -punctured disc $\Delta(A, B)^c$. It follows that $X = \Delta(A, B)^c$. Moreover, ∂X is the simple closed curve given by the union of the two arcs $s(A, B) \cup t(B, A)$.

For the second part We have

$$i(\zeta, \partial\Delta(A, B)) = i(\zeta, s(A, B)) + i(\zeta, t(B, A)) = i(\zeta, s(A, B))$$

since $t(B, A) \subseteq \partial B$ and ∂B and ζ are disjoint. \square

We are now ready to start approaching the proof of the “if” direction in Proposition 2.5.1 for the case regular hexagons. In doing so we will consider non-standard regular alternating hexagons. Since these hexagons are non-standard Proposition 2.5.2 implies that there exist two 2-distant vertices whose boundaries intersect more that twice, but the topological type of those discs is not specified. In order to simplify our arguments we show that it is not restrictive to assume those two vertices to be n -punctured discs.

Lemma 2.5.17. *Let $k \geq 2$, $n \geq k$. Let H be a non-standard regular alternating hexagon in $\mathcal{D}_k(\Sigma_{k+n+1})$. Then there exists two 2-distant vertices representing n -punctured discs A, B such that $i(\partial A, \partial B) > 2$.*

Proof. The statement follows from the same argument as in the proof of Proposition 2.5.2. Indeed, we can argue by contradiction and the argument we used there can be replicated to prove that the n -punctured vertices can be simultaneously identified with the three n -punctured vertices of the standard alternating hexagon in Definition 2.2.7, hence applying Lemma 2.5.3 to conclude that the hexagon is standard. In fact, the only hypothesis needed for the aforementioned argument would be that for each pair of n -punctured discs in the hexagon their boundaries intersect exactly twice.

An alternative direct argument is the following. Let $H = (D_1, \dots, D_6)$ be a nonstandard hexagon and let $\gamma_i = \partial D_i$. Let us argue by contradiction and suppose that for every two n -punctured discs in H their boundaries intersect exactly twice. Since the hexagon was not standard, due to Proposition 2.5.2 we can assume, without loss of generality, that odd vertices are k -punctured, and that $i(\gamma_1, \gamma_3) > 2$. An application of Lemma 2.5.16 with $A = D_1$, $B = D_3$ and $\zeta = \gamma_4$ yields:

$$2 = i(\gamma_2, \gamma_4) = i(s(D_1, D_3), \gamma_4) = \frac{1}{2} i(\gamma_1, \gamma_3) i(\gamma_1, \gamma_4) \geq i(\gamma_1, \gamma_3) > 2$$

and we have found a contradiction. In the third equality we have used the fact that $\gamma_1 \cap D_4 \subseteq \gamma_1 \setminus D_3$, and the fact that all the arcs of $\gamma_1 \setminus D_3$ are isotopic to $s(D_1, D_3)$, since they are TD_3 -arcs in $P(D_3, D_2)$, thanks to Lemma 2.5.12. That, combined with the fact that the number of those arcs is half the intersection number between γ_1 and γ_3 , proves the equality we were looking for. Lastly, for the second to last inequality we use the fact that $i(\gamma_1, \gamma_4) \geq 2$, as the discs D_1 and D_4 are not disjoint as, for instance, they have intersecting projections to the extended Kneser graph, thanks to Proposition 2.4.3. \square

Thanks to the previous lemma, the following proposition proves Proposition 2.5.1 for regular hexagons, and we will devote the rest of this section to its proof.

Proposition 2.5.18. *Let $k \geq 2$ and $n \geq k$. Let (D_1, D_2, D_3, D_4) be four vertices of $\mathcal{D}_k(\Sigma_{k+n+1})$ such that, for $i = 1, 2, 3$, D_i is adjacent to D_{i+1} . Moreover, let D_1 be a k -punctured disc, and assume that $i(\partial D_2, \partial D_4) > 2$. Then there exists at most one pair of discs (D_5, D_6) such that (D_1, \dots, D_6) is a regular nonstandard alternating hexagon.*

From now on let (D_1, D_2, D_3, D_4) be as in the hypotheses of the previous proposition, and let (D_5, D_6) be such that (D_1, \dots, D_6) is a non-standard alternating regular hexagon. Moreover let $\gamma_i = \partial D_i$. We will now study how rigid this structure has to be: our final goal is to prove that, given the first four curves, the other two are uniquely determined.

For the remain of this section we will simplify the notation we employ by writing $d = l(D_5, D_3) \subseteq D_5$ and $z = s(D_1, D_3) \subseteq \gamma_1$.

Moreover, we will write $P = P(D_2, D_3)$. Since the hexagon (D_1, \dots, D_6) is regular Proposition 2.4.3 implies that $p(D_5, D_3) = p(D_5, D_1)$: we will denote this puncture with p .

Moreover, we recall that, thanks to Proposition 2.4.3, there exists a partition of the set of marked points as $\Pi = U \sqcup V \sqcup \{p, q, r\}$ where $|U| = k - 1$ and $|V| = n - 1$ such that the following holds:

$$\begin{aligned} \pi(D_1) &= U \cup \{r\}; \\ \pi(D_2) &= V \cup \{p\}; \\ \pi(D_3) &= U \cup \{q\}; \\ \pi(D_4) &= V \cup \{r\}; \\ \pi(D_5) &= U \cup \{p\}; \\ \pi(D_6) &= V \cup \{q\}. \end{aligned}$$

Throughout the rest of the section we will often use this information about the projection

to the extended Kneser graph.

We will now define a subdivision for the arc d .

Definition 2.5.19. We now fix an orientation on d by having the first endpoint being on γ_3 and p as the second one. We define a subdivision of the arc d as

$$d = d_0 \cup \dots \cup d_j$$

where the d_i 's are closures of connected components of $d \setminus \gamma_2$ ordered using the orientation of d , and such that the intersection of two consecutive arcs is a single point.

The number j is function of the discs D_2 and D_4 only, as the following lemma shows.

Lemma 2.5.20. *We have that $j = i(d, \gamma_2) = \frac{i(s(D_5, D_3), \gamma_2)}{2} = \frac{i(\gamma_4, \gamma_2)}{2} \geq 2$.*

Proof. The second equality follows from the fact that

$$\partial M(D_5, D_3) = s(D_5, D_3) \cup t(D_3, D_5),$$

and for every intersection of d with γ_2 the boundary of its regular neighbourhood $M(D_5, D_3)$ will have two.

Moreover, since $t(D_3, D_5) \subseteq \gamma_3$ and γ_2 are disjoint, the third equality follows from Lemma 2.5.16 with $A = D_5$, $B = D_3$ and $\zeta = \gamma_2$. \square

Since $j = i(d, \gamma_2) \geq 2$ it follows that the arc d_1 has both endpoints on γ_2 . The set $D_2 \setminus d_1$ is hence composed of two components whose closures are punctured discs. We will denote those discs with L_1, L_2 .

We will now start studying how marked points are distributed between the aforementioned regions, and will later study which are the regions whose interior can intersect the arc d . The final goal is to prove that every arc d_i with $i \geq 2$ lies in a region which is topologically sufficiently simple to uniquely determine the isotopy type of the arc. Our final goal will be to prove that the entire arc d is uniquely determined by the first four vertices of the hexagon, that is, by (D_1, \dots, D_4) .

First, we will state a trivial technical observation that will be at the core of many arguments to follow.

Lemma 2.5.21. *Let $X, Y \subseteq \Sigma_{k+n+1}$ be two connected closed subsets such that $X \cap Y \neq \emptyset$ and $\partial X \cap \partial Y = \emptyset$. Then either $X \subseteq Y$ or $Y \subseteq X$. In particular, if X contains strictly fewer marked points than Y , We have that $X \subseteq Y$.*

A more directly applicable statement we will use in many upcoming arguments is the following corollary. We recall that in all of our applications the technical condition about the existence of the common neighbour $\Delta(A, B)$, as in Definition 2.5.4, is automatically satisfied for any two 2-distant vertices in an alternating hexagon.

Corollary 2.5.22. *Let $k \geq 2$, $m \geq 2k + 1$, and let h be such that $k \leq h \leq m - k - 1$. Let A, B be two 2-distant h -punctured discs such that $\pi(A) \neq \pi(B)$ and there exists a $(m - h - 1)$ -punctured disc D such that $D = \Delta(A, B)$. Let $X \subseteq \Sigma_m$ be a connected closed subsurface such that $\partial X \subseteq A \cup B$. If $X \cap \Delta(A, B) \neq \emptyset$ then $\Delta(A, B) \subseteq X$. In particular X contains every marked point in $\pi(\Delta(A, B))$.*

Proof of Corollary 2.5.22. Since $\partial\Delta(A, B)$ is disjoint from $\partial X \subseteq A \cup B$ it follows from Lemma 2.5.21 that either $X \subseteq \Delta(A, B)$ or $\Delta(A, B) \subseteq X$. The former case is impossible since $\Delta(A, B)$ is disjoint from $A \cup B$. \square

We are now ready to study how the marked points in $\pi(D_2)$ are split among the two components L_1, L_2 . Let us recall that $\pi(D_2) = V \cup \{p\}$.

Lemma 2.5.23. *One of the regions L_1, L_2 contains every marked points in V , while the other only contains $p = p(\gamma_5, \gamma_3)$.*

Proof. Let us suppose by contradiction that both L_1 and L_2 contain at least one marked point of V . It follows that both $L_1 \cap D_6$ and $L_2 \cap D_6$ are not empty. In particular it follows that there exists an arc $c \subseteq \gamma_6$ with one endpoint on $\partial L_1 \cap \gamma_2$ and the other on $\partial L_2 \cap \gamma_2$, and whose interior does not intersect D_2 . The subspace $(D_2^c \cap D_3^c) \setminus c$ has at least two connected components, one of which contains d_0 , another one of which contains d_2 . We will denote the component containing d_0 by Z , and the one containing d_2 by W .

We claim that the component Z must intersect D_1 . Indeed, if c is disjoint from D_3 it follows directly that $D_3 \subseteq Z$, as an immediate consequence of Lemma 2.5.21 applied with $X = D_3$ and $Y = Z$. Since $D_1 \cap D_3 \neq \emptyset$ it follows that $D_1 \cap Z \neq \emptyset$.

Let us now suppose that c intersects D_3 . Every arc of $\gamma_5 \cap D_3^c$ is isotopic to $s(D_5, D_3)$, which is a consequence of Lemma 2.5.12 as they are all TD_3 -arcs in $P(D_3, D_4)$. Thanks to this observation it follows from Lemma 2.5.21, applied with $X = W \cap D_3$ and $Y = D_3 \cap D_5$, that $D_3 \cap D_5 \subseteq Z$. In particular $\pi(D_3) \cap \pi(D_5) = \pi(D_1) \cap \pi(D_5) \subseteq Z$, hence $Z \cap D_1 \neq \emptyset$.

It follows from Corollary 2.5.22, applied with $A = D_2$, $B = D_6$ and $X = Z$, that Z contains all the marked points in $\pi(D_1)$. Given that the only marked point in $D_2^c \cap D_3^c$ is $p(D_1, D_3) \in \pi(D_1)$, it follows that $W \cap D_3^c$ contains no marked points. Since $d_2 \subset W \cap D_3^c$ it follows that $s(D_5, D_2)$ forms an empty bigon with γ_2 , hence γ_5 does. We have found a contradiction to minimal position, and hence proven the lemma. \square

With the previous lemma in mind, from now on we will call L_1 is the region containing the marked points in V , while L_2 will be the one only containing p .

We will now study the intersection of the arc d with L_1 .

Lemma 2.5.24. *The interior of the region L_1 (containing the punctures in V) does not intersect d .*

Proof. Let us argue by contradiction. Let us consider the set of arcs d_i 's whose interior is contained in the interior of L_1 (so it must hold $i \geq 3$), and suppose this set is non empty. Let d_h be an outermost one among those arcs, that is, such that one of the regions of $L_1 \setminus d_h$ does not intersect the interior of any of the d_j arcs. We will denote this outermost region by Y . Let x, x' be the endpoints of d_h , and let a be the arc of γ_2 between them whose interior does not intersect d . These and the following constructions are pictured in Figure 2.11.

At least one of the marked points in V must be contained in Y , otherwise the arcs d_h and a would form an empty bigon, contradicting minimal position. We will denote the set of marked points in this region by $Q = Y \cap \Pi$. Now, let y, y' be the intersection of d_{h-1}, d_{h+1} with z , such that the interior of the subarc of d_{h-1} (resp. d_{h+1}) between x

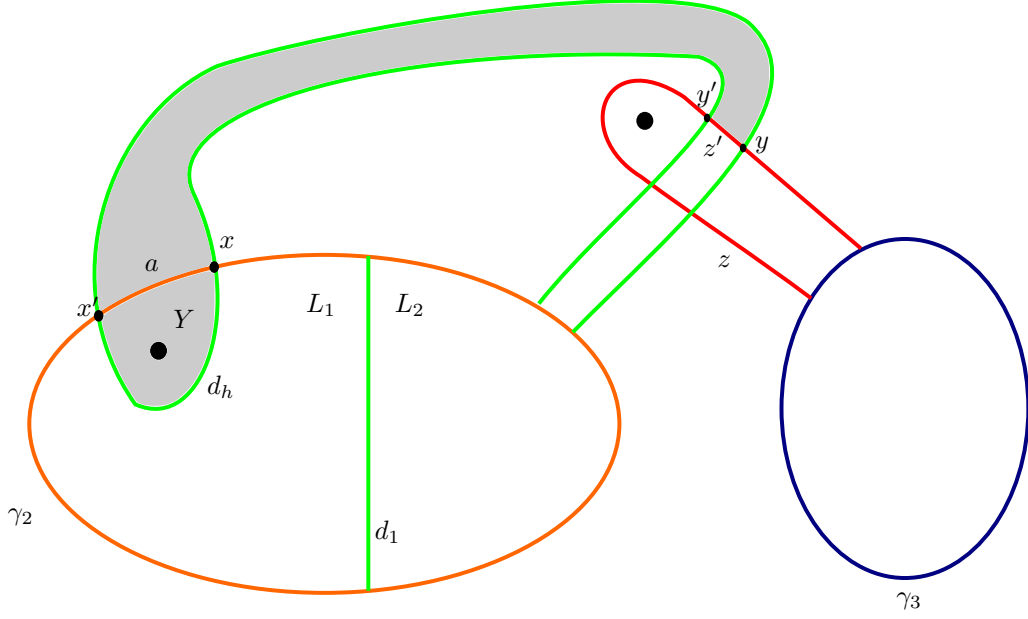


Figure 2.11: The gray region is Z .

and y (resp. x' and y') does not intersect z . Let z' be the subarc of z between y and y' . Let

$$\beta = d_{i-1} \cup d_i \cup d_{i+1} \cup z'.$$

Let Z be the connected component of $\Sigma_m \setminus \beta$ containing Y . Let us observe that We have $\partial Z \subseteq \beta \subseteq D_5 \cup D_1$. From Corollary 2.5.22, applied with $A = D_1$, $B = D_5$ and $X = Z$, it follows that $\pi(D_6) \subseteq Q$. But this is absurd since We have $Q \subseteq \pi(D_2)$, and hence we would have $\pi(D_6) \subseteq \pi(D_2)$, which is impossible since the hexagon (D_1, \dots, D_6) is regular, hence $\pi(D_2) = V \cup \{p\}$ and $\pi(D_6) = V \cup \{q\}$ thanks to Proposition 2.4.3. We have reached a contradiction and the lemma is proven. \square

We have just proven that, for $i \geq 2$, the arc d_i lies in P if i is even, and it lies in L_2 when i is odd. From now on let $\tilde{d} = d_2 \cup \dots \cup d_m$.

Let N be a small regular neighbourhood of d_0 not containing any marked point and not intersecting any arc d_i with $i \geq 2$. Let X be the closure of the region

$$(P \cup L_2) \setminus N.$$

The region X , represented in Figure 2.12, only contains two marked points, which are $p = p(D_5, D_3)$ and $r = p(D_1, D_3)$ and its boundary is connected and contained in $D_3 \cup N \cup L_1$, hence X is a pair of pants (in particular a surface of type S_2^1). It follows that \tilde{d} is an arc in the pair of pants X with one endpoint $(d_1 \cap d_2)$ on the boundary and p as the other endpoint.

We will now use this information to study the isotopy type of d , and hence the isotopy type of $M(D_5, D_3)$. That is, the content of the following.

Lemma 2.5.25. *Let $k \geq 2$ and $n \geq k$. Let (D_1, D_2, D_3, D_4) be four vertices of $\mathcal{D}_k(\Sigma_{k+n+1})$ such that, for $i = 1, 2, 3$, D_i is adjacent to D_{i+1} . Suppose that D_1 a k -*

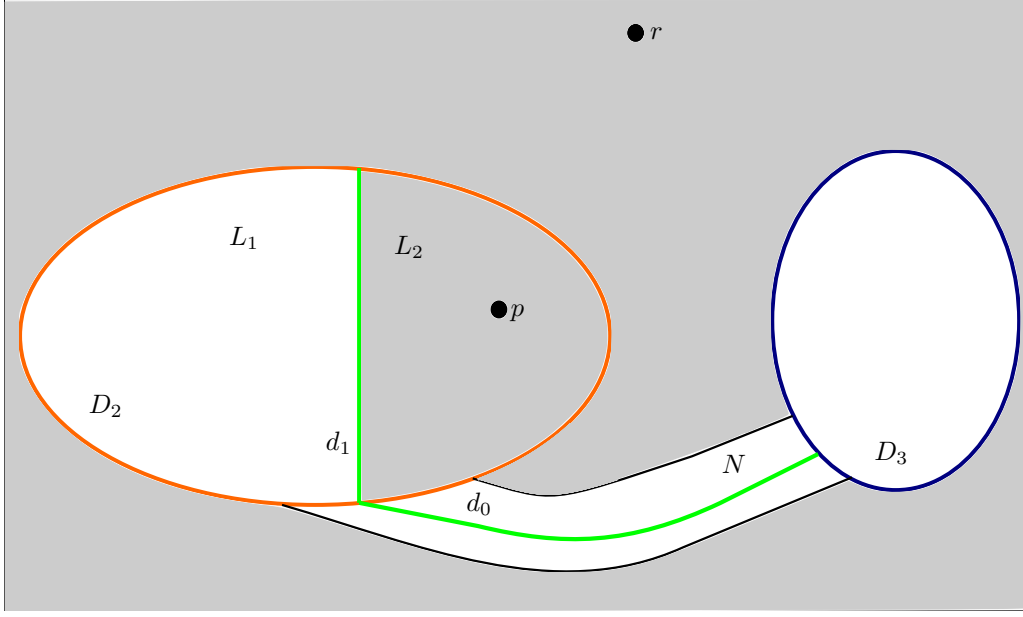


Figure 2.12: The gray region is X .

punctured disc, and let $\gamma_i = \partial D_i$. Let us suppose that $i(\gamma_2, \gamma_4) > 2$. Let (D_5, D_6) and (D'_5, D'_6) be such that both (D_1, \dots, D_5, D_6) and (D_1, \dots, D'_5, D'_6) are both non-standard alternating regular hexagons. Let $d = l(D_5, D_3)$ and $d' = l(D'_5, D_3)$ be arcs as in Definition 2.5.15. Then the arc d is isotopic to the arc d' .

Given that, as previously observed, the subsurface $M(A, B)$ as in Definition 2.5.9 is isotopic to a small regular neighbourhood of the arc $l(A, B)$, the following corollary is immediate.

Corollary 2.5.26. *Under the hypotheses of the previous lemma the two components $M(D_5, D_3)$ and $M(D'_5, D_3)$ are isotopic.*

Before we move into the proof of the previous lemma we will state without proof a technical classification of isotopy classes of some arcs in a pair of pants, whose proof is straightforward, but is nonetheless an useful precise statement to give.

Lemma 2.5.27. *Let S be a surface of type S_2^1 . Let $x \in \partial S$ be a point of the boundary and let y, z be the two marked points. Let l be an essential simple arc in $S \setminus \{y, z\}$ with x as one of the endpoints and the other on ∂S , and whose interior does not intersect either ∂S nor any of the marked points. Let a, b be two properly embedded simple oriented parametrised arcs in $S \setminus \{z\}$ with the first endpoint on x and y as the second endpoint, and whose interior does not intersect either ∂S nor any of the marked points. Suppose that $i(a, l) = i(b, l)$. Then there exists an isotopy relative to the endpoints between the arcs a and b .*

We can now move to the proof of Lemma 2.5.25.

Proof of Lemma 2.5.25. Let $d' = d'_0 \cup \dots \cup d'_j$ be defined analogously to the subdivision of the arc d given in Definition 2.5.19. We fix an orientation on d' by having the first

endpoint being on γ_3 and p as the second one and let us subdivide it as

$$d' = d'_0 \cup \dots \cup d'_{j'},$$

where the d_i 's are closures of connected components of $d \setminus \gamma_2$ ordered using the orientation such that the intersection of two consecutive arcs is a point. Lemma 2.5.20 proves that $j' = j$. Moreover let L'_1, L'_2 the closure of the two components of $D_2 \setminus d'_1$. Thanks to Lemma 2.5.23 we will assume L'_1 to contain all the marked points in V and L'_2 to only contain p .

We first want to show that d_0 is isotopic to d'_0 . Lemma 2.5.12 implies that all the mixed arcs in P are parallel so, up to an isotopy, we have that $d_0 = d'_0$.

We claim that the arc d_1 is isotopic to d'_1 . This naively follows from the fact that those are arcs in the disc D_2 which split the marked points in the same way. The following construction can be visualised in Figure 2.13. More formally speaking let us argue by contradiction assuming the two arcs to be distinct and in minimal position to each other and, up to an isotopy, we can assume they share the endpoints. Let u be the first endpoint, that is, $\{u\} = d_0 \cap d_1 = d'_0 \cap d'_1$ and let v be the first intersection, that is, the point of $d_1 \cap d'_1$ such that the interior of the subarc of d between u and v does not intersect d' . Let $f \subseteq d_1$ (resp. $f' \subseteq d'_1$) be the subarc of d_1 (resp. d'_1) between u and v . Because of how the marked points are split it follows that, $f' \subset L_i$ if and only if $f \subset L'_j$ for $i \neq j$, with $i, j \in \{1, 2\}$. It follows that f and f' cobound a region of either $L_1 \cap L'_2$ or $L_2 \cap L'_1$. Neither of these regions contains any marked point, hence f, f' form an empty bigon contradicting minimal position. It follows that, up to an isotopy, $d_1 = d'_1$.

We have already observe that the arcs $\tilde{d} = d_2 \cup \dots \cup d_j$ and $\tilde{d}' = d'_2 \cup \dots \cup d'_j$ are arcs in the pair of pants X both with w as their first endpoint, that is,

$$\{w\} = d_1 \cap d_2 = d'_1 \cap d'_2 \subset \partial X,$$

and p as the second one, and such that their interiors do not contain r . Let now g be the arc $\gamma_2 \cap L_2 \cap X$ between w and u , where $\{u\} = d_0 \cap d_1$. We have $i(\tilde{d}, g) = j - 1 = i(\tilde{d}', g)$. An application of Lemma 2.5.27 with $l = g$, $a = \tilde{d}$ and $b = \tilde{d}'$ now proves that \tilde{d} is isotopic to \tilde{d}' relatively to the endpoints. It follows that such an isotopy can be glued with the previous ones so the entire arc d is isotopic to \tilde{d} . \square

Thanks to the fact that the isotopy type of the arc d only depends on (D_1, \dots, D_4) we can now prove that the disc D_6 fitting in the alternating hexagon is also uniquely determined.

Lemma 2.5.28. *Let $k \geq 2$ and $n \geq k$. Let (D_1, D_2, D_3, D_4) be four vertices of $\mathcal{D}_k(\Sigma_{k+n+1})$ such that, for $i = 1, 2, 3$, D_i is adjacent to D_{i+1} . Suppose that D_1 a k -punctured disc, and let $\gamma_i = \partial D_i$. Let us suppose that $i(\gamma_2, \gamma_4) > 2$. Let (D_5, D_6) and (D'_5, D'_6) be such that both (D_1, \dots, D_5, D_6) and (D_1, \dots, D'_5, D'_6) are both non-standard alternating regular hexagons. Then the disc D_6 is isotopic to D'_6 .*

Proof. Let $Y = D_1 \cup M(D_5, D_3)$. Let Z be the closure of $\Sigma_{k+n+1} \setminus Y$. The constructions in this paragraph are pictured in Figure 2.14. Since the hexagons are nonstandard the arc d_2 is a nontrivial TD_2 -arc in $P = P(D_2, D_3)$: it follows that the component $M(D_5, D_3)$

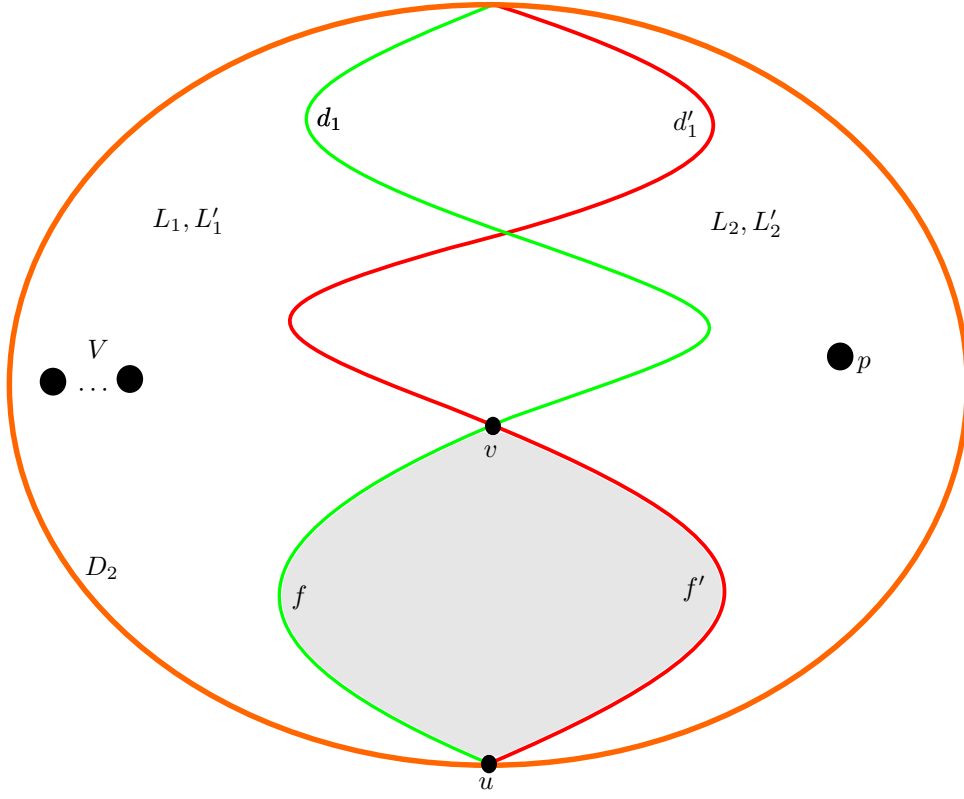


Figure 2.13: The gray component is the empty bigon cobounded by f and f' .

must intersect the arc $z = s(D_1, D_3) \subset D_1$, hence Y is connected. It follows that every connected component of Z can only have one boundary component. An application of Corollary 2.5.22 with $A = D_1$, $B = D_5$ and any component of Z as X proves that every component of Z is either disjoint from D_6 or contains it. Let W be the component such that $D_6 \subseteq W$. Since the hexagon (D_1, \dots, D_6) is regular Proposition 2.4.3 implies that the marked points in Z are exactly the ones of $\pi(D_1)^c \cap \pi(D_5)^c = V \cup \{q\} = \pi(D_6)$. It follows that Z is composed of a number of discs and a single n -punctured disc, that is, W . Since D_6 is a n -punctured disc itself it follows that, up to isotopy, $D_6 = W$.

We can now perform the same construction for the pair of discs (D'_5, D'_6) obtaining a new region W' isotopic to D'_6 . In order to construct W (resp. W') we have just used the disc D_1 and the region $M(D_5, D_3)$ (resp. $M(D'_5, D_3)$) and, thanks to Corollary 2.5.26, We have that, up to isotopy, $M(D_5, D_3) = M(D'_5, D_3)$. From this it also follows that $W = W'$ and hence D_6 is isotopic to D'_6 . \square

We are now ready to complete the proof of Proposition 2.5.18.

Proof of Proposition 2.5.18. Let both $(D_1, \dots, D_4, D_5, D_6)$ and $(D_1, \dots, D_4, D'_5, D'_6)$ be two regular alternating hexagons. Lemma 2.5.28 proves that $D_6 = D'_6$, and an application of Lemma 2.5.3 with $A = D_1$ and $B = D_6 = D'_6$ now proves that it also holds $D_5 = D'_5$, and the proof is complete. \square

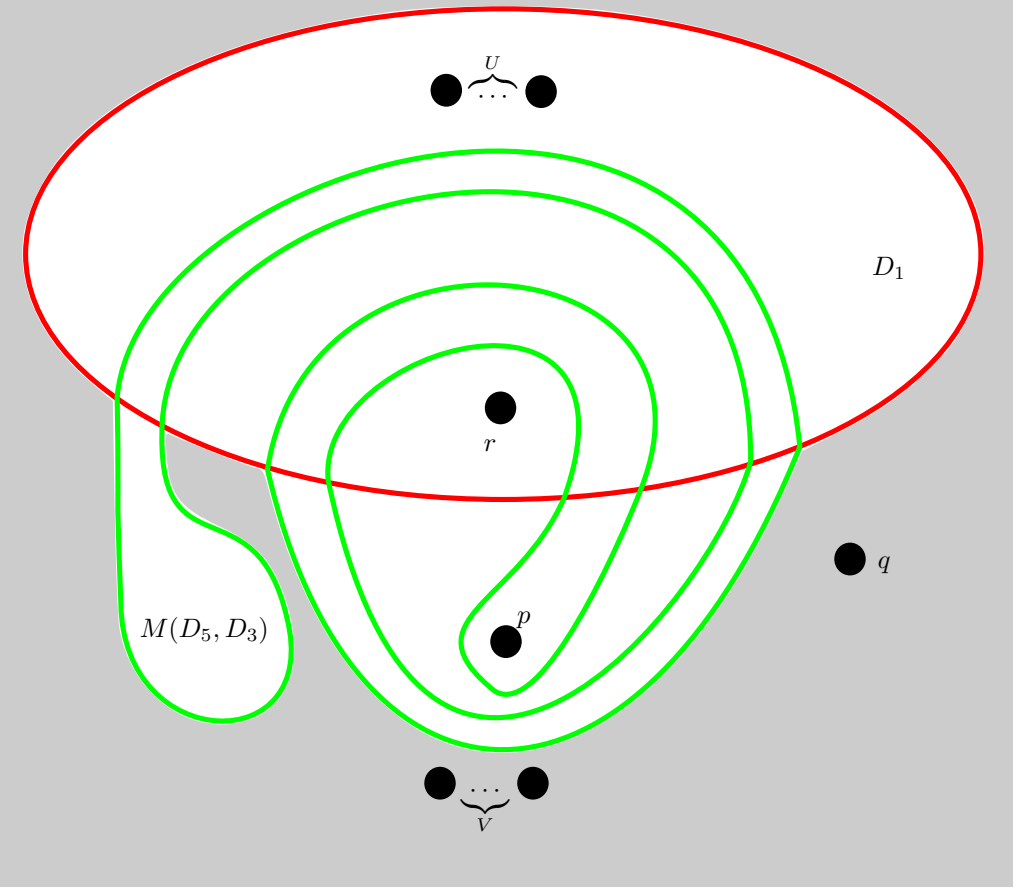


Figure 2.14: The gray region is W .

2.6 Irregular Alternating Hexagons

In this section we will complete the proof of Proposition 2.5.1 by dealing with the case of irregular alternating hexagons, that is,, hexagons whose projection to the extended Kneser graph is not an hexagon, as in Definition 2.4.5. An example of an irregular alternating hexagon, alongside its projection to the extended Kneser graph, is shown in Figure 2.15.

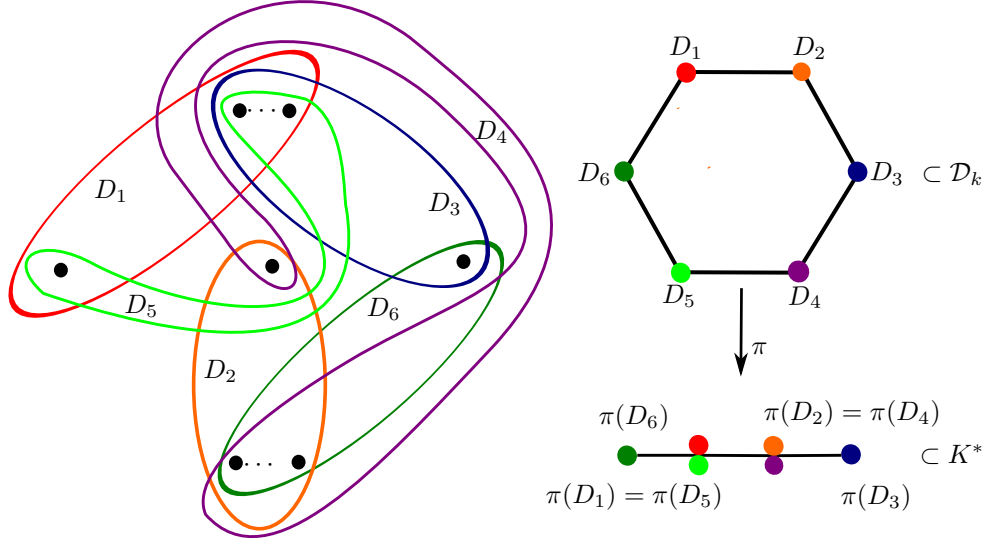


Figure 2.15: An irregular alternating hexagon.

We start studying irregular hexagons from their projections to the extended Kneser graph. In Proposition 2.4.6 we proved that the image of such a projection must be a tree. In Corollary 2.4.7 we have identified the possible configurations for the image of this projection. These are a single edge, two consecutive edges, three edges with a vertex in common (a “tripod”), or three consecutive edges, as in Figure 2.15. When $k \geq 3$, the last case is the only one which can be realised. However, as already stated, the proof of this fact cannot rely on combinatorial arguments only, but requires a good amount of topology, and will take most of this section.

We will begin with a statement implying that an irregular alternating hexagon cannot project to either a single edge or two adjacent edges. We remark that the following result still holds for $k = 2$.

Proposition 2.6.1. *Let $k \geq 2$ and $n \geq k$. Then there exists no alternating hexagon (D_1, \dots, D_6) in $\mathcal{D}_k(\Sigma_{k+n+1})$ such that $\pi(D_1) = \pi(D_3) = \pi(D_5)$ and $\pi(D_2) = \pi(D_4)$.*

In order to prove the previous proposition we will argue by contradiction. Let (D_1, \dots, D_6) be an alternating 6-cycle in $\mathcal{D}_k(\Sigma_{k+n+1})$ such that $\pi(D_1) = \pi(D_3) = \pi(D_5)$ and $\pi(D_2) = \pi(D_4)$. Moreover, let us assume $D_1 \neq D_3$ and $D_2 \neq D_4$. Let $\gamma_i = \partial D_i$. We start by defining some objects that will be useful later in the proof. A picture of the following constructions is provided in Figure 2.16.

Let p be the only marked point in $\pi(D_1) \setminus \pi(D_2)$.

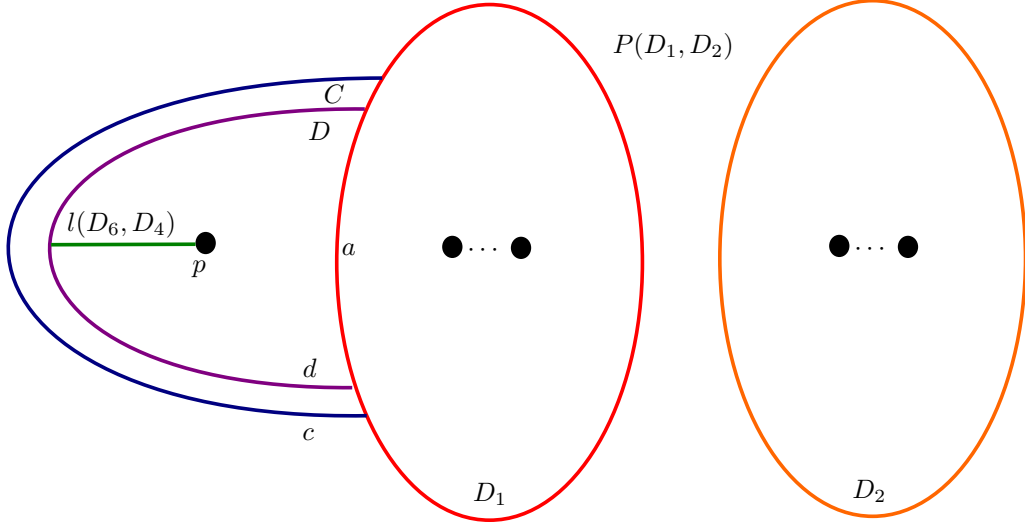


Figure 2.16: We recall that the arc $l(D_6, D_4)$ is only defined when $\pi(D_4) \neq \pi(D_6)$.

All the arcs of $\gamma_3 \cap D_1^c$ must be parallel to each other due to Lemma 2.5.12 since they are nontrivial TD_1 -arcs in the pair of pants $P(D_1, D_2) = D_1^c \cap D_2^c$ (see Definition 2.5.10). The nontriviality follows from the fact that $D_1 \neq D_3$.

Let c be the innermost of these arcs, that is,, such that the interior the component of $P(D_1, D_2) \setminus c$ containing p is disjoint from γ_3 . Let us denote this component, which is a once-punctured disc, with C .

Let a be the arc of γ_1 between the endpoints of c whose interior does not intersect γ_3 , that is,, such that $\partial C = a \cup c$.

The region $Y = D_3 \cup C$ is a disc containing the punctures

$$\pi(D_3) \cup \{p\} = \pi(D_1) \cup \{p\} = \pi(D_2)^c = \pi(D_2^c).$$

We claim that $C \cap D_2 = \emptyset$: if not Corollary 2.5.22, applied with $A = D_1$, $B = D_3$ and $X = C$ would imply that $D_2 \subseteq C$, which is clearly impossible. It follows that the disc Y is contained in D_2^c and it contains all the marked points in it, hence Y is isotopic to D_2^c .

We claim that $D_4 \cap D_2^c \neq \emptyset$. Indeed, if such intersection were empty it would hold that $D_4 \subseteq D_2$, but since those discs have the same number of marked points it would follow that $D_2 = D_4$, contradicting the 6-cycle being embedded.

Moreover, since D_4 is disjoint from D_3 it follows that $D_4 \cap C \neq \emptyset$. Hence, there exists an arc of γ_4 inside C with both endpoints on a . This arc is isotopic to c , as it would otherwise form an empty bigon with a . Let d be the innermost one among these arcs, that is,, the arc such that the interior of the component of $C \setminus d$ containing p , which we will denote with D , does not intersect γ_4 .

We are now ready to prove Proposition 2.6.1. In order to do so we need to consider two cases: the first one is when $\pi(D_6) = \pi(D_2) = \pi(D_4)$, while the second one is for $\pi(D_6) \neq \pi(D_2) = \pi(D_4)$. We start by proving that the former is impossible.

Lemma 2.6.2. *Let $k \geq 2$ and $n \geq k$. Let (D_1, \dots, D_6) be an alternating hexagon in $\mathcal{D}_k(\Sigma_{k+n+1})$ such that $\pi(D_1) = \pi(D_3) = \pi(D_5)$ and $\pi(D_2) = \pi(D_4)$. Then we have that*

$$\pi(D_6) \neq \pi(D_2).$$

Proof. The constructions in this proof are pictured in Figure 2.17

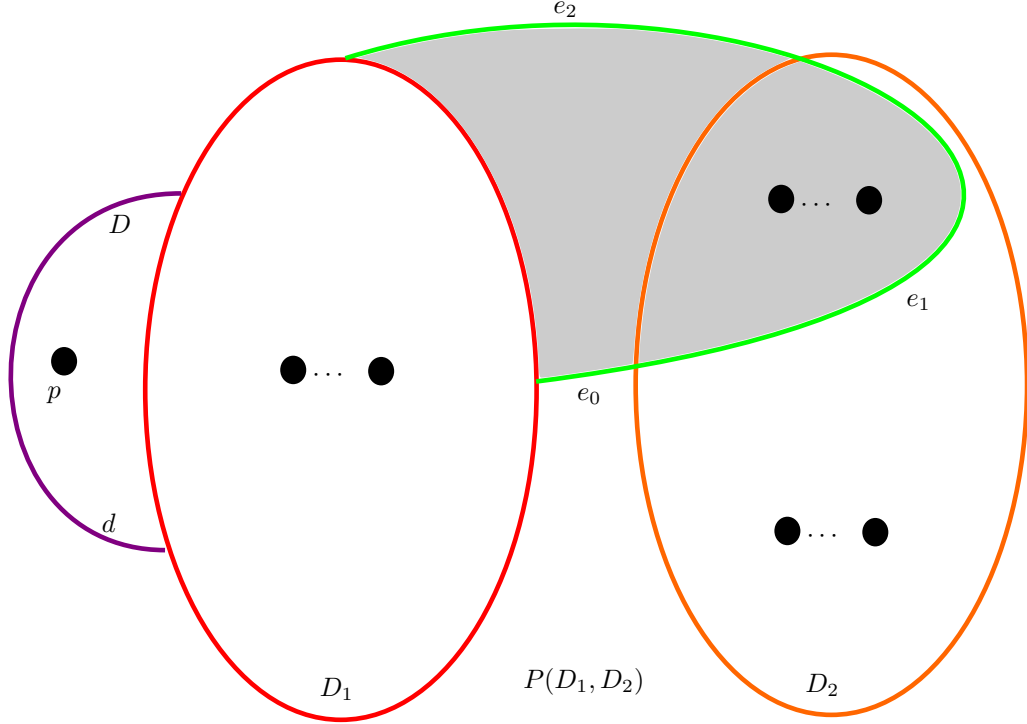


Figure 2.17: The gray region is a possible choice for Z .

We claim that the discs D_2 and D_5 intersect. Indeed, if they were disjoint, due to the uniqueness of common neighbours property stated in Lemma 2.5.3, we would have $D_2 = \Delta(D_1, D_5) = D_6$, where $\Delta(D_1, D_5)$ is as in Definition 2.5.4, contradicting the fact that (D_1, \dots, D_6) is an alternating hexagon. The claim is proven. In particular it follows that D_2 and γ_5 intersect.

Let e be an arc of $\gamma_5 \cap D_1^c$ intersecting D_2 . We can write $e = e_0 \cup \dots \cup e_m$ where the e_i 's are closure of connected components of $e \setminus \gamma_2$ such that $e_i \cap e_j$ is nonempty if and only if $|i - j| = 1$ and in that case it is a single point. Every component of $e \cap D_2^c$ with both endpoints on γ_2 cannot intersect either D_1 nor D_4 , and hence d . From an application of Lemma 2.5.21 with $A = D_1$, $B = D_4$ and $X = D$ it follows that such arc must lie in the annulus $P(D_1, D_2) \setminus D$: since such subspace contains no marked points, then the arc must be trivial. From minimal position it follows that every arc of $e \cap D_2^c$ has one endpoint on γ_2 and the other on γ_1 . It follow that $m = 2$.

It follow that the set $D_1^c \setminus e$ has two connected components, and both of them intersect D_2 : let Z be one of them. Since the arc e_1 does not form a bigon with γ_2 it follows that both components of $D_2 \setminus e$ contain some of the marked points in $\pi(D_2)$, but cannot contain all. In particular, Z contains a strict subset of the marked points in $\pi(D_2)$. We have $\partial(Z \cup D_1) \subseteq D_1 \cup D_5$, hence $\partial(Z \cup D_1)$ is disjoint from γ_6 . Lemma 2.5.21, applied with $X = Z \cup D_1$ and $Y = D_6$ proves that either $(Z \cup D_1) \subseteq D_6$ or $D_6 \subseteq (Z \cup D_1)$. In either case it is impossible to have $\pi(D_6) = \pi(D_2)$. Indeed, if $(Z \cup D_1) \subseteq D_6$ a

contradiction follows from the fact that $Z \cup D_1$ contains some marked points not in $\pi(D_2)$; if $D_6 \subseteq (Z \cup D_1)$ contradiction descends from the fact that $Z \cup D_1$ does not contain all the marked points in $\pi(D_2)$. \square

We are now left proving that the case of hexagons where $\pi(D_6) \neq \pi(D_2) = \pi(D_4)$ is impossible, as well.

Lemma 2.6.3. *Let $k \geq 2$ and $n \geq k$. Let (D_1, \dots, D_6) be a 6-cycle in $\mathcal{D}_k(\Sigma_{k+n+1})$ such that $\pi(D_1) = \pi(D_3) = \pi(D_5)$ and $\pi(D_2) = \pi(D_4) \neq \pi(D_6)$. Then such cycle is not embedded.*

Proof. If the first four vertices (D_1, \dots, D_4) are not pairwise distinct we are clearly done. We are hence going to assume that they are, and prove that $D_3 = D_5$. Notice that $D_6 \neq D_2$ and $D_6 \neq D_4$ since they have different projections to the extended Kneser graph.

Since $\pi(D_6) \subseteq \pi(D_1)^c = \pi(D_2) \cup \{p\}$ but $\pi(D_6) \neq \pi(D_2)$ it follows that $p \in \pi(D_6)$. We recall that $M(D_6, D_4)$, defined as in Definition 2.5.9, is the component of $D_6 \cap D_4^c$ containing the puncture p . It follows that $M(D_6, D_4) \cap D \neq \emptyset$. We claim that this component $M(D_6, D_4)$ is contained in D . If this were not true we would have that the interior of the arc $l(D_6, D_4)$, which is the only arc in $M(D_6, D_4)$ between p and γ_4 as in Definition 2.5.15, intersects $\partial D \subseteq d \cup \gamma_1$. It would follow that the arc $l(D_6, D_4)$ intersects $a \subset D_1$, which is absurd since $l(D_6, D_4) \subset D_6$. The claim is proven.

In particular, we have that $M(D_6, D_4)$ is disjoint from D_3 . The region

$$Z = D_4 \cup M(D_6, D_4)$$

contains exactly all the marked points of $\pi(D_4) \cup \{p\} = \pi(D_3^c)$ and it is contained in D_3^c : it follows that $Z = D_3^c$. From Lemma 2.5.16, applied with $A = D_6$, $B = D_4$ and it follows that Z is isotopic to D_5^c , too. From this it follows that $D_3 = D_5$, hence the cycle is not embedded. \square

This completes the proof of Proposition 2.6.1.

We can now move to proving that there exists no “tripods” in the graph of discs, that is,, the content of the following proposition. For this result, we will need the stronger hypothesis that $k \geq 3$. Indeed, an example of an alternating hexagon in $\mathcal{C}_2(\Sigma_5)$ which projects to a “tripod” is provided in Figure 2.18. This is the only reason why we need such hypothesis in Corollary 2.6.10 and Proposition 2.6.11, although the proof of the latter could probably be adapted to the case of $k = 2$ with some extra work. However, we have already noticed how the result with $k \geq 3$ will still suffice our needs.

Proposition 2.6.4. *Let $k \geq 3$ and $n \geq k$. There does not exists any alternating hexagon (D_1, \dots, D_6) in $\mathcal{D}_k(\Sigma_{k+n+1})$ such that $\pi(D_1) = \pi(D_3) = \pi(D_5)$ and such that we have $\pi(D_i) \neq \pi(D_j)$ for every $i, j \in \{2, 4, 6\}$ with $i \neq j$.*

We have the following information about the projection to the extended Kneser graph.

Lemma 2.6.5. *Let (D_1, \dots, D_6) be an alternating hexagon in the graph $\mathcal{D}_k(\Sigma_{k+n+1})$ such that we have $\pi(D_1) = \pi(D_3) = \pi(D_5)$ and such that $\pi(D_i) \neq \pi(D_j)$ for every $i, j \in \{2, 4, 6\}$ with $i \neq j$. Let $\gamma_i = \partial D_i$.*

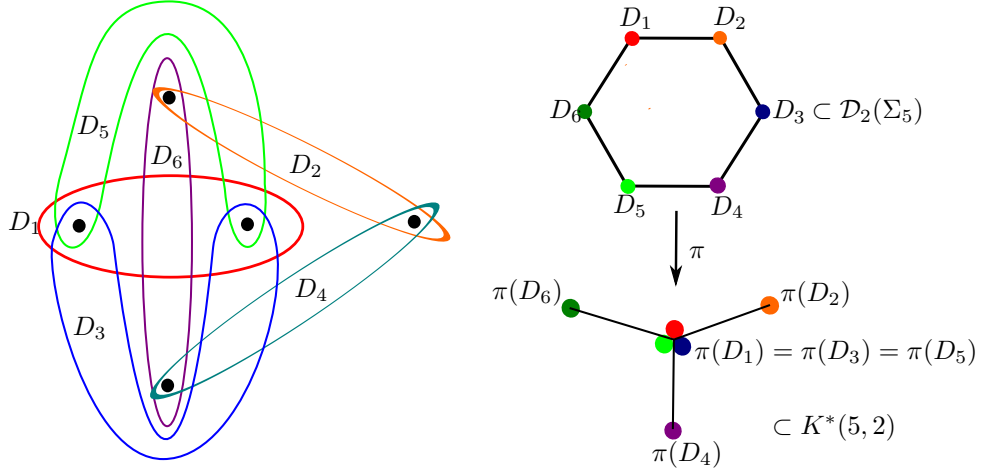


Figure 2.18: A hexagon in $\mathcal{D}_2(\Sigma_5)$ projecting onto a “tripod” in $K^*(5, 2)$.

There exists a partition of the marked points as $\Pi = U \cup V \cup \{p, q, r\}$ such that

$$\begin{aligned}\pi(D_1) &= \pi(D_3) = \pi(D_5) = U; \\ \pi(D_2) &= V \cap \{p, q\}; \\ \pi(D_4) &= V \cap \{p, r\}; \\ \pi(D_6) &= V \cap \{q, r\}\end{aligned}$$

where either $|U| = k$ and $|V| = n - 2$ or $|U| = n$ and $|V| = k - 2$.

We recall that we are assuming $n \geq k \geq 3$, hence all the sets mentioned in the aforementioned lemma are nonempty.

Proof of Lemma 2.6.5. Since $(\pi(D_1), \dots, \pi(D_6))$ is an alternating 6-cycle in the extended Kneser graph $K^*(k + n + 1, k)$, for i even we have that $\pi(D_i)$ and $\pi(D_{i \pm 2})$ are distinct and have $\pi(D_{i \pm 1})$ as a common neighbour. We can now apply Lemma 2.4.1 proving that

$$|\pi(D_i) \cap \pi(D_{i \pm 2})| = |\pi(D_i)| - 1.$$

Since $\pi(D_2) \cup \pi(D_4) \cup \pi(D_6) = \pi(D_1)^c$, an application of Inclusion-Exclusion Principle yields

$$|\pi(D_2) \cap \pi(D_4) \cap \pi(D_6)| = |\pi(D_2)| - 2,$$

and the lemma follows. \square

The following constructions, which are somewhat similar to the ones performed before, are pictured in Figure 2.19. We retain the hypotheses of the previous lemma.

The marked point $r = p(D_4, D_2) \in \pi(D_4) \setminus \pi(D_2)$, as in Definition 2.5.9, does not belong to either $\pi(D_3) = \pi(D_1)$ or to $\pi(D_2)$, hence it is the only marked point in the pair of pants $P(D_1, D_2)$, as in Definition 2.5.10.

All the arcs of $\gamma_3 \cap D_1^c$ must be parallel due to Lemma 2.5.12 since they are nontrivial TD_1 -arcs in the pair of pants $P(D_1, D_2)$. Analogously to what we have done before, let

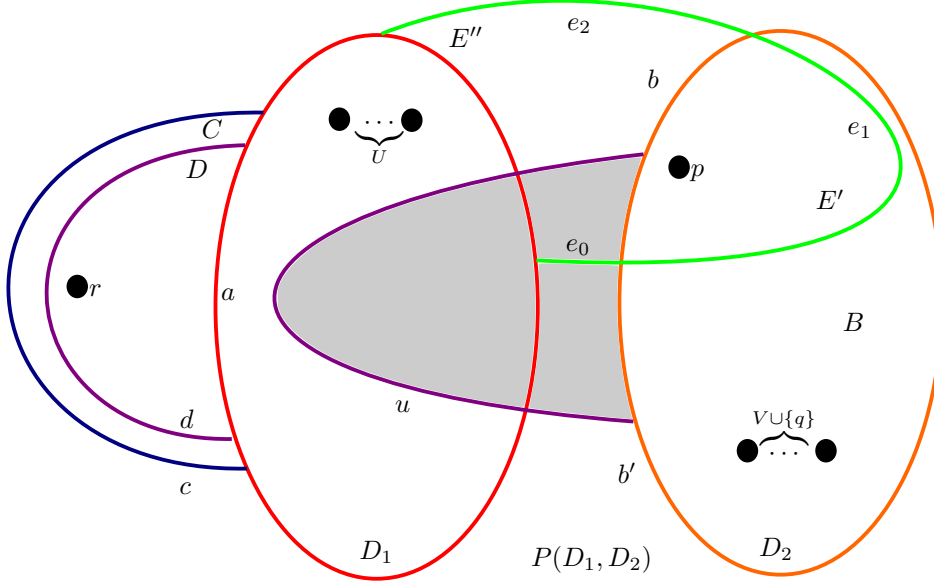


Figure 2.19: The gray region is Z . We recall that $E = E' \cup E''$.

c be the innermost of these arcs, that is,, such that the interior of the component of $P(D_1, D_2) \setminus c$ containing r is disjoint from D_3 . Let us denote this component, which is a once-punctured disc, with C . Let a be the arc of γ_1 between the endpoints of c whose interior does not intersect D_3 , that is,, such that $\partial C = a \cup c$.

Since $r \in \pi(D_4)$ it obviously follows that $M(D_4, D_2) \cap C \neq \emptyset$, where $M(D_4, D_2)$ is the component of $D_4 \cap D_2^c$ containing r , as in Definition 2.5.9. Let D be the component of $M(D_4, D_2) \cap C$ containing r . The region D is a once-punctured disc and ∂D is composed of a single arc of γ_4 , which we will denote with d , and a single arc of γ_1 .

Similarly, the only marked point which is contained in the pairs of pants $P(D_1, D_6)$ is $p = p(D_2, D_6) = p(D_4, D_6)$. All the arcs of $\gamma_5 \cap D_1^c$ are isotopic thanks to Lemma 2.5.12, since they are all TD_1 -arcs in $P(D_1, D_6)$. Let e be the innermost one, that is,, such that the interior of the component of $D_1^c \setminus e$ containing p , whose closure we will denote by E , does not intersect γ_5 . Similarly to what we have done before, let $e = e_0 \cup \dots \cup e_m$, where e_i is the closure of a component of $e \setminus \gamma_2$ such that $e_i \cap e_j$ is nonempty if and only if $|i - j| = 1$ and in that case it is a single point. We have the following.

Lemma 2.6.6. *We have that $m = 2$. The arc $e_1 \subset D_2$ is such that $D_2 \setminus e_1$ is composed of exactly two connected component, one of which is $E' = E \cap D_2$ and only contains the marked point p . The closure of the other component, which is $D_2 \setminus E'$, contains all the marked points in $\pi(D_2) \cap \pi(D_6) = V \cup \{q\}$.*

We will denote with B the closure of the component $D_2 \setminus E'$.

Proof of Lemma 2.6.6. If it held $m > 2$ it would follow that the arc e_2 would be a TD_2 -arc in $P(D_1, D_2)$. Since $e \subset D_5$ it does not intersect D_4 , so the arc e_2 would actually lie in the annulus $P(D_1, D_2) \setminus D$, hence it would form a bigon with γ_2 , contradicting minimal position.

The arc e_1 separates D_2 into two components. The region $E' = E \cap D_2$ is a subset of E so it can only contain p . All the marked points in the other component are those of $\pi(D_2) \cap \pi(D_6) = V \cup \{q\}$. \square

The arcs e_0 and e_2 are mixed arcs in $P(D_1, D_2)$, hence E is a punctured disc. We observe that its boundary is composed of e and a single arc of γ_1 . We claim the following.

Lemma 2.6.7. *The component D is disjoint from E . In particular the subspace defined by $E \cap P(D_1, D_2)$ is a disc with no marked points. Moreover, the boundary of said subspace is composed of the arcs e_0, e_2 , a single arc of γ_1 and a single arc of γ_2 .*

We will denote the disc $E \cap P(D_1, D_2)$ as E'' .

Proof of Lemma 2.6.7. Since $D \subseteq D_4$ we have $D \cap e = \emptyset$, hence either $D \subseteq E$ or they are disjoint. If it held $D \subseteq E$ we would have that $r = p(D_6, D_2) = p(D_4, D_2) \in E$. In particular, E would contain exactly the two marked points p, r . Since $\partial E \subseteq e \cup \gamma_1$, Lemma 2.5.21, applied with $X = E$ and $Y = D_6$ would imply that either $E \subseteq D_6$ or $D_6 \subseteq E$. The former is impossible since E contains $p = p(D_2, D_6) \notin \pi(D_6)$. The latter is also impossible, since D_6 contains at least $k \geq 3$ marked points, while E only contains the marked points r, p and $p \notin \pi(D_6)$.

For the second part it is enough to notice that E only contains the marked point $p = p(D_2, D_6) \in D_2$ so $p \notin P(D_1, D_2)$. It follows that $E \cap P(D_1, D_2)$ contains no marked points. \square

Since the arcs c and d are isotopic we can assume that c is contained in a small regular neighbourhood of d , and hence disjoint from the arc e . Moreover, every arc of $\gamma_5 \cap D_1^c$ is isotopic to e , so we can assume it to be contained in a small regular neighbourhood of e , and in particular it to be disjoint from C .

Since $p(D_2, D_6) = p(D_4, D_6) = p$, it follows that $p \in D_4$ and hence $D_4 \cap E' \neq \emptyset$. We claim that γ_4 intersects the arc $b = \gamma_2 \cap E$. If not, Lemma 2.5.21, applied with $X = E'$ and $Y = D_4$ would imply that $E' \subseteq D_4$ which is impossible since $\partial E' \cap D_5 = e_1 \neq \emptyset$.

We recall that B denotes the closure of $D_2 \setminus E'$. Let $b' = \gamma_2 \cap B$. Since $k \geq 3$ the set V of Lemma 2.6.5 is not empty, hence it follows from Lemma 2.6.6 that at least one puncture in B must belong to $\pi(D_4)$. In particular we have that $D_4 \cap B \neq \emptyset$ and $\gamma_4 \cap b' \neq \emptyset$. The second claim follows with an argument very similar to the one in the previous paragraph. Indeed, either $\gamma_4 \cap b' \neq \emptyset$ or Lemma 2.5.21, applied with $X = B$ and $Y = D_4$ implies that $B \subseteq D_4$, which is impossible since $q = \pi(D_2, D_4) \in B$ but $q \notin D_4$.

Since γ_4 intersects both b and b' it follows that there must exist an arc of γ_4 with one endpoint on b and the other on b' . Since this arc cannot intersect $e \subset D_5$ it follows that it must intersect $D_1 \cup C$. We will eventually prove this causes a contradiction. We first start by proving that such an arc must intersect C , that is,, the following.

Lemma 2.6.8. *Any arc of γ_4 with one endpoint on b and the other on b' intersects C .*

Proof. For a picture of the constructions to follow we still refer to Figure 2.19.

We argue by contradiction: let u be an arc of γ_4 with one endpoint on b and the other on b' not intersecting C . Up to substituting u with a subarc, we can assume without loss

of generality that the interior of u does not intersect γ_2 . Since an arc from b to b' inside D_2 must intersect $e_1 \subset \gamma_5$ it follows that $u \subset D_2^c$.

The space $D_2^c \setminus u$ has exactly two connected components: let Z be the closure of the one not containing C . We have that $\partial Z \subseteq \gamma_2 \cup \gamma_4$ and $D_3 \subseteq Z$, since Z does not contain the arc c . It follows from Corollary 2.5.22, applied with $A = D_2$, $B = D_4$, and $X = Z$, that $Z \subseteq D_3^c$. Since $Z \cap D_2$ has empty interior the only puncture in $\pi(D_3^c) = U \cup \{p, q, r\}$ which can be contained in Z is r , but that puncture is contained in the interior of C , which is disjoint from Z by definition. It follows that Z cannot contain any puncture, hence u forms an empty bigon with γ_2 . We have reached a contradiction to minimal position and hence proven the lemma. \square

The last ingredient we are now missing for the proof of Proposition 2.6.4 is the following lemma. It implies that there cannot exist an arc of γ_4 with one endpoint on b and the other on b' intersecting C , providing the contradiction we need.

Lemma 2.6.9. *There exists no pair of disjoint arcs $v, w \subseteq \gamma_4$, such that v has one endpoint on a and the other on b , and w has one endpoint on a and the other on b' .*

Proof. A picture of the constructions used in this proof is provided in Figure 2.20.

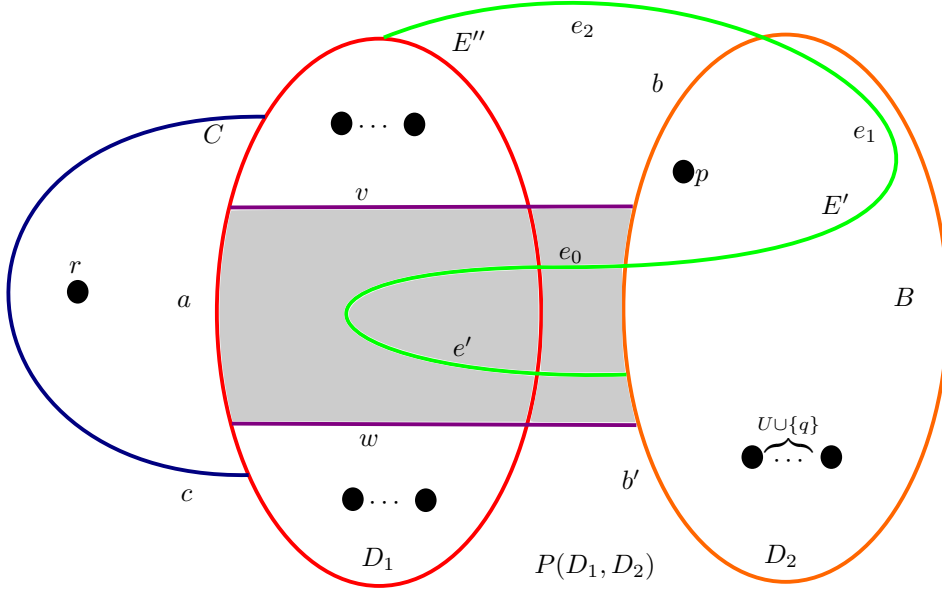


Figure 2.20: The gray region is Z .

Up to considering subarcs, if needed, we can assume, without loss of generality, that the interior of neither v nor w intersects neither C nor D_2 . Let $a' \subseteq a$ be the arc between the endpoints of v and w on a . Since $\partial a \subseteq \gamma_3$ we have that a' is contained in the interior of the arc a . In particular, we have $a' \cap D_3 = \emptyset$.

Let $z = v \cup w \cup a'$: this is an arc with one endpoint on $b \subseteq \gamma_2$ and the other on $b' \subseteq \gamma_2$. It follows that the space $D_2^c \setminus z$ has exactly two connected components. Let Z be the closure of the component of $D_2^c \setminus z$ not containing the interior of C .

We have $\partial Z \subseteq \gamma_2 \cup \gamma_4 \cup a'$, hence ∂Z is disjoint from γ_3 . Moreover, since $Z \cap D_2 \neq \emptyset$, we have that $Z \not\subseteq D_3$, hence $Z \cap D_3^c \neq \emptyset$. Lemma 2.5.21, applied with $X = Z$ and

$Y = D_3^c$ implies that either $D_3^c \subseteq Z$ or $Z \subseteq D_3^c$. The former option is impossible since $r \in \pi(D_3^c) = U \cup \{p, q, r\}$ but $r \in C$ so $r \notin Z$. It follows that $Z \subseteq D_3^c$. Since Z is disjoint from the interior of D_2 it cannot contain any of the marked points in $\pi(D_2) = U \cup \{p, q\}$, and it does not contain $r \in C$ by definition. It follows that Z does not contain any marked point.

We claim that $w \cap E'' = \emptyset$, where $E'' = E \cap P(D_1, D_2)$. Indeed, since we assumed the interior of w not to intersect D_2 , and the endpoints of w are on a and b' , it follows that every component of $w \cap E''$ must have both endpoints on $\gamma_1 \cap E''$. From this it follows that such a component arc would form an empty bigon with γ_1 , since the region E'' does not contain any marked point thanks to Lemma 2.6.7, and $\partial E'' \cap \gamma_1$ is connected. We have found a contradiction and proven the claim.

We claim that the intersection $v \cap E$ is composed of a single arc inside E'' with one endpoint on γ_1 and the other on $b \subseteq \gamma_2$. Indeed, the intersection $v \cap E''$ is a collection of arcs, one of which has one endpoint on γ_1 and the other on γ_2 . Any other component would be an arc in E'' with both endpoints on γ_1 . This follows from the fact that $\partial E'' \subseteq \gamma_1 \cup \gamma_2 \cup \gamma_5$, as we noticed in Lemma 2.6.7, alongside the fact that $v \subset \gamma_4$, hence it is disjoint from γ_5 , and the hypothesis that the interior of v does not intersect D_2 . Once again, since E'' is a disc with no marked points thanks to Lemma 2.6.7, and $\partial E'' \cap \gamma_1$ is connected, it follows that any arc with both endpoints on γ_1 would form an empty bigon with that curve, contradicting minimal position. The claim is proven.

It follows from the previous two claims that the space $E'' \setminus (v \cup w) = E'' \setminus v$ has two connected components and one of these components is contained in Z . Up to switching e_0 with e_2 we will assume, without loss of generality, that the aforementioned component contains the arc e_0 . Let e' be the arc of $\gamma_5 \cap D_2^c$ containing e_0 .

We have previously noticed that we can assume γ_5 to be disjoint from C . From this and the fact that γ_5 does not intersect the arcs $u, v \subseteq \gamma_4$ we have that $e' \subset Z$ and it has both endpoints on γ_2 . Since Z does not contain any marked points, and $\partial Z \cap D_2$ is connected, the arc e' forms an empty bigon with γ_2 , hence contradicting minimal position. We have reached a contradiction and proven the lemma. \square

We are now ready to conclude the proof of Proposition 2.6.4.

Proof of Proposition 2.6.4. Let us argue by contradiction and assume that a hexagon (D_1, \dots, D_6) such as in the statement exists, so all the previous constructions apply.

Since both $\gamma_4 \cap b \neq \emptyset$ and $\gamma_4 \cap b' \neq \emptyset$ we have already noticed that there exists an arc $t \subseteq \gamma_4$ from b to b' . Lemma 2.6.8 implies that $t \cap C \neq \emptyset$, so in particular $t \cap a \neq \emptyset$.

By considering the subarcs between any of the two endpoints and the closest intersection with a it follows that there exists two arcs $v, w \subseteq t$ with one endpoint on a and the other on b (resp. b'). These arcs can intersect at most in their endpoint on a , but in that case $t \subseteq \gamma_4$ and $a \subseteq \gamma_1$ would not be transverse, contradicting general position. It follows that the arcs v, w are disjoint. Since Lemma 2.6.9 proves such arcs cannot exist, we have reached a contradiction and concluded the proof. \square

Proposition 2.6.1 and Proposition 2.6.4, alongside Corollary 2.4.7 complete the study the projection of irregular hexagons to the extended Kneser graph, and proving that the only possible configuration was the “folding”.

Corollary 2.6.10. *Let $k \geq 3$ and $n \geq k$. Let (D_1, \dots, D_6) be an irregular alternating hexagon in $\mathcal{D}_k(\Sigma_{k+n+1})$. Then, up to a translation on the indices, we have that*

$$\pi(D_1) = \pi(D_5) \neq \pi(D_3)$$

and

$$\pi(D_2) = \pi(D_4) \neq \pi(D_6).$$

Moreover the projections $\pi(D_1), \pi(D_2), \pi(D_3), \pi(D_6)$ are all different from each other.

We now have complete knowledge about the possible combinatorial configurations which can arise from the projection of an irregular alternating hexagon to the extended Kneser graph. We are then ready to prove a criterion to combinatorially distinguish irregular hexagons from standard ones, hence completing the proof of Proposition 2.5.1.

We would like to have a combinatorial criterion to tell the irregular hexagons apart from all the regular ones, that is,, proving that these classes are preserved under graph automorphism. Once we have proven the rigidity of the graphs this must be true, *a posteriori*, but unfortunately we have not been able to find a simple direct *a priori* proof of that fact. However, since our goal is to show that the standard hexagons are preserved by graph automorphisms, being able to only combinatorially distinguish between standard hexagon and irregular ones will still suffice our needs, since we have proven how to distinguish regular hexagons from standard irregular ones in the previous section. This is the content of the next proposition.

Proposition 2.6.11. *Let $k \geq 3$ and $n \geq k$. Let (D_1, D_2, D_3, D_4) be four vertices of $\mathcal{D}_k(\Sigma_{k+n+1})$ such that D_i is adjacent to D_{i+1} for $i = 1, 2, 3$. Moreover, suppose that $\pi(D_2) = \pi(D_4)$ and such that $\pi(D_1), \pi(D_2), \pi(D_3)$ are distinct. Then there exists at most one pair of discs (D_5, D_6) such that (D_1, \dots, D_6) is an alternating hexagon.*

Proof. Let $(D_1, \dots, D_4, D_5, D_6)$ be an alternating hexagons with the required properties. Let $\gamma_i = \partial D_i$. Let us first notice that, due to Corollary 2.6.10, it follows that $\pi(D_1) = \pi(D_5)$ and $\pi(D_6)$ must be different from all the other projections. Moreover $\pi(D_4) \cap \pi(D_1) = \emptyset$.

The discs D_1 and D_4 are not disjoint. If not, both D_2 and D_4 would be disjoint from both D_1 and D_3 , hence the uniqueness of common neighbours stated in Lemma 2.5.3 would imply that $D_4 = D_2$: this is a contradiction to the injectivity of the 6-cycle. It follows that the set $D_1 \cap D_4$ is nonempty but it contains no punctures. Hence, if the boundary of a component of $D_1 \cap D_4$ were composed of a single arc of γ_1 and a single arc of γ_4 the two curves would form an empty bigon, contradicting minimal position. It follows that we must have $i(\gamma_1, \gamma_4) \geq 4$.

We claim the set $D_4 \cap D_1^c$ is composed of at least two connected components. Indeed, any connected component of $\gamma_1 \cap D_4$ disconnects D_4 . Since we have just proven that $i(\gamma_1, \gamma_4) \geq 4$ it follows that $\gamma_1 \cap D_4$ and $\gamma_4 \cap D_1^c$ both have at least two connected components. If $D_4 \cap D_1^c$ were connected then it would contain all the marked points in $\pi(D_4)$. Moreover, Jordan Curve Theorem would immediately imply that $D_1^c \cap D_4^c$ is not connected, and that the boundary of every component is composed of a single arc of γ_1 and a single arc of γ_4 . Since there exists only one marked point in $\pi(D_1^c) \cap \pi(D_4^c)$ it

follows that γ_1 and γ_4 would form an empty bigon, contradicting minimal position. The claim is proven.

The constructions performed in the next part of the proof are pictured in Figure 2.21.

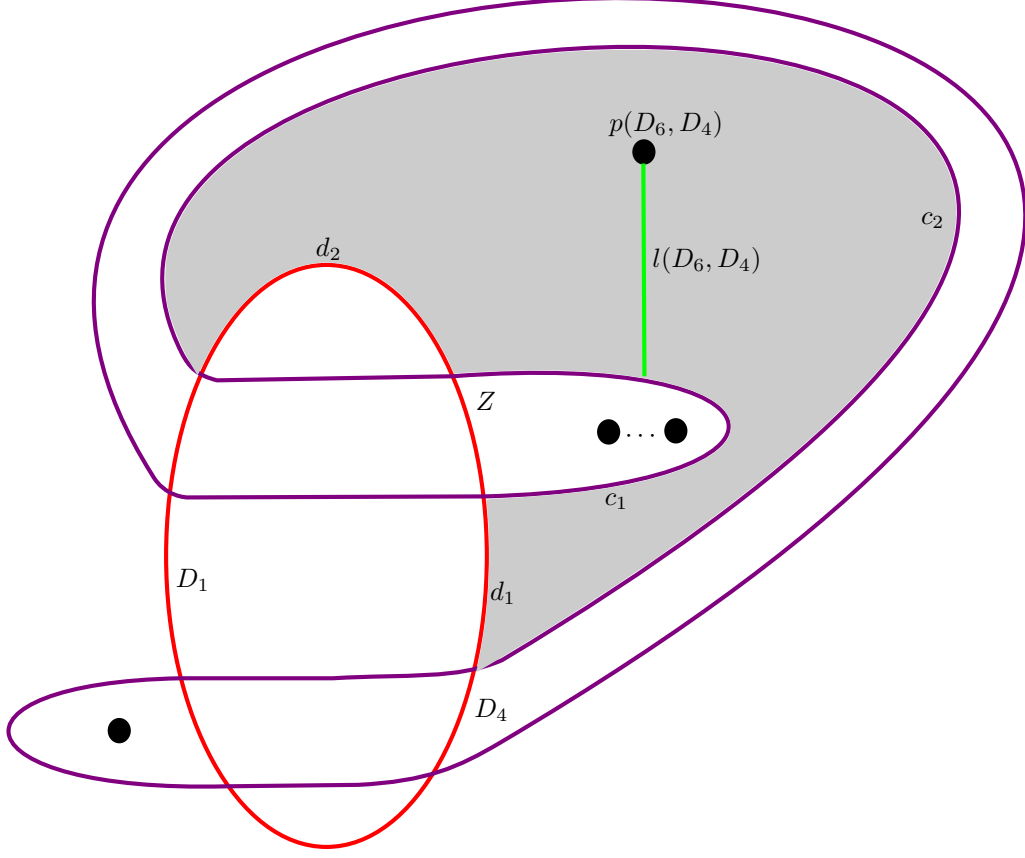


Figure 2.21: The gray region is Y .

The set $D_1^c \cap D_4^c$ only contains the puncture $p(D_6, D_4)$, hence it is composed of a certain number of discs with no marked points and a single once-punctured disc, which we will denote with Y . The boundary of Y can be written as $c_1 \cup d_1 \cup \dots \cup c_m \cup d_m$ where the c_i 's are pairwise disjoint arcs of γ_4 and the d_i 's are pairwise disjoint arcs of γ_1 , such that $c_i \cap d_i$ and $d_i \cap c_{i+1}$ are composed of a single endpoint. Distinct arcs c_i 's belong to the boundaries of distinct connected component of $D_4 \cap D_1^c$.

The arc $l(D_6, D_4)$, defined as in Definition 2.5.15 must be contained in the once-punctured disc Y . It follows that the isotopy type of the arc $l(D_6, D_4)$, and hence of the region $M(D_6, D_4)$ is now uniquely determined by the arc c_i containing the endpoint that is, not $p(D_6, D_4)$, hence by the component of $D_4 \cap D_1^c$ containing it, which we will denote with Z . Without loss of generality let us assume that the arc c_1 is the one containing one endpoints of $l(D_6, D_4)$.

All the connected components of $\gamma_6 \cap D_4^c$ are TD_4 -arcs in the pair of pants $P(D_4, D_5)$, as in Definition 2.5.10. Lemma 2.5.12 hence proves that all of these arcs are isotopic to each other, and in particular are isotopic to $s(D_6, D_4)$, as in Definition 2.5.14. It follows that, up to isotopy, all those arcs are contained in a small neighbourhood of $s(D_6, D_4)$, hence in a small neighbourhood of $l(D_6, D_4)$. In particular, we can without loss of

generality assume that every component of $D_6 \cap D_4^c$ is contained in Y and its intersection with γ_4 is contained in c_1 .

We have just shown that we have $D_6 \subseteq Y \cup Z$. Let h be the number of marked points in D_6 , and hence in D_4 (this is either k or n). Since D_6 is a h -punctured disc it follows that the region $Y \cup Z$ must contain at least h marked points, hence the region Z must contain at least $h - 1$ marked points. We claim the region Z contains exactly $h - 1$ marked points. In fact since $Z \subseteq D_4$ it contains at most h , but if it contained h then it would follow that $Z = D_4$ and hence D_1 and D_4 would be disjoint, which is impossible. Since $h \geq k > 2$ at most one component of $D_4 \cap D_1^c$ containing $(h - 1)$ punctures can exist. It follows that the region $M(D_6, D_4)$ is uniquely determined by D_1 and D_4 .

Let now $(D_1, D_2, D_3, D_4, D'_5, D'_6)$ be another alternating hexagons with the required properties. The previous construction proves that we have $M(D_6, D_4) = M(D'_6, D_4)$, since the isotopy type of that region only was only dependent on D_1 and D_4 . We can now apply Lemma 2.5.16 to prove that D_5^c is isotopic to $D_4 \cup M(D_6, D_4) = D_4 \cup M(D'_6, D_4)$, and that the right-hand side is isotopic to $(D'_5)^c$. Since we now have $D_5 = D'_5$ it follows from Lemma 2.5.3 that $D_6 = D'_6$, and the proof is complete. \square

Proposition 2.5.18, and Proposition 2.6.11 now conclude the proof of the recognisability of standard hexagons in the k -separating disc graph, that is, Proposition 2.5.2. Indeed, the first result is the proof of recognisability of standard hexagons from regular nonstandard ones, while the second is the proof of the same property for standard and irregular hexagons.

2.7 Disc Graphs and Curve Graphs

Before we start approaching the proof of rigidity for the graphs of k -separating discs, we will relate them and their automorphism groups to graph of curves. This will also provide us with the important base case for the inductive proof of Theorem *B* in the next section.

Practically speaking, our goal for this section will be to construct a group isomorphism

$$\text{Aut}(\mathcal{D}_k(\Sigma_m)) \longrightarrow \text{Aut}(\mathcal{C}_k(\Sigma_m)),$$

where $\mathcal{D}_k(\Sigma_m)$ (resp. $\mathcal{C}_k(\Sigma_m)$) is the k -separating disc (resp. curve) graph, as defined in Definition 1.4.1 (resp. Definition 1.3.6).

In order to be able to define such a homomorphism we will start by giving a description of k -separating curves, that is, vertices of $\mathcal{C}_k(\Sigma_m)$, in terms of combinatorial objects constructed from the graph $\mathcal{D}_k(\Sigma_m)$.

Definition 2.7.1. Let $k \geq 2$ and $m \geq 2k + 1$. Let

$$\mathcal{PC}_k(\Sigma_m) = \{\{D, D^c\} \text{ with } D \in \mathcal{D}_k(\Sigma_m)\}$$

be the set of *pairs of complementary discs* in $\mathcal{D}_k(\Sigma_m)$.

Any automorphism $g \in \text{Aut}(\mathcal{D}_k(\Sigma_m))$ preserves pairs of disjoint discs thanks to Proposition 2.2.4, hence acts on the set $\mathcal{PC}_k(\Sigma_m)$ by $g \cdot \{D, D^c\} = \{g(D), g(D^c)\}$.

We can now define the map

$$\mu : \mathcal{PC}_k(\Sigma_m) \longrightarrow \mathcal{V}(\mathcal{C}_k(\Sigma_m))$$

given by $\mu(\{D, D^c\}) = \partial D$. This map is well-defined, since ∂D is isotopic to ∂D^c . Moreover, if D is a h -punctured disc and D^c is a l -punctured disc the curve $\mu(\{D, D^c\})$ is a $(\min\{h, l\})$ -separating curve. Since $D, D^c \in \mathcal{V}(\mathcal{D}_k(\Sigma_m))$ we have that $h, l \leq k$, hence $\mu(\{D, D^c\}) \in \mathcal{C}_k(\Sigma_m)$.

We observe that the map μ is a bijection.

Lemma 2.7.2. Let $k \geq 2$ and $m \geq 2k + 1$. The previously defined map

$$\mu : \mathcal{PC}_k(\Sigma_m) \longrightarrow \mathcal{V}(\mathcal{C}_k(\Sigma_m))$$

is a bijection.

Proof. We will first prove injectivity. Let $\{A, A^c\}, \{B, B^c\}$ be such that we have $\mu(\{A, A^c\}) = \mu(\{B, B^c\})$. Up to switching A with A^c , if need be, this means that either $A \subseteq B \subseteq A$ or $A \subseteq B \subseteq A^c$. The latter case is clearly impossible, hence it follows that $A = B$, so $\{A, A^c\} = \{B, B^c\}$ and the map is injective.

To prove surjectivity let us consider a simple essential curve $\gamma \in \mathcal{V}(\mathcal{C}_k(\Sigma_m))$. Let D, D^c be the two complementary discs of γ . Since γ is h -separating with $h \geq k$ each of those discs must contain at least $h \geq k$ marked points, so the other one cannot contain more than $m - h \leq m - k$. It follows that $D, D^c \in \mathcal{V}(\mathcal{D}_k(\Sigma_m))$ and hence $\gamma = \mu(\{D, D^c\})$, proving that the map is surjective. \square

The inverse map

$$\mu^{-1} : \mathcal{V}(\mathcal{C}_k(\Sigma_m)) \longrightarrow \mathcal{PC}_k(\Sigma_m)$$

sends a curve γ to the pair of its complementary discs, that is the closures of the two components of $\Sigma_m \setminus \gamma$.

We now want to express edges in $\mathcal{C}(\Sigma_m)$, that is, disjointness between curves, in terms of properties of discs. In order to do so we give the following definition.

Definition 2.7.3. Let $k \geq 2$ and $m \geq 2k + 1$. We say two pairs of complementary discs $\{A, A^c\}, \{B, B^c\} \in \mathcal{PC}_k(\Sigma_m)$ are *nested* if either $A \subsetneq B$ or $A \subsetneq B^c$.

Expressing disjointness between curves in terms of nesting of the complementary discs is the content of the following result.

Lemma 2.7.4. Let $k \geq 2$ and $m \geq 2k + 1$. Let $\alpha, \beta \in \mathcal{C}_k(\Sigma_m)$. It holds that $d(\alpha, \beta) = 1$ if and only if $\mu^{-1}(\alpha)$ and $\mu^{-1}(\beta)$ are nested.

Proof. Let $\alpha, \beta \in \mathcal{C}_k(\Sigma_m)$ be such that $d(\alpha, \beta) = 1$, that is, they are two distinct disjoint curves. It follows that α would be contained in one of the complementary discs of β . It follows that one of the complementary discs of α must be contained in such a disc, as well. Moreover the inclusion is proper since α is not isotopic to β .

The converse follows the observation that if α and β intersect then each pair of complementary discs intersects as well, so it is impossible to have any inclusion between those. \square

We want to describe the k -separating curve graph $\mathcal{C}_k(\Sigma_m)$ in terms of combinatorial objects in $\mathcal{D}_k(\Sigma_m)$. In order to do so we will now define the *graph of pairs of complementary discs*, which is naturally isomorphic to the k -separating curve graph $\mathcal{C}_k(\Sigma_m)$.

Definition 2.7.5. Let $k \geq 2$ and $m \geq 2k + 1$. We define the *graph of pairs of complementary discs* $\mathcal{GPC}_k(\Sigma_m)$ as the abstract graph such that:

Vertices There is one vertex for every element of $\mathcal{PC}_k(\Sigma_m)$, that is, for every pair of complementary discs $\{D, D^c\}$ of $\mathcal{D}_k(\Sigma_m)$;

Edges There is an edge between two vertices corresponding to two pairs of complementary discs if they are nested.

Lemma 2.7.2 and Lemma 2.7.4 immediately prove the following.

Proposition 2.7.6. Let $k \geq 2$ and $m \geq 2k + 1$. The map

$$\mu : \mathcal{PC}_k(\Sigma_m) \longrightarrow \mathcal{V}(\mathcal{C}_k(\Sigma_m))$$

defined by $\mu(\{D, D^c\}) = \partial D$ extends to a $\text{Mod}^\pm(\Sigma_m)$ -equivariant graph isomorphism

$$\mu : \mathcal{GPC}_k(\Sigma_m) \xrightarrow{\sim} \mathcal{C}_k(\Sigma_m).$$

In particular we have $\text{Aut}(\mathcal{GPC}_k(\Sigma_m)) \cong \text{Aut}(\mathcal{C}_k(\Sigma_m))$.

We make the following simple remark about the relation between topological types of curves and the topological types of the pairs of discs they are identified with. Given a h -separating curve $\gamma \in \mathcal{C}_k(\Sigma_m)$, let $\gamma = \mu(\{D, D^c\})$. Then, up to exchanging D with D^c the disc D contains h marked points, and the disc D^c contains $m - k$.

Since automorphisms of $\mathcal{D}_k(\Sigma_m)$ preserve inclusion between discs, due to Proposition 2.2.1, they preserve nestedness, as well, hence induce automorphisms of the graph $\mathcal{GPC}_k(\Sigma_m)$. Such action induces a group homomorphism

$$\chi : \text{Aut}(\mathcal{D}_k(\Sigma_m)) \longrightarrow \text{Aut}(\mathcal{GPC}_k(\Sigma_m)).$$

We are now left to prove this is an isomorphism, which in turn proves that $\text{Aut}(\mathcal{D}_k(\Sigma_m))$ and $\text{Aut}(\mathcal{C}_k(\Sigma_m))$ are $\text{Mod}^\pm(\Sigma_m)$ -equivariantly isomorphic. that is, the content of the following.

Proposition 2.7.7. *Let $k \geq 2$ and $m \geq 2k + 1$. The map*

$$\chi : \text{Aut}(\mathcal{D}_k(\Sigma_m)) \longrightarrow \text{Aut}(\mathcal{GPC}_k(\Sigma_m))$$

is a $\text{Mod}^\pm(\Sigma_m)$ -equivariant isomorphism.

In the previous proposition surjectivity is the harder property to prove, and is the content of the following lemma.

Lemma 2.7.8. *Let $k \geq 2$ and $m \geq 2k + 1$. Let $f \in \text{Aut}(\mathcal{GPC}_k(\Sigma_m))$ be a graph automorphism. Then f is induced by an automorphism of $\mathcal{D}_k(\Sigma_m)$, that is, there exists a graph automorphism $\varphi \in \text{Aut}(\mathcal{D}_k(\Sigma_m))$ such that for every pair of complementary discs $\{D, D^c\} \in \mathcal{PC}_k(\Sigma_m)$ we have $f(\{D, D^c\}) = \{\varphi(D), \varphi(D^c)\}$.*

In order to prove the previous lemma, we first need to ensure that a map such as in the previous lemma preserves the topological types of the pairs of discs.

Thanks to the connection between the k -separating curve graph $\mathcal{C}_k(\Sigma_m)$ and the graph $\mathcal{GPC}_k(\Sigma_m)$, we will prove the equivalent property for the k -separating curve graph, since it will lead to much less cumbersome notation.

The key ingredient is to prove an analogous of Proposition 2.2.3 for the k -separating curve graphs, that is, topological types of curves are preserved under graph automorphisms. In order to do so we give a definition of maximal chains, similar to Definition 2.2.2, and prove an analogous of Proposition 2.2.3. This entire argument will closely follow the one in [B2, §5].

We begin with the following.

Definition 2.7.9. Let $\alpha, \beta, \gamma \in \mathcal{V}(\mathcal{C}_k(\Sigma_m))$ be three distinct curves. We say that γ *separates* α from β if every curve $\delta \in \mathcal{V}(\mathcal{C}_k(\Sigma_m))$ which is disjoint from γ is disjoint from either α or β .

Since it is defined only in terms of disjointness between curves, separation is preserved by graph automorphisms of $\mathcal{C}_k(\Sigma_m)$.

Topologically speaking, we have the following characterisation of separation.

Lemma 2.7.10. *Let $k \geq 2$ and $m \geq 2k + 1$. Let $\alpha, \beta, \gamma \in \mathcal{V}(\mathcal{C}_k(\Sigma_m))$ be three distinct curves. Let D, D^c be the two complementary discs of γ . Then γ separates α, β if and only if, up to switching the two discs, we have that $\alpha \subset D$ and $\beta \subset D^c$.*

Proof. The “if” implication is immediate.

For the converse we argue by contradiction. If, without loss of generality, it held $\alpha, \beta \subseteq D$ then there would exist a k -separating curve $\gamma \subseteq D$ intersecting both α and β , contradicting separation. \square

We can now define the concept of *maximal chains of curves*.

Definition 2.7.11. A *chain* of curves is a collection of curves $(\gamma_0, \dots, \gamma_j) \in \mathcal{C}_k(\Sigma_{k+n+1})$ such that for three indices $0 \leq h, i, l \leq j$ the curve γ_i separates γ_h from γ_l if and only if $h < i < l$.

A chain is said to be *maximal* if it is not a proper subchain of any other chain.

Since separation is recognisable in the k -separating curve graph then chains of curves are recognisable too.

Differently to what happens to discs we observe that, given a chain $(\gamma_0, \dots, \gamma_j)$ then the $(j + 1)$ -tuple $(\gamma_j, \dots, \gamma_0)$ is also a chain.

Given a curve in $\mathcal{C}_k(\Sigma_{k+n+1})$, there exists a maximal chain containing it. The following proposition, analogous to Proposition 2.2.3, gives us the characterisation of topological types of curves in terms of maximal chains.

Proposition 2.7.12. *Let $k \geq 2$ and $n \geq k$. Let $(\gamma_0, \dots, \gamma_j)$ be a maximal chain of curves in $\mathcal{C}_k(\Sigma_{k+n+1})$. Then $j = n - k + 1$ and γ_i is a $(\min\{(k + i), (n + 1 - i)\})$ -separating curve.*

Proof. Let D_0 be the complementary disc of γ_0 not containing any other curve in the chain. We know that such a disc exists thanks to Lemma 2.7.10. Let D_i be the complementary disc of γ_i containing D_{i-1} . These discs form a chain $D_0 \subsetneq \dots \subsetneq D_j$, as in Definition 2.2.2.

We claim the aforementioned chain is maximal. If not, suppose there exists a disc D and an index h such that $D_h \subsetneq D \subsetneq D_{h+1}$. Since both discs D_h and D_{h+1} contain at least k marked points, so do both D and D^c , hence the curve ∂D belongs to $\mathcal{C}_k(\Sigma_{k+n+1})$. It would follow that $(\gamma_0, \dots, \gamma_h, \partial D, \gamma_{h+1}, \dots, \gamma_j)$ would be a chain of curves strictly containing the original chain, hence contradicting maximality.

We can now apply Proposition 2.2.3 to deduce that the disc D_h contains $k + h$ marked points, hence D_h^c contains $n + 1 - h$. It follows that γ_i is a $(\min\{(k + i), (n + 1 - i)\})$ -separating curve. \square

Since maximal chains are preserved by graph automorphisms we can now deduce how to recognise topological types by automorphisms of the k -separating curve graph.

Corollary 2.7.13. *Let $k \geq 2$ and $m \geq 2k + 1$. Let h be such that $k \leq h \leq m - k$. Let γ be a vertex of $\mathcal{C}_k(\Sigma_m)$. Then γ is a h -punctured disc if and only if there exists a maximal chain $(\gamma_0, \dots, \gamma_{m-2k})$ such that $\gamma = \gamma_{h-k}$. In particular, the topological type of a curve is recognisable in the graph.*

We can now state an analogous property of preservation of topological types for $\mathcal{GPC}_k(\Sigma_m)$.

Corollary 2.7.14. *Let $k \geq 2$ and $m \geq 2k + 1$. Let $\{D, D^c\} \in \mathcal{PC}_k(\Sigma_m)$. Then one of the discs contains h -punctures and the other contains $m - h$ if and only if the curve $\mu(\{D, D^c\}) \in \mathcal{V}(\mathcal{C}_k(\Sigma_{k+n+1}))$ is $(\min\{h, m - h\})$ -separating. In particular the topological types of a pair of complementary discs is preserved by automorphisms of $\mathcal{GPC}_k(\Sigma_m)$.*

Proof. Since $\mu(\{D, D^c\}) = \partial D$ the first part comes directly from Corollary 2.7.13.

For the second part let $f \in \text{Aut}(\mathcal{GPC}_k(\Sigma_m))$. Then $\mu \circ f \circ \mu^{-1} \in \text{Aut}(\mathcal{C}_k(\Sigma_m))$. It follows that if $\mu(\{D, D^c\})$ is h -separating then also

$$\mu \circ f \circ \mu^{-1}(\mu(\{D, D^c\})) = \mu(f(\{D, D^c\}))$$

is, so the discs $f(\{D, D^c\})$ has the same topological type of those in $\{D, D^c\}$. \square

We are now ready to prove that automorphisms of the graph $\mathcal{GPC}_k(\Sigma_m)$ induce automorphisms of the graph of discs.

Proof of Lemma 2.7.8. Our goal is to define a map $\varphi \in \text{Aut}(\mathcal{D}_k(\Sigma_m))$ such that for every $\{D, D^c\} \in \mathcal{PC}_k(\Sigma_m)$ we have $f(\{D, D^c\}) = \{\varphi(D), \varphi(D^c)\}$. Given a pair of complementary discs $\{D, D^c\} \in \mathcal{PC}_k(\Sigma_m)$ such that they are not $\frac{m}{2}$ -punctured we can define $\varphi(D)$ to be the unique disc in $f(\{D, D^c\}) \in \mathcal{PC}_k(\Sigma_m)$ of the same topological type of D , which exists since f preserves topological types, thanks to Corollary 2.7.14.

Given a pair of complementary $\frac{m}{2}$ -punctured discs $\{A, A^c\}$ there exists a pair of complementary discs $\{B, B^c\}$ with B a h -punctured disc and $k \leq h < \frac{m}{2}$ such that $B \subsetneq A$, that is, such that $\{A, A^c\}$ and $\{B, B^c\}$ are nested. We define $\varphi(A)$ to be the unique disc in $f(\{A, A^c\})$ containing $\varphi(B)$, which exists since f preserves nesting. We will now prove that this is a good definition. Moreover, we observe that, by definition, we have $\varphi(D^c) = \varphi(D)^c$.

We now claim that if B, B' are a h -punctured and a h' -punctured discs respectively, with $h, h' < \frac{m}{2}$, and such that $B, B' \subsetneq A$, such that $\{B, B^c\}$ and $\{B', (B')^c\}$ are not nested, then $\varphi(B) \subsetneq \varphi(A)$ if and only if $\varphi(B') \subsetneq \varphi(A)$. If not, without loss of generality we can assume we have $\varphi(B) \subseteq \varphi(A)$, while $\varphi(B') \subseteq \varphi(A)^c$. It follows that $\varphi(B') \subsetneq \varphi(B)^c$, hence the pairs $\{\varphi(B), \varphi(B)^c\}$ and $\{\varphi(B'), \varphi(B')^c\}$ are nested. Since f^{-1} preserves nesting, this is a contradiction to the fact that $\{B, B^c\}$ and $\{B', (B')^c\}$ were not nested. The claim is proven.

Let now A be a $\frac{m}{2}$ -punctured disc, and let $B, B' \subsetneq A$ be two discs. Then there exists a disc $C \subsetneq A$ such that the pair $\{C, C^c\}$ is not nested with either $\{B, B^c\}$ or $\{B', (B')^c\}$. It follows from the previous claim that $\varphi(B) \subsetneq \varphi(A)$ if and only if $\varphi(C) \subsetneq \varphi(A)$, which happens if and only if $\varphi(B') \subsetneq \varphi(A)$. From this, it follows that the map φ we have previously define is well-defined for $\frac{m}{2}$ -punctured discs, too.

We have a permutation of vertices of $\mathcal{D}_k(\Sigma_m)$ preserving topological types and inducing the map f . We are now left to prove that this map φ is edge-preserving or, equivalently, it preserves inclusion between discs.

First, let $\{A, A^c\}, \{B, B^c\}$ be two nested pairs of complementary discs such that A has the same topological type of B . Then the discs contain a number of marked

points which is different from $\frac{m}{2}$, otherwise we would have either $A = B$ or $A = B^c$, contradicting nesting. For the same reason, it follows that $A \subset B^c$. Moreover, the disc $\varphi(A)$ contains the same number of marked points as A by definition of φ , and hence the same number of marked points of B and $\varphi(B)$. It follows that if it held $\varphi(A) \subseteq \varphi(B)$ we would have equality, contradicting nestedness, hence $\varphi(A) \subsetneq \varphi(B)^c$.

We will now prove that φ preserves inclusion between every pair of discs, independently of their topological type. Let us argue by contradiction and suppose there exist two disjoint discs $A, B \in \mathcal{V}(\mathcal{D}_k(\Sigma_m))$ such that $A \subsetneq B^c$ but $\varphi(A) \subsetneq \varphi(B)$. If A contains $\frac{m}{2}$ marked points, then B does not, hence $B \subsetneq A^c$ and the inclusion $\varphi(B) \subseteq \varphi(A^c) = \varphi(A)^c$ follows from the previous discussion about the good definition of $\varphi(A)$. Then, up to considering A^c instead of A , if need be, we can assume A to be h -punctured with $h < \frac{m}{2}$. Up to considering φ^{-1} instead of φ we can assume B to be the disc of minimum complexity among $\{B, B^c\}$, and it to be l -punctured with $l < \frac{m}{2}$. Indeed, the case $l = \frac{m}{2}$ follows once again from the discussion about the good definition of $\varphi(B)$. Moreover, the case $h = l < \frac{m}{2}$ has been proven in the previous paragraph, hence we can assume that $h < l$. Since B^c contains $m - l > l$ marked points there exists a l -punctured disc $C \subsetneq B^c$ such that the pairs $\{A, A^c\}$ and $\{C, C^c\}$ are not nested. Since $\{B, B^c\}$ and $\{C, C^c\}$ are nested it follows from previous arguments that $\varphi(C) \subsetneq \varphi(B)^c$, that is, $\varphi(B) \subsetneq \varphi(C)^c$. If we had $\varphi(A) \subsetneq \varphi(B)$ then it would hold that $\varphi(A) \subsetneq \varphi(C)^c$, hence $\{\varphi(A), \varphi(A^c)\}$ and $\{\varphi(C), \varphi(C^c)\}$ would be nested and, since f^{-1} preserves nestedness, it follows that $\{A, A^c\}$ and $\{C, C^c\}$ would be nested, which is a contradiction. It follows that $\varphi(A) \subsetneq \varphi(B^c)$, hence $\varphi(A)$ is disjoint from $\varphi(B)$.

This proves that φ is a well-defined graph automorphism of $\mathcal{D}_k(\Sigma_m)$ inducing the map f . \square

We are now ready to prove bijectivity of the previously defined map

$$\chi : \text{Aut}(\mathcal{D}_k(\Sigma_m)) \longrightarrow \text{Aut}(\mathcal{GPC}_k(\Sigma_m)).$$

Proof of Proposition 2.7.7. We will start proving injectivity. Let $\varphi \in \text{Aut}(\mathcal{D}_k(\Sigma_m))$ be a graph automorphism such that $\chi(\varphi) = \text{Id}$. We claim that $\varphi = \text{Id}_{\mathcal{D}_k(\Sigma_m)}$. Let $\{D, D^c\}$ be a pair of complementary disc which are not $\frac{m}{2}$ -punctured, hence they have two different topological types. Since $m \geq 2k + 1$ such a pair of complementary discs exists. It holds $\{D, D^c\} = \chi(\varphi)(\{D, D^c\}) = \{\varphi(D), \varphi(D^c)\}$. Since topological types of discs are preserved by graph automorphisms, due to Proposition 2.2.3 it follows that $\varphi(D) = D$ and $\varphi(D^c) = D^c$. So, φ acts as the identity on every disc which is not $\frac{m}{2}$ -punctured.

Let now $\{A, A^c\}$ be a pair of complementary $\frac{m}{2}$ -punctured discs. Let us argue by contradiction and assume $\varphi(A) = A^c$. In particular there exists a h -punctured disc $B \subsetneq A$ with $k \leq h < \frac{m}{2}$. From the argument in the previous paragraph it follows $\varphi(B) = B$. Since inclusion between discs is preserved by graph automorphism due to Proposition 2.2.1 it follows that $B \subseteq A^c$, which is absurd since $B \subseteq A$. We have reached a contradiction and proven that $\varphi = \text{Id}_{\mathcal{D}_k(\Sigma_m)}$, hence the homomorphism χ is injective.

Surjectivity was proven in Lemma 2.7.8. \square

2.8 Rigidity of Graphs of Discs

Now that we have proven the core technical results of the chapter, we can proceed to the proof of first the two main theorems. The core of the argument is straightforward: starting with the graph \mathcal{D}_k we will use surrounding pairs and triples, which we can recognise thanks to Corollary 2.2.10, to reconstruct the graph \mathcal{D}_{k-1} . We will use that construction to prove that there exists an injection

$$\text{Aut}(\mathcal{D}_k) \hookrightarrow \text{Aut}(\mathcal{D}_{k-1})$$

between the automorphism groups. At that point, an inductive argument will let us prove there exists an injection of the automorphism group into the group of automorphisms of the graph \mathcal{D}_2 . This is the reason why, in Corollary 2.2.10, the hypothesis $k \geq 3$ sufficed our needs. From such map we will construct an injection into the automorphism group of the regular curve graph, allowing us to apply Ivanov's Theorem to prove rigidity. This approach closely mimics the ones used by Bowditch in [B2, §4] and by McLeay in [Mc1].

For the definition of surrounding pairs and triples we refer to Definition 2.1.1 and Definition 2.1.2. We begin with the following natural definition.

Definition 2.8.1. Let $k \geq 3$ and $m \geq 2k + 1$. Let O be a (isotopy class of) $(k - 1)$ -punctured disc in Σ_m . We say that a surrounding pair $\{A, B\}$ surrounds O if $O \subseteq A \cap B$ up to isotopy. We say that a surrounding triple $\{A, B, C\}$ surrounds O if $O \subseteq A \cap B \cap C$ up to isotopy.

We introduce the following notation.

Definition 2.8.2. Let $k \geq 3$ and $m \geq 2k + 1$. Let $\mathcal{H}(\Sigma_m)$ be the set of surrounding pairs in $\mathcal{D}_k(\Sigma_m)$. Let O be a (isotopy class of) $(k - 1)$ -punctured disc in Σ_m . We denote by $\mathcal{H}_O(\Sigma_m)$ the set of surrounding pairs surrounding O .

We have already mentioned that we will use surrounding pairs in the graph \mathcal{D}_k to represent $(k - 1)$ -punctured disc. However, different pairs can represent the same curve, which happens for example when three curves form a surrounding triple. Hence we need to mod out by an appropriate equivalence relation, which is defined in the following

Definition 2.8.3. Let $k \geq 3$ and $m \geq 2k + 1$. We denote by \sim the minimal equivalence relation on the set of surrounding pairs $\mathcal{H}(\Sigma_m)$ generated by $\{A, B\} \sim \{B, C\}$ if $\{A, B, C\}$ is a surrounding triple.

Since a surrounding pair only surrounds a unique isotopy class of $(k - 1)$ -punctured discs, there exists a natural bijection between the isotopy classes of $(k - 1)$ -punctured discs and the family of pairs in $\mathcal{H}_O(\Sigma_m)$. In order to be able to use this identification for our future purposes, we need to describe these sets in terms of the combinatorial structure of the curve graphs only. This is not obvious *a priori*, and it is the content of the following proposition.

Proposition 2.8.4. Let $k \geq 3$ and $k \geq k$. Let O be a $(k - 1)$ -punctured disc in Σ_{k+n+1} . The set $\mathcal{H}_O(\Sigma_{k+n+1})$ is an equivalence class of $\mathcal{H}(\Sigma_{k+n+1})$ under the relation \sim .

The previous result could be expressed in terms of connectedness of graphs, in a way which would be analogous to [B2, Lemma 4.2] and [Mc1, Lemma 3.3]. In this case, a nice and quick proof, although non constructive could follow from an application of [Pu, Lemma 2.1] (that is, Lemma 2.3.6), a lemma by Andrew Putman. We will however not follow such an approach, which requires knowledge about generating sets of the mapping class group, and give a more elementary and constructive proof.

Before we prove the previous proposition we will follow [B2, Lemma 4.2] and introduce a nice new description of the k -punctured discs containing a $(k - 1)$ -punctured disc O , which will really simplify our arguments. We start by giving the following definition.

Definition 2.8.5. Let O be a $(k - 1)$ -punctured disc in Σ_m . An O -arc will be an isotopy class of arcs in O^c with one endpoint on ∂O and the other on a marked point, such that its interior does not intersect either ∂O or any of the marked points. The isotopies considered have to keep the first endpoint on ∂O at any time and the second endpoint fixed, and at any time the interior of the arc cannot intersect either ∂O or any of the marked points.

Given a parametrised arc a representing an O -arc let $\mathcal{N}(O \cup a)$ be a sufficiently small regular neighbourhood of $O \cup a$, that is, such that it does not contain any marked point not in $O \cup a$. This is a k -punctured disc containing O . Moreover, if a' is a parametrised arc isotopic to a then the two neighbourhoods $\mathcal{N}(O \cup a)$ and $\mathcal{N}(O \cup a')$ are isotopic. This means that, given an (isotopy class of) O -arc $[a]$ its *regular neighbourhood* $N_O([a]) = \mathcal{N}(O \cup \bar{a})$ is well defined as an isotopy class of k -punctured discs. We have hence defined a map

$$N_O : \{O\text{-arcs}\} \longrightarrow \{k\text{-punctured discs containing } O\}/\text{isotopy}.$$

The map we have just defined is actually a bijection, as shown in the following.

Lemma 2.8.6. *Let $k \geq 3$ and $m \geq 2k + 1$. Let O be a $(k - 1)$ -punctured disc in Σ_m . The map N_O , defined as above, is a bijection.*

Proof. Let A be a k -punctured disc containing O . Up to isotopy, there exists a unique arc in the pair of pants (to be precise a surface of type S_1^2) $A \cap O^c$ with one endpoint on ∂O and the marked point as the other one. Clearly $\mathcal{N}(O \cup a)$ is isotopic to the disc A , hence the map N_O is surjective.

Let now $a \neq b$ be two distinct O -arcs such that $N_O(a) = N_O(b)$, that is, such that, up to isotopy, $\mathcal{N}(O \cup a) = \mathcal{N}(O \cup b)$. In particular, up to isotopy, we have that $b \subset \mathcal{N}(O \cup a)$, hence it is an arc in the pair of pants $\mathcal{N}(O \cup a) \cap O^c$ from ∂O to the only marked point. We have already noticed that such an arc is unique up to isotopy, hence $a = b$ and the map N_O is injective, as well. \square

Given O -arcs a, b, c they can always be realised in a way such that they are in minimal and general position and their endpoints on ∂O are pairwise distinct, hence from now on we will always assume so, unless otherwise stated.

We are now interested in characterising surrounding pairs and triple in terms of O -arcs, which is the content of the following.

Lemma 2.8.7. *Let $k \geq 3$ and $m \geq 2k + 1$. Let O be a $(k - 1)$ -punctured disc in Σ_m . Let a, b, c be O -arcs, let $A = N_O(a)$, $B = N_O(b)$ and $C = N_O(c)$. The pair $\{A, B\}$ is a surrounding pair surrounding O if and only if the arcs a and b are disjoint. Moreover, the triple $\{A, B, C\}$ is a surrounding triple surrounding O if and only if the arcs a, b, c are pairwise disjoint.*

Proof. If the three arcs a, b, c are disjoint the boundaries of the three discs A, B, C only intersect twice. Moreover, both the triple intersection and any pairwise intersection contains exactly the $k - 1$ marked points in O . It follows from Lemma 2.5.6 that any pair of such arcs form a surrounding pair, and the three of them constitute a surrounding triple.

Conversely, if $\{A, B\}$ is a surrounding pair we can assume, up to isotopy, that we have $A \cap B = O$. It follows that a (resp. b), which is the unique arc from the boundary to the marked point in the once-punctured discs $A \cap B^c$ (resp. $A^c \cap B$) is disjoint from B (resp. A): in particular the arcs a and b are disjoint. The argument for surrounding triples is completely analogous. \square

Now that we have introduced the definition of O -arcs, and hence simplified our notation, we can move to the proof of Proposition 2.8.4.

Proof of Proposition 2.8.4. First, we claim that the set of pairs $\mathcal{H}_O(\Sigma_{k+n+1})$ is closed under the relation \sim , so it is entirely contained in a unique equivalence class. Indeed, let $\{A, B\} \in \mathcal{H}_O(\Sigma_{k+n+1})$ and $\{B, C\}$ be a surrounding pair such that $\{A, B\} \sim \{B, C\}$. Up to isotopy, we have that

$$B \cap C = A \cap B \cap C = A \cap B = O,$$

that is, both pairs $\{A, B\}$ and $\{B, C\}$ surround O , where the first two equalities follows from the fact that the triple $\{A, B, C\}$ is a surrounding triple. It follows that we have $\{B, C\} \in \mathcal{H}_O(\Sigma_{k+n+1})$. Since the equivalence \sim is generated by the above relation then it follows that for every pair $\{D, E\}$ such that $\{A, B\} \sim \{D, E\}$ it then holds $\{D, E\} \in \mathcal{H}_O(\Sigma_{k+n+1})$. Our claim is proven.

During the remain of the proof we will abuse notation by omitting the map N_O and directly identifying k -punctured discs containing O with O -arcs. In particular, due to Lemma 2.8.7, an element in $\mathcal{H}_O(\Sigma_{k+n+1})$ will be a pair $\{a, b\}$ of distinct disjoint O -arcs.

We now left to prove that, given two surrounding pairs $\{a, b\}, \{c, d\} \in \mathcal{H}_O$ then $\{a, b\} \sim \{c, d\}$. We recall that we will always assume all of our O -arcs to be in minimal and general position with each other and to have different endpoints on O . Moreover we will write $\omega = \partial O$.

First, we claim that if u, v are two O -arcs there exists an O -arc w which intersects u and v at most once, and when it does the intersection point is the endpoint of the O -arcs which is not on ω . The following constructions (in the case where u and v share an endpoint) are pictured in Figure 2.22. First, let x_0 be the endpoint of u on ω . For i even, let x_{i+1} be the intersection point of $u \cap v$ such that x_{i+1} belongs to the component of $u \setminus \{x_i\}$ disjoint from ω , and the interior of the subarc of u between x_i and x_{i+1} does not intersect v . If such a point does not exists let x_{i+1} be equal to the second endpoint

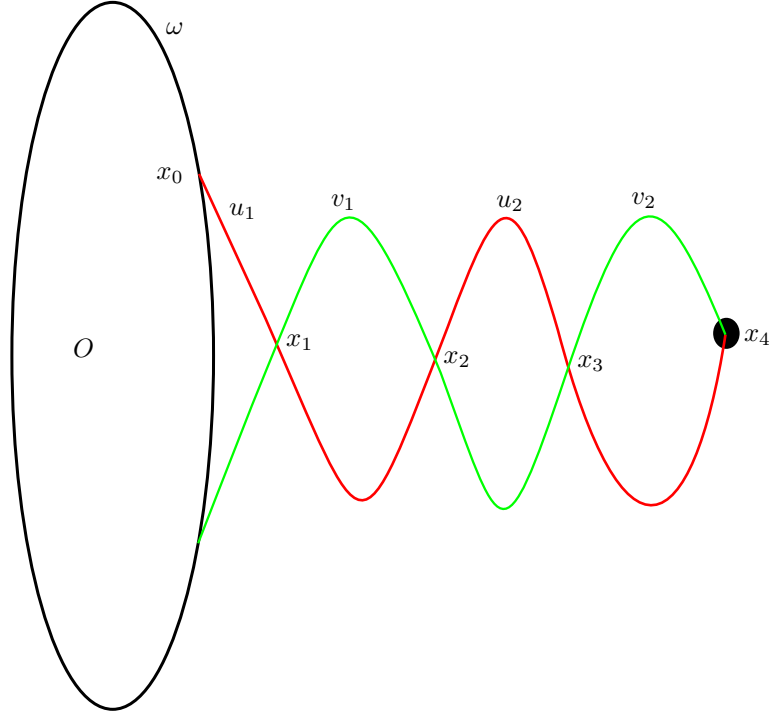


Figure 2.22: The red arc is u , while the green one is v .

of u and let $N = i + 1$. Let u_h be the subarc of u between x_{2h} and x_{2h+1} . For i even, if x_i is not the second endpoint of u , let x_{i+1} be the intersection point of $u \cap v$ such that x_{i+1} belongs to the component of $v \setminus \{x_i\}$ disjoint from ω the interior of the subarc of v between x_i and x_{i+1} does not intersect u . If such a point does not exist let x_{i+1} be equal to the second endpoint of v and let $N = i + 1$. Let v_h be the subarc of v between x_{2h+1} and x_{2h+2} . The arc

$$w = \bigcup_{i=0}^N u_i \cup v_i$$

can be isotoped to only intersect $u \cup v$ at most in its endpoint not on ω , so the claim is proven.

Second, we claim that if u, v are two O -arcs intersecting only once, then there exists an O -arc w which is disjoint from both. Indeed, the space $O^c \setminus (u \cup v)$ only has two connected components, hence both component intersects ω , and the space contains at least $n \geq 1$ marked points, hence at least one component contains at least a marked points. It follows that such component contains a nontrivial O -arc, which is disjoint from both u and v . The claim is proven.

We claim that it is sufficient to prove that any two pairs in $\mathcal{H}_O(\Sigma_{k+n+1})$ of the form $\{u, w\}$, $\{v, w\}$ are \sim -equivalent. Indeed, let $\{a, b\}, \{c, d\} \in \mathcal{H}_O$. Let e be an O -arc intersecting a and c at most once, which exists by the first claim. Let f (resp. g) be an O -arc disjoint from both e and a (resp. c), which exists due to a previous claim. If our last claim is true than the following holds:

$$\{a, b\} \sim \{a, f\} \sim \{f, e\} \sim \{e, g\} \sim \{e, c\} \sim \{c, d\}$$

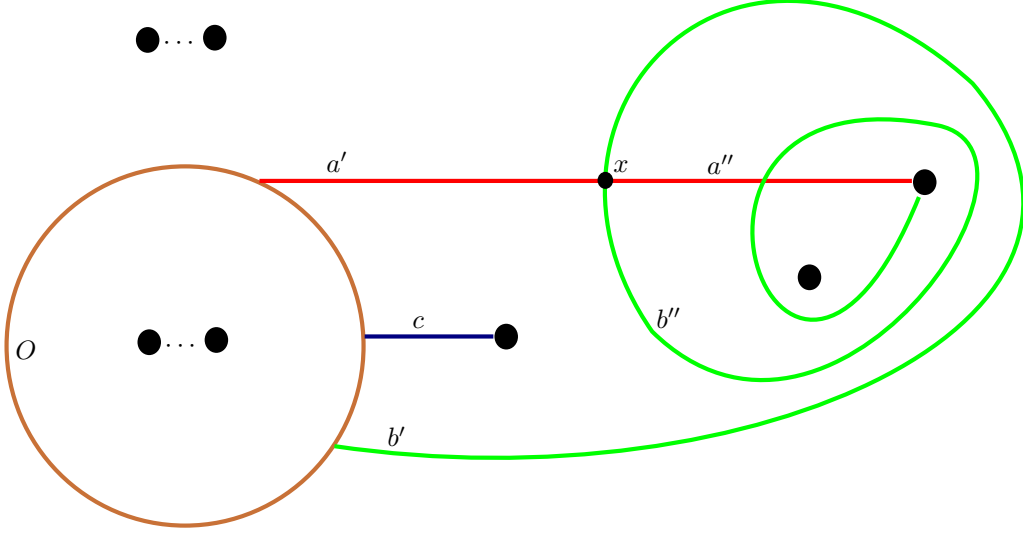


Figure 2.23: One of the possible cases of the construction (note that the arcs a, b do not share an endpoint in general).

hence $\{a, b\} \sim \{c, d\}$ due to the transitivity of \sim .

We now move to the proof of the last claim. Let $\{a, c\}, \{b, c\} \in \mathcal{H}_O(\Sigma_{k+n+1})$. We will argue by induction on the intersection number $i(a, b)$. If that number is zero the triple $\{a, b, c\}$ is a surrounding triple, so $\{a, c\} \sim \{a, b\}$ by definition.

If $i(a, b) = 1$ the subspace $O^c \setminus (a \cup b \cup c)$ has exactly two connected components, since c is disjoint from $a \cup b$, hence both intersect ω . Moreover, such space contains at least $k + n + 1 - (k - 1) - 3 = n - 1 \geq 1$ marked points, hence it contains a nontrivial O -arc d . The arc d is disjoint from all the three arcs a, b, c . It follows that

$$\{a, c\} \sim \{a, d\} \sim \{d, b\} \sim \{b, c\}$$

hence $\{a, c\} \sim \{b, c\}$.

We will now assume $i(a, b) > 1$: a picture of the following construction is provided in Figure 2.23. Let $x \in a \cap b$ be the intersection point which is closest to ω , that is, such that the interior of the arc of a between ω and x does not intersect b . Since $i(a, b) > 1$ the point x is not marked. Let a' (resp. b') be defined the arcs of a (resp. b) between ω and x . Let $a'' = a \setminus a'$. We can now consider the arc $\tilde{a} = b' \cup a''$. This arc is still disjoint from c and such that $i(\tilde{a}, b) = i(a, b) - 1$ and $i(\tilde{a}, a) = 1$. We can now use the previous case and the inductive hypothesis to prove that $\{a, c\} \sim \{\tilde{a}, c\} \sim \{b, c\}$, and complete the proof. \square

From Proposition 2.8.4 we can immediately deduce the following.

Corollary 2.8.8. *Let $k \geq 3$ and $m \geq 2k + 1$. The map $O \mapsto \mathcal{H}_O(\Sigma_m)$, is a bijective correspondence between the collection of sets $\mathcal{H}_O(\Sigma_m)$, O varying among $(k-1)$ -punctured discs on Σ_m , and the quotient set \mathcal{H}/\sim .*

Moreover, given a k -punctured disc A , we have $O \subseteq A$ if and only if A is an element of a surrounding pair in $\mathcal{H}_O(\Sigma_m)$.

We have now proven that we can naturally describe $(k-1)$ -punctured discs, alongside their inclusion inside k -punctured discs, just in terms of objects deriving from the graph $\mathcal{D}_k(\Sigma_m)$. Indeed, $(k-1)$ -punctured discs correspond to equivalence classes of surrounding pairs, which we can recognise in the the graph $\mathcal{D}_k(\Sigma_m)$, mod out by a relation defined uniquely in terms of belonging to surrounding triples, which is another property that can be detected in the graph of discs. In order to be able to reconstruct the entire graph $\mathcal{D}_{k-1}(\Sigma_m)$, hence proceed with the induction, we still need to be able to recognise disjointness between a $(k-1)$ -punctured disc and another disc. This is the content of the following lemma.

Lemma 2.8.9. *Let $k \geq 3$ and $n \geq k$. Let O, O' be two $(k-1)$ -punctured discs in Σ_{k+n+1} . Let A be a vertex of $\mathcal{D}_k(\Sigma_{k+n+1})$. Then the discs O and O' are disjoint if and only if there exist two k -punctured discs C, D with $d(C, D) = 1$ such that $O \subseteq C$ and $O' \subseteq D$. The discs O and A are disjoint if and only if there exists a k -punctured disc B such that $d(A, B) = 1$ and $O \subseteq B$.*

Proof. We will first prove the first part on the lemma. Let O, O' be two disjoint $(k-1)$ -punctured discs in Σ_{k+n+1} . The region $X = O^c \cap O'^c$ contains

$$k + n + 1 - 2(k - 2) = (n - k) + 3 \geq 3$$

marked points. It follows that there exist an O -arc a and an O' -arc b in X which are disjoint and such that $a \cap O' = b \cap O = \emptyset$. These arcs represent two disjoint k -punctured discs C, D such that $O \subseteq C$ and $O' \subseteq D$. The converse easily follows from the fact that $O \cap O' \subseteq C \cap D = \emptyset$.

For the second statement let A be a h -punctured disc disjoint from O . The region $Y = A^c \cap O^c$ contains

$$k + n + 1 - (k - 1) - h = n + 2 - h \geq 1$$

marked points, since $h \leq n + 1$. It follows that there exists an O -arc b in Y , and $B = N_O(b)$ is a k -punctured disc containing O which is disjoint from A . The converse immediately follows from the fact that $O \cap A \subseteq B \cap A = \emptyset$. \square

We are now ready to describe a graph, constructed by adding some new vertices and edges to graph $\mathcal{D}_k(\Sigma_m)$, and we will later prove that this new graph is isomorphic to the graph $\mathcal{D}_{k-1}(\Sigma_m)$.

Definition 2.8.10. Let $k \geq 3$ and $m \geq 2k + 1$. We define the *augmentation* of the graph of k -separating discs, which we will denote with $\mathcal{D}_k^+(\Sigma_m)$ as the abstract graph defined by

Vertices There is a vertex for each vertex of $\mathcal{D}_k(\Sigma_m)$ and for every element of $\mathcal{H}/\sim \times \{0, 1\}$;

Edges Edges are defined as follows:

- Between two vertices of \mathcal{D}_k if they are adjacent in that graph;

- Between a disc A and an equivalence class $\mathcal{H}_O(\Sigma_m) \times \{0\}$ if there exists a disc B adjacent to A in $\mathcal{D}_k(\Sigma_m)$ and appearing in a surrounding pair in $\mathcal{H}_O(\Sigma_m)$;
- Between two equivalence classes $\mathcal{H}_O(\Sigma_m) \times \{0\}$ and $\mathcal{H}_{O'}(\Sigma_m) \times \{0\}$ if there exist discs C and D appearing in a surrounding pair in $\mathcal{H}_O(\Sigma_m)$ and $\mathcal{H}_{O'}(\Sigma_m)$, respectively, and adjacent in $\mathcal{D}_k(\Sigma_m)$;
- Between two equivalence classes $\mathcal{H}_O(\Sigma_m) \times \{0\}$ and $\mathcal{H}_O(\Sigma_m) \times \{1\}$.

We recall that the equivalence classes on \mathcal{H}/\sim represent $(k-1)$ -separating curves, hence for each there are two discs in $\mathcal{D}_{k-1}(\Sigma_m)$ having it as boundary curve. In the previous definition a vertices of type $\mathcal{H}_O(\Sigma_m) \times \{0\}$ represent the $(k-1)$ - punctured discs bounded by $\mathcal{H}_O(\Sigma_m)$, while the vertex $\mathcal{H}_O(\Sigma_m) \times \{1\}$ stands for its complementary $(m-k+1)$ -separating disc. The latter is adjacent to only one disc in $\mathcal{D}_{k-1}(\Sigma_m)$, it being its complementary disc.

It is now possible to define a map

$$f : \mathcal{D}_{k-1}(\Sigma_m) \longrightarrow \mathcal{D}_k^+(\Sigma_m)$$

defined by

$$f(D) = \begin{cases} D & \text{if } D \in \mathcal{D}_k(\Sigma_m) \\ \mathcal{H}_D \times \{0\} & \text{if } D \text{ is a } (k-1)\text{-punctured disc} \\ \mathcal{H}_D \times \{1\} & \text{if } D \text{ is a } (m-k+1)\text{-punctured disc.} \end{cases}$$

Moreover, there are two natural embeddings

$$\mathcal{D}_k(\Sigma_m) \hookrightarrow \mathcal{D}_{k-1}(\Sigma_m)$$

and

$$\mathcal{D}_k(\Sigma_m) \hookrightarrow \mathcal{D}_k^+(\Sigma_m).$$

Both of these embeddings commute with the map f .

The following lemma holds.

Lemma 2.8.11. *Let $k \geq 3$ and $m \geq 2k+1$. The map*

$$f : \mathcal{D}_{k-1}(\Sigma_m) \longrightarrow \mathcal{D}_k^+(\Sigma_m)$$

is a $\text{Mod}^\pm(\Sigma_m)$ -equivariant graph isomorphism.

Proof. Corollary 2.8.8 implies that the map f induces a bijection on vertices. Moreover, Lemma 2.8.9, alongside the definition of edges in $\mathcal{D}_k^+(\Sigma_m)$, immediately implies that the map is a graph homomorphism and that the edges are in bijective correspondance, too. \square

From the previous result it is possible to derive an inductive proof of the connectedness of the k -separating disc graph, that is, Theorem 2.3.1. We will discuss this line of proof, alongside a detailed discussion about connectedness results, in Section 2.3.

Corollary 2.2.10 ensures that automorphisms of $\mathcal{D}_k(\Sigma_m)$ induce automorphism of the augmented graph, since they preserve surrounding pairs and triples, hence equivalence classes of \mathcal{H} . This induces the group homomorphism

$$\varphi : \text{Aut}(\mathcal{D}_k(\Sigma_m)) \longrightarrow \text{Aut}(\mathcal{D}_k^+(\Sigma_m)).$$

This homomorphism is defined as follows. Given an automorphism $g \in \text{Aut}(\mathcal{D}_k(\Sigma_m))$ and a vertex $D \in \mathcal{D}_k^+(\Sigma_m)$ which corresponds to a disc, that is, a vertex of $\mathcal{D}_k(\Sigma_m)$, we simply define $\varphi(g)(D) = g(D)$. For one of the other vertices, say $\mathcal{H}_O \in \mathcal{H}/\sim$ we define $\varphi(g)(\mathcal{H}_O) = g \cdot \mathcal{H}_O$, where the action on a surrounding pair $\{A, B\}$ is given by $g \cdot \{A, B\} = \{g(A), g(B)\}$. The homomorphism φ is $\text{Mod}^\pm(\Sigma_m)$ -equivariant.

We can now conjugate φ with f to define a second $\text{Mod}^\pm(\Sigma_m)$ -equivariant group homomorphism

$$\psi : \text{Aut}(\mathcal{D}_k(\Sigma_m)) \longrightarrow \text{Aut}(\mathcal{D}_{k-1}(\Sigma_m))$$

given by

$$\psi(g) = f^{-1} \circ \varphi(g) \circ f.$$

We can now prove the key reduction step for our argument.

Lemma 2.8.12. *Let $k \geq 3$ and $m \geq 2k + 1$. The group homomorphism*

$$\varphi : \text{Aut}(\mathcal{D}_k(\Sigma_m)) \longrightarrow \text{Aut}(\mathcal{D}_{k-1}(\Sigma_m))$$

is injective.

Proof. Let $g \in \text{Aut}(\mathcal{D}_k(\Sigma_m))$ be such that $\psi(g) = \text{Id} \in \text{Aut}(\mathcal{D}_{k-1}(\Sigma_m))$. By conjugation with f^{-1} this implies that $\varphi(g) = \text{Id} \in \text{Aut}(\mathcal{D}_k^+(\Sigma_m))$. In particular, for every vertex v of $\mathcal{D}_k^+(\Sigma_m)$ representing a h -punctured disc, with $k \leq h \leq n + 1$, we have $v = \varphi(g)(v) = g(v)$. It follows that $g = \text{Id} \in \text{Aut}(\mathcal{D}_k(\Sigma_m))$, and the homomorphism is injective. \square

A repeated application of the previous lemma proves the following.

Corollary 2.8.13. *Let $k \geq 2$ and $m \geq 2k + 1$. There exists an injective $\text{Mod}^\pm(\Sigma_m)$ -equivariant group homomorphism*

$$\text{Aut}(\mathcal{D}_k(\Sigma_m)) \hookrightarrow \text{Aut}(\mathcal{D}_2(\Sigma_m)).$$

Thanks to the link with automorphisms of the curve graph introduced in the previous section we are now ready to prove Theorem B.

Theorem B. *Let $k \geq 2$ and $m \geq 2k + 1$. Then the k -separating disc graph is rigid, that is, the group homomorphism*

$$\rho : \text{Mod}^\pm(\Sigma_m) \longrightarrow \text{Aut}(\mathcal{D}_k(\Sigma_m))$$

induced by the natural action is an isomorphism.

Proof. Thanks to Corollary 2.8.13 and Proposition 2.7.7, applied with $k = 2$, alongside the observation that $\mathcal{C}_2(\Sigma_m) = \mathcal{C}(\Sigma_m)$, we have the following commutative diagram:

$$\begin{array}{ccc}
\mathrm{Mod}^{\pm}(\Sigma_m) & \xrightarrow{\mathrm{Id}} & \mathrm{Mod}^{\pm}(\Sigma_m) \\
\downarrow \rho & & \downarrow \cong \\
\mathrm{Aut}(\mathcal{D}_k(\Sigma_m)) & \xrightarrow{\bar{\chi}} & \mathrm{Aut}(\mathcal{C}(\Sigma_m))
\end{array}$$

where the two descending arrows are the group homomorphisms induced by the natural actions of the extended mapping class group. The map $\bar{\chi}$ is injective. The right descending arrow is an isomorphism thanks to Ivanov's Theorem (Theorem 1.3.2). In particular it follows that the map $\bar{\chi} \circ \rho$ is an isomorphism, in particular it is surjective. It follows that the map $\bar{\chi}$ must also be surjective, hence it is an isomorphism. From this fact we can deduce that the map ρ is an isomorphism as well. \square

Theorem *B* and Proposition 2.7.7 immediately imply Theorem *A*, that is, the rigidity for the k -separating curve graphs.

As another immediate corollary of the proof of the previous theorem it follows that the group homomorphisms of Lemma 2.8.12, and hence those of Corollary 2.8.13 are not only injective maps, but actual isomorphisms.

2.9 Rigidity of Graphs of Regions

In this section we will prove Theorem *D*. For the definition of graphs of regions, which will be the main objects of this section, we refer to Definition 1.4.3.

Theorem C. *Let $m \geq 5$. Let Σ_m be the m -punctured sphere. Let $\mathcal{A} \subseteq \mathcal{R}(\Sigma_m)$ be a $\text{Mod}^\pm(\Sigma_m)$ -invariant collection of regions, and let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be the associated graph of regions. Then the graph is rigid, i.e. the natural homomorphism*

$$\rho : \text{Mod}^\pm(\Sigma_m) \longrightarrow \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$$

is an isomorphism, if and only if $m \geq 2\nu_{\mathcal{A}} + 1$, and the graph has no vertices with holes and no cork pairs.

The first part of the section will be devoted to prove two necessary conditions in order for a graph of regions to be rigid: these conditions will be an upper bound on the complexity (for whose definition we refer to Definition 1.4.4), and the absence of exchange automorphisms (which we will introduce in Definition 2.9.2) or, equivalently, the absence of regions with holes and cork pairs (see Definition 1.4.5 and Definition 1.4.6).

Based on the complexity, with respect to the number of marked points on the sphere, and the presence of exchange automorphisms the following cases are possible, as we will prove throughout the section:

$m < 2\nu_{\mathcal{A}}$ The graph $\mathcal{G}_{\mathcal{A}}$ is discrete, hence it cannot be rigid;

$m = 2\nu_{\mathcal{A}}$ The graph $\mathcal{G}_{\mathcal{A}}$ is either disconnected or it admits exchange automorphisms, so it is never rigid;

$m > 2\nu_{\mathcal{A}}$ The graph $\mathcal{G}_{\mathcal{A}}$ either admits exchange automorphisms or it is rigid.

We will later characterise the pair of vertices which give rise to exchange automorphisms in purely topological terms. Such vertices will turn out to be exactly regions with holes and cork pairs. We will then use this characterisation to give a complete topological description of all the graph of regions which are rigid, that is Theorem *D*, whose proof occupies the latter part of the section.

We begin by proving a first bound on the complexity.

Proposition 2.9.1. *Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions with $m < 2\nu_{\mathcal{A}}$. Then the graph $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is discrete. In particular*

$$\text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m)) \not\cong \text{Mod}^\pm(\Sigma_m),$$

hence the graph is not rigid.

Proof. Let us argue by contradiction and assume that there exist two disjoint regions $P, Q \in \mathcal{A}$. The region P is contained in a complementary disc of Q , say D . It follows that D contains at least $\nu(P)$ marked points, hence the disc D^c contains at most $m - \nu_{\mathcal{A}} < \nu_{\mathcal{A}}$ marked points. It follows that D^c cannot contain any region in \mathcal{A} , but $Q \subseteq D^c$, which is a contradiction.

It follows that $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is a discrete graph, hence its automorphism group is isomorphic to a permutation group. As remarked in §1.1.7, such a group is either finite or uncountable, so it is never isomorphic to any extended mapping class group. \square

We have just shown that a bound on the complexity of the graph is a necessary condition for rigidity. Such a bound, however, differently from the cases of the graph of discs or curves, will not be a sufficient condition by itself. Indeed, the existence of a specific class of automorphisms, called *exchange automorphisms*, is another obstruction to rigidity. The definition of those automorphisms, first described by John McCarthy and Athanase Papadopoulos in [MP], is the following.

Definition 2.9.2. Let Γ be a graph. An automorphism $\varphi \in \text{Aut}(\Gamma)$ is said to be an *exchange automorphism* if there exist two distinct vertices $v, w \in \mathcal{V}(\Gamma)$ such that $\varphi(v) = w$, $\varphi(w) = v$ and φ fixes every other vertex of Γ . In this case we say that φ exchanges v with w .

Let us remark that two vertices v, w admit a graph automorphism exchanging them if and only if, for every vertex $z \notin \{v, w\}$, the vertex z is adjacent to v if and only if it is adjacent to w . This property can be expressed in graph-theoretical terms as $\text{link}(v) \setminus \{w\} = \text{link}(w) \setminus \{v\}$.

We will now prove that a graph of regions which admits exchange automorphisms cannot be rigid. Although this fact is known in the literature, see for instance [BM2, Theorem 1.8], this proof is direct and self-contained, avoiding the use of vertices with holes or cork pairs at any time.

Proposition 2.9.3. *Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions with $m \geq 2\nu_{\mathcal{A}}$. Let us suppose that there exists an exchange automorphism $\varphi \in \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$. Let η be the group homomorphism*

$$\eta : \text{Mod}^{\pm}(\Sigma_m) \longrightarrow \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$$

induced by the natural action. Then η is not surjective, and in particular $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is not rigid.

The hypothesis $m \geq 2\nu_{\mathcal{A}}$ may apparently seem restrictive but is actually not. Indeed, if $m < 2\nu_{\mathcal{A}}$ we already know that the graph of regions $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is never rigid, thanks to Proposition 2.9.1. As a side note we can observe that, since such a graph of regions is discrete, its automorphism group contains every permutation of the vertices, so in particular any two vertices can be exchanged.

We will now loosely present a proof of Proposition 2.9.3 which is quite short, but at the cost of using an extremely powerful tool in Nielsen Realisation, that is, the fact that, given a finite subgroup of the mapping class group there exists an isomorphic group of homeomorphisms of the surface inducing it.

Proof of Proposition 2.9.3. If there existed two regions exchanged by an exchange automorphism then, up to conjugating with appropriate mapping classes, there would exist infinitely many distinct regions which come in pairs each exchanged by an exchange automorphism. In particular, all of these exchange automorphisms would commute with each other, hence the group $\text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$ would contain an infinite torsion group isomorphic to $(\mathbb{Z}/2\mathbb{Z})^{\mathbb{N}}$. However, thanks to Nielsen Realisation (see [FM, Theorem 7.1] for

a statement for closed surfaces, and [Wh] for a proof which works in the case of punctured surfaces), there exists an upper bound on the order of finite subgroups of the extended mapping class group $\text{Mod}^\pm(\Sigma_m)$ (indeed the proof of [FM, Theorem. 7.4] can be adapted by choosing the opportune minimal area quotient orbifold), hence the automorphisms in the group could not be all induced topologically. \square

As we have seen, Nielsen Realisation Theorem provides a quick proof of Proposition 2.9.3, but at the price of being a very complicated result itself. This is the reason why we are interested in presenting a direct constructive proof of Proposition 2.9.3. This proof will be much longer than the one we have just given, but much more elementary and hopefully more transparent. We will argue by proving that a mapping class that fixes every region apart from at most two must actually be the identity. This is the content of the following proposition, from which Proposition 2.9.3 immediately follows.

Proposition 2.9.4. *Let $\mathcal{G}_A(\Sigma_m)$ be a graph of regions with $m \geq 2\nu_A$. Let $P, Q \in \mathcal{A}$ be two regions. Let $f \in \text{Mod}^\pm(\Sigma_m)$ be an extended mapping class such that for every region $R \in \mathcal{A}$ with $R \neq P, Q$ we have $f(R) = R$. Then $f = \text{Id}$.*

We remark that the previous result immediately imply that the natural group action $\text{Mod}^\pm(\Sigma_m) \curvearrowright \mathcal{G}_A(\Sigma_m)$ is faithful.

The idea behind the proof of the previous proposition is to prove that an extended mapping class which fixes every region in $\mathcal{G}_A(\Sigma_m)$ apart from two must also fix every curve in $\mathcal{C}_{\nu_A}(\Sigma_m)$ apart from two, and directly prove that this implies the mapping class is the identity. Indeed, as the two next lemmas show, mapping classes which fix large enough collections of curves must actually be the identity. We begin with the following, which is Proposition 2.9.4 for $\mathcal{C}_k(\Sigma_m)$. Indeed, it states that there are no exchange automorphisms in $\mathcal{C}_k(\Sigma_m)$.

Lemma 2.9.5. *Let $k \geq 2$ and $m \geq 2k + 1$. Let $\alpha, \beta \subset \Sigma_m$ be two different k -separating curves. Let $f \in \text{Mod}^\pm(\Sigma_m)$ be an extended mapping class such that for every curve $\gamma \in \mathcal{C}_k(\Sigma_m)$ with $\gamma \neq \alpha, \beta$ we have $f(\gamma) = \gamma$. Then $f = \text{Id}$.*

Proof of Lemma 2.9.5. Let us argue by contradiction and assume there exists a mapping class f as in the statement. Let D, D^c be the two complementary discs of α , and assume that D^c contains k -marked points. Since every k -separating curve contained in D^c is isotopic to α , and $\alpha \neq \beta$, it follows that $\beta \subset D$. Since D has at least $k + 1$ marked points it follows that there exists a k -separating curve $\gamma \subset D$ not disjoint from β . Since $f(\gamma) = \gamma$ and f preserves disjointness it is now impossible that $f(\alpha) = \beta$, hence f acts on $\mathcal{C}_k(\Sigma_m)$ as the identity. Since that graph is rigid (Theorem D) we deduce that $f = \text{Id}$. \square

The key ingredient for the proof will actually be the following corollary, which is a strengthened version of the previous lemma, stating that for a mapping class to be the identity, it suffices that it fixes all k -separating curves but two, with no need to check the other curves. We remind the reader that throughout the entire thesis we have always been assuming $m \geq 5$, so that hypothesis is not restrictive.

Lemma 2.9.6. *Let $k \geq 2$, $m \geq 2k$ and $m \geq 5$. Let $\alpha, \beta \in \Sigma_m$ be two different k -separating curves. Let $f \in \text{Mod}^\pm(\Sigma_m)$ be such that for every k -separating curve $\gamma \in \Sigma_m$ with $\gamma \neq \alpha, \beta$ we have $f(\gamma) = \gamma$. Then $f = \text{Id}$.*

Proof. We first deal with the case in which $m \geq 2k + 1$. Let f be as in the statement. Thanks to Lemma 2.9.5 if f also fixes every h -separating curve with $h > k$ then $f = \text{Id}$. It follows that, without loss of generality, we can assume there exists a h -separating curve δ , with $k < h \leq \frac{m}{2}$, such that $f(\delta) \neq \delta$.

Let D be a complementary disc of δ with nonempty intersection with $f(\delta)$. Such a disc contains at least $h > k$ marked points. It follows that there exist infinitely many k -separating curves in D with nontrivial intersection with $f(\delta) \cap D$. In particular there exist one such curve which is fixed by f , which is a contradiction, since such a curve is disjoint from δ but not from $f(\delta)$.

We are now left to the case when $m = 2k$: since $m \geq 5$ then $k \geq 3$. We first claim that f fixes every $(k - 1)$ -separating curve. Indeed, let ω be a $(k - 1)$ -separating curve, and let O be its $(k - 1)$ -punctured complementary disc. The disc O^c contains $l = m - (k - 1) \geq k + 1 \geq 3$ marked points, hence there exist l disjoint O -arcs, which correspond to k -separating curves $\gamma_1, \dots, \gamma_l$, such that none of these curves is equal to either α or β . We observe that there exists a unique $(k - 1)$ -separating curve disjoint from all these curves, and such a curve is ω . Since $f(\gamma_i) = \gamma_i$ it follows that $f(\omega)$ must be disjoint from all the curve γ_i , as well, hence $f(\omega) = \omega$. The claim is proven.

Let now α, β be the two different k -separating curves as in the statement. There exists a $(k - 1)$ -separating curve ω that is disjoint from α but not from β . Since $\omega = f(\omega)$, due to the previous claim, it follows that $f(\alpha)$ has to be disjoint from it, so it cannot be equal to β . It follows that $f(\alpha) = \alpha$ and $f(\beta) = \beta$, hence f acts on $\mathcal{C}_{k-1}(\Sigma_m)$ as the identity: an application of Lemma 2.9.5 yields that $f = \text{Id}$. The proof is complete. \square

We can now prove Proposition 2.9.4.

Proof of Proposition 2.9.4. Let $P, Q \in \mathcal{A}$. Let $f \in \text{Mod}^\pm(\Sigma_m)$ be a mapping class such that for every region $R \in \mathcal{A}$ with $R \neq P, Q$ we have $f(R) = R$.

Let $R \in \mathcal{A}$ be a region such that $\nu(R) = \nu_{\mathcal{A}}$. If R is an annulus that then the two curves in ∂R are $\nu_{\mathcal{A}}$ -separating and isotopic to each other. If R is nonannular, since R is contained in a $\nu(R)$ -punctured disc and we have that $m \geq 2\nu(R)$, every boundary curve of R is a h -separating curve with $h \leq \nu(R) = \nu_{\mathcal{A}}$. In particular there exists a unique $\nu_{\mathcal{A}}$ -separating curve in ∂R .

In either of the two cases there exists a unique isotopy class of $\nu_{\mathcal{A}}$ -separating curves in ∂R , say α_R , and it follows that $f(\alpha_R)$ is the unique $\nu_{\mathcal{A}}$ -separating curve in $\partial f(R)$.

Since every $\nu_{\mathcal{A}}$ -separating curve is in the $\text{Mod}^\pm(\Sigma_m)$ -orbit of α_R it follows that, for every k -separating curve γ , there exists a region $Z \in \mathcal{A}$ such that $\gamma = \alpha_Z$. In particular if $Z \neq P, Q$ we have $f(\alpha_Z) = \alpha_{f(Z)} = \alpha_Z$. It follows that f fixes all but at most two $\nu_{\mathcal{A}}$ -separating curves, hence Lemma 2.9.6 proves that $f = \text{Id}$. \square

We have now proved that the presence of exchange automorphisms constitutes an obstruction to the rigidity of graphs of regions. We will now prove the characterisation in topological terms of the vertices giving rise to such automorphisms. As already mentioned

these two types of vertices are regions with holes (see Definition 1.4.5) and cork pairs (as in Definition 1.4.6).

We will start by proving that vertices with holes produce exchange automorphisms. We recall that an annular region cannot be a vertex with a hole.

Lemma 2.9.7. *Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions. Let P be a vertex with a hole. Then there exists an exchange automorphism $\varphi \in \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$ exchanging P with another vertex.*

Proof. Let D be a complementary disc of P which does not contain any subsurface representing a region in \mathcal{A} . Let $X = P \cup D$: this region is a surface of type S_p^b with $p + b \geq 4$. In fact the disc D contains at least two marked points, since its boundary has to be essential, by definition of region. Since P is non annular then it must contain at least one marked point, so $p \geq 3$. If $b = 0$ that means that ∂P is connected, hence P is a punctured disc. In that case then P must contain at least two marked points, so $p \geq 4$ and our condition holds. It follows that $\text{Mod}(X) \neq \{\text{Id}\}$, hence there exists a mapping class $f \in \text{Mod}(X) < \text{Mod}^{\pm}(\Sigma_m)$ such that $Q = f(P) \neq P$. It follows that $Q \in \mathcal{A}$, that $f(D)$ is a complementary disc of Q not containing any subsurface representing a region in \mathcal{A} , and such that $X = Q \cup f(D)$.

We claim that a region $R \in \mathcal{A}$ is disjoint from P if and only if it is disjoint from Q . This will prove that there exists an exchange automorphism exchanging P and Q . If $R \in \mathcal{A}$ is disjoint from P it follows that it is contained in a complementary disc of P , which cannot be D , since D does not admit subsurfaces in \mathcal{A} . It follows that $R \subseteq X^c$, hence it is disjoint from Q , too. The other implication is completely analogous, and the lemma is proven. \square

We will now prove that the existence of cork pairs produce exchange automorphisms, as well.

Lemma 2.9.8. *Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions. Let (P, D) be a cork pair. Then there exists an exchange automorphism $\varphi \in \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$ exchanging P with D .*

Proof. We only have to prove that a region in \mathcal{A} , which is neither P nor D , is disjoint from P if and only if it is disjoint from D . Let $Q \in \mathcal{A}$ be a region disjoint from P . Since P is an annulus it follows that either $Q \subseteq D$ or $Q \subseteq D^c$. In the latter case Q is disjoint from D and we are done. In the former one, by definition of cork pair, we have that either $Q = D$ or $Q = P$.

Since there exists a representative on P which is contained inside D it follows immediately that every region disjoint from D must be disjoint from Q , as well. The lemma is proven. \square

Lemma 2.9.7 and Lemma 2.9.8, alongside Proposition 2.9.3 immediately imply the following.

Corollary 2.9.9. *Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions admitting either a vertex with a hole or a cork pair. Then group homomorphism*

$$\eta : \text{Mod}^{\pm}(\Sigma_m) \longrightarrow \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$$

induced by the natural action is not surjective. In particular the graph is not rigid.

Since we are interested in rigid graphs, from now on we will only consider graphs without regions with holes and cork pairs. We start with a very simple lemma giving an upper bound on the complexity of regions in such graphs.

Lemma 2.9.10. *Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions without any region with a hole. Let $R \in \mathcal{A}$. Then $\nu(R) \leq m - \nu_{\mathcal{A}}$.*

Proof. Let us first suppose that R is not an annulus with no marked points. Since $R \neq \Sigma_m$ it admits a complementary disc D , hence $R \subseteq D^c$. Since there exists a region in \mathcal{A} represented in D , which would otherwise be a hole, it follows that D contains at least $\nu_{\mathcal{A}}$ marked points, hence the bound on $\nu(R)$.

If R is an annulus with no marked points, let D, D' be its complementary discs. Let h be the number of marked points in D and h' the number of marked points in D' . It follows that $h + h' = m$ and hence

$$\nu(R) = \min\{h, h'\} = m - \max\{h, h'\} \leq m - \nu(R) \leq m - \nu_{\mathcal{A}}. \quad \square$$

Proposition 2.9.1 proves that the bound $m \geq 2\nu_{\mathcal{A}}$ is a necessary condition for the rigidity of a graph $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$. This bound however is not sharp: indeed, as we will see, even though there exist graphs $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ such that $m = 2\nu_{\mathcal{A}}$ and no vertices with holes or cork pairs, those graphs are unfortunately not rigid. It turns out that these graphs are exactly the graph of k -separating discs \mathcal{D}_k and the k -separating curve graph \mathcal{C}_k . This is the content of the following lemma.

Lemma 2.9.11. *Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions without vertices with holes and cork pairs, and such that $m = 2\nu_{\mathcal{A}}$. Then either $\mathcal{G}_{\mathcal{A}}(\Sigma_m) = \mathcal{D}_{\frac{m}{2}}(\Sigma_m)$ or $\mathcal{G}_{\mathcal{A}}(\Sigma_m) = \mathcal{C}_{\frac{m}{2}}(\Sigma_m)$.*

Instead of directly proving the previous lemma we will immediately deduce it from the following result, which is a classification of the elements of minimal complexity, which will be also a technical fact that will turn out to be crucial later. We remark that this is the only result for which the absence of cork pairs is required, as all the other results we will state will only require the absence of vertices with holes.

Lemma 2.9.12. *Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions without vertices with holes such that $m \geq 2\nu_{\mathcal{A}}$. Then any region $R \in \mathcal{A}$ such that $\nu(R) = \nu_{\mathcal{A}}$ is either a $\nu_{\mathcal{A}}$ -punctured disc or an annulus homotopic to a $\nu_{\mathcal{A}}$ -separating curve.*

Moreover, if $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is also without any cork pair, then every two regions $P, Q \in \mathcal{A}$ such that $\nu(P) = \nu(Q) = \nu_{\mathcal{A}}$ have the same topological type.

Proof. Let $R \in \mathcal{A}$ be such that $\nu(R) = \nu_{\mathcal{A}}$. By definition of complexity, there exists a $\nu_{\mathcal{A}}$ -punctured disc D containing R . If $R \neq D$ there would exist a h -punctured disc $D' \subseteq D$ such that $R \cap D' = \emptyset$. If $h = \nu_{\mathcal{A}}$ it follows that D' would be isotopic to D , hence R would be an annulus. Otherwise if $h < \nu_{\mathcal{A}}$ no region in \mathcal{A} could be represented in D' , hence the region R would have a hole. This proves the first part of the lemma.

For the second statement, if there existed both an annulus $P \in \mathcal{A}$ representing a $\nu_{\mathcal{A}}$ -separating curve and a $\nu_{\mathcal{A}}$ -punctured disc $Q \in \mathcal{A}$, then, up to the action of the mapping class group, we could assume Q to be one of the complementary discs of P . However, this would mean that the pair (P, Q) would be a cork pair, providing a contradiction. \square

Our last step in the study of all the graph of regions which are not rigid is to prove that $\mathcal{D}_k(\Sigma_{2k})$ and $\mathcal{C}_k(\Sigma_{2k})$ belong to that class. We have already noticed this behaviour in §1.3.2 and §1.4.1, but we will restate it formally in the following propositions. We start with the case of the k -separating curve graph $\mathcal{C}_k(\Sigma_{2k})$.

Proposition 2.9.13. *Let $k \geq 2$. The graph $\mathcal{C}_k(\Sigma_{2k})$ is discrete. In particular*

$$\text{Aut}(\mathcal{C}_k(\Sigma_{2k})) \not\cong \text{Mod}^\pm(\Sigma_{2k}),$$

and the graph is not rigid.

We now deal with the case of the k -separating disc graph.

Proposition 2.9.14. *Let $k \geq 2$. The graph $\mathcal{D}_k(\Sigma_{2k})$ is a disjoint union of pairs of vertices joined by an edge, each corresponding to a pair of complementary discs. In particular*

$$\text{Aut}(\mathcal{D}_k(\Sigma_{2k})) \not\cong \text{Mod}^\pm(\Sigma_{2k}),$$

and the graph is not rigid.

We can now sum up all the necessary condition for the rigidity of graphs of region we have discovered so far in the following statement, which is exactly the “only if” direction of Theorem *D*.

Corollary 2.9.15. *Let $\mathcal{G}_A(\Sigma_m)$ be a graph of regions which is rigid. Then $m \geq 2\nu_A + 1$ and the graph does not contain any region with holes or any cork pairs.*

From now on we will prove the “if” direction of Theorem *D*. Such a proof will closely resemble that of Theorem *B*. We will define appropriate objects in the graph $\mathcal{G}_A(\Sigma_n)$, which will be *maximal perfect joins*, and prove they are in bijective correspondence with $\mathcal{PC}_{\nu(A)}(\Sigma_m)$, the set of pairs of complementary discs in $\mathcal{D}_{\nu_A}(\Sigma_m)$, as in Definition 2.7.1. This will let us define a $\text{Mod}^\pm(\Sigma_m)$ -equivariant group homomorphism

$$\text{Aut}(\mathcal{G}_A(\Sigma_m)) \longrightarrow \text{Aut}(\mathcal{D}_k(\Sigma_m))$$

which we will prove to be injective. From that fact we can use the rigidity result for the graph of discs to deduce rigidity of the starting graph of regions.

We will start with the following graph-theoretical definition. See [Hara, p.21] and [Mc1, §4].

Definition 2.9.16. Let Γ be a graph. Let $\{V_i\}_1^n$ be a collection of pairwise disjoint subgraphs of Γ such that, for $i \neq j$, every vertex of V_i is adjacent to every vertex of V_j . We define the *join* $V = V_1 * \cdots * V_n$ as the full subgraph of Γ induced by $\bigcup_1^n V_i$. The join is said to be *nontrivial* if at least two of the subgraphs V_i are nonempty. If one of the V_i is non empty and is not a nontrivial join itself, then it is said to be a *component* of the join.

Not having been able to find adequate references in the literature, we will now state and prove a useful technical lemma concerning the behaviour of joins and components under the action of graph automorphisms.

Lemma 2.9.17. *Let Γ be a graph. Let $f : \Gamma \rightarrow \Gamma$ be a graph automorphism. Let $V = V_1 * \dots * V_h$ and $W = W_1 * \dots * W_l$ be joins such that $f(V) \subseteq W$. Then for every $i = 1, \dots, h$, if V_i is a component of V , there exists some j such that $f(V_i) \subseteq W_j$. Moreover, if it also holds that $f(V) = W$ then $f(V_i)$ is a component of W .*

Proof. We argue by contradiction. Since f is invertible we have

$$V_i = f^{-1} \left(\bigcup_{j=1}^l (f(V_i) \cap W_j) \right) = (V_i \cap f^{-1}(W_1)) * \dots * (V_i \cap f^{-1}(W_l)).$$

It follows that if $f(V_i)$ had nonempty intersection with two different W_j then V_i would be a nontrivial join, hence not a component.

For the second part of the statement we notice that, for every $i, i' \in \{1, \dots, h\}$ and $j, j' \in \{1, \dots, l\}$ such that $(i, j) \neq (i', j')$ two vertices $w \in f(V_i) \cap W_j$ and $w' \in f(V_{i'}) \cap W_{j'}$ are adjacent. It follows that we can write W as a join as follows:

$$W = *_{i,j} (f(V_i) \cap W_j).$$

If V_i is a component then there exists a unique j such that $f(V_i) \cap W_j$ is nonempty, hence $f(V_i) \cap W_j = f(V_i)$. Since V_i is a nontrivial join it follows that $f(V_i)$ is also nontrivial, hence a component of W . \square

From the previous lemma, uniqueness of components of a join, which is the content of the following, follows immediately.

Corollary 2.9.18. *Let Γ be a graph. Let V_1, \dots, V_j and W_1, \dots, W_l be nonempty subgraphs such that $V_1 * \dots * V_j = W_1 * \dots * W_l$. Suppose that all the V_i and W_i are components of the join. Then $j = l$ and, up to a permutation of the indices, we have $V_i = W_i$ for every $i = 1, \dots, j$.*

We are now ready to define the class of joins we will use as a bridge between the graph of regions and the graph of discs.

Definition 2.9.19. Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions. A subgraph $V \subseteq \mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is said to be a *perfect join* if it is a join with either two or three components, and is such that at least one but no more than two components are infinite, and at most one component is composed of a single vertex.

A perfect join is said to be *maximal* if it is not properly contained in any other perfect join.

We will denote the set of maximal perfect joins of $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ with $\mathcal{MPJ}_{\mathcal{A}}(\Sigma_m)$.

Since the previous definition is entirely combinatorial, maximal perfect joins are preserved under graph automorphism of $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$, that is $\text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$ acts on $\mathcal{MPJ}_{\mathcal{A}}(\Sigma_m)$.

We recall that $\mathcal{PC}_k(\Sigma_m)$ is the set of pair of complementary discs of $\mathcal{D}_k(\Sigma_m)$, as defined in Definition 2.7.1.

We give the following definitions.

Definition 2.9.20. Let $\{D, D^c\} \in \mathcal{PC}_{\nu(A)}(\Sigma_m)$ be a pair of complementary disc. We define \mathcal{A}_D (resp. \mathcal{A}_{D^c}) as the subgraph of $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ induced by regions $R \in \mathcal{A}$ such that $R \subseteq D$ (resp. $R \subseteq D^c$), and which are not homotopic to ∂D .

We define $\mathcal{A}_{\partial D}$ to be the set of regions $R \in \mathcal{A}$ homotopic to ∂D .

The set $\mathcal{A}_{\partial D}$ is either empty or it only contains an annular region which we will denote, with a slight abuse of notation, as ∂D . Moreover, we notice that if D (resp. D^c) contains more than $\nu_{\mathcal{A}}$ marked points then the set \mathcal{A}_D (resp. \mathcal{A}_{D^c}) is infinite, since it contains infinitely many elements of the $\text{Mod}^{\pm}(\Sigma_m)$ -orbit of any region of complexity $\nu_{\mathcal{A}}$. Lastly, if D contains exactly $\nu_{\mathcal{A}}$ marked points and the graph $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ does not have vertices with holes or corks, then it follows from Lemma 2.9.12 that either $\mathcal{A}_D = \{D\}$ and $\mathcal{A}_{\partial D} = \emptyset$ or $\mathcal{A}_D = \emptyset$ and $\mathcal{A}_{\partial D} = \{\partial D\}$.

Given $\{D, D^c\} \in \mathcal{PC}_{\nu(\mathcal{A})}(\Sigma_m)$ we can define $\lambda(\{D, D^c\})$ to be the complete subgraph of $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ induced by vertices contained in either D or D^c or, equivalently, disjoint from ∂D . Equivalently, we have that $\lambda(\{D, D^c\}) = \mathcal{A}_D * \mathcal{A}_{D^c} * \mathcal{A}_{\partial D}$. The following proposition proves that the subgraph we have just defined is a maximal perfect join.

Lemma 2.9.21. *Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions without vertices with holes, and such that $m \geq 2\nu_{\mathcal{A}} + 1$. Then the subgraph $\lambda(\{D, D^c\}) \subseteq \mathcal{G}_{\mathcal{A}}(\Sigma_m)$ is a maximal perfect join.*

Proof. Without loss of generality let us assume that D is a h -punctured disc, with $\nu_{\mathcal{A}} \leq h \leq \frac{m}{2}$. Since $\lambda(\{D, D^c\})$ can be written as a join as $\mathcal{A}_D * \mathcal{A}_{D^c} * \mathcal{A}_{\partial D}$ we claim that the three subgraphs $\mathcal{A}_D, \mathcal{A}_{D^c}, \mathcal{A}_{\partial D}$ are either empty or components of the join satisfying the conditions of Definition 2.9.19. In particular we have noticed that \mathcal{A}_{D^c} is always infinite, since it contains $m - h \geq \frac{m}{2} > \nu_{\mathcal{A}}$ marked points, while $\mathcal{A}_{\partial D}$ is never infinite. Moreover, if $h > \nu_{\mathcal{A}}$ the subgraph \mathcal{A}_D is infinite as well, whereas in the case when $h = \nu_{\mathcal{A}}$ exactly one subgraph between \mathcal{A}_D and $\mathcal{A}_{\partial D}$ is nonempty, and it only contains a single vertex.

To prove that $\mathcal{A}_D * \mathcal{A}_{D^c} * \mathcal{A}_{\partial D}$ is a perfect join we are left to prove that the three subgraphs which are not empty are not nontrivial joins themselves, hence they are components of the join. Let $P, Q \in \mathcal{A}_{D^c}$ be two distinct regions: we want to show that there exists a region $R \in \mathcal{A}_{D^c}$ which intersects both, thus proving that \mathcal{A}_{D^c} cannot be a nontrivial join. Indeed, if $\mathcal{A}_{D^c} = B * B'$, with $P \in B$ and $Q \in B'$ then every vertex in \mathcal{A}_C would always be adjacent to at least one of P, Q . Lemma 2.9.12 implies that either $\nu_{\mathcal{A}}$ -punctured discs or annuli representing $\nu_{\mathcal{A}}$ -separating curves are always contained in \mathcal{A} . There always exists one of either such discs or curves intersecting both P and Q , and R can be chosen to be such a region. If $h > \nu_{\mathcal{A}}$, that is \mathcal{A}_D is not a singleton, the same argument applies there, otherwise \mathcal{A}_D , as well as $\mathcal{A}_{\partial D}$, is either empty or just a singleton, so in the latter case it cannot be a nontrivial join. Our claim is proven and we can conclude that $\lambda(\{D, D^c\})$ is a perfect join.

We are now left to prove that the perfect join $\lambda(\{D, D^c\}) = \mathcal{A}_D * \mathcal{A}_{D^c} * \mathcal{A}_{\partial D}$ is maximal. Let $V = V_1 * V_2 * V_3$ be a perfect join strictly containing V , with V_i components. Thanks to Lemma 2.9.17 we know that there exist indices $i, j, h \in \{1, 2, 3\}$ such that $\mathcal{A}_{D^c} \subseteq V_i$, $\mathcal{A}_D \subseteq V_j$ and $\mathcal{A}_{\partial D} \subseteq V_h$. Without loss of generality, we can assume $\mathcal{A}_{D^c} \subseteq V_1$. In particular every vertex of V_2 or V_3 must be disjoint from every region in D^c which is represented in \mathcal{A} , and in particular from every $\nu_{\mathcal{A}}$ -punctured disc in D^c , hence it must be contained in D . It follows that $V_2 * V_3 \subseteq \mathcal{A}_D * \mathcal{A}_{\partial D}$.

Since we assumed that $\lambda(\{D, D^c\}) \subsetneq V$ it follows that there exists a region $P \in V_1$ such that $P \notin \mathcal{A}_{D^c}$. It follows that such a region P must intersect D nontrivially. This

means that there exists a region $Q \in \mathcal{A}$ such that $Q \subseteq D$, hence such that we have $Q \in \mathcal{A}_D * \mathcal{A}_{\partial D}$ and which, which is not disjoint from P . From this it follows that $Q \in V_1$. Due to Lemma 2.9.17 it now follows that either $\mathcal{A}_D \subseteq V_1$ or $\mathcal{A}_{\partial D} \subseteq V_1$.

In the former case it follows that $V_2 * V_3 \subseteq \mathcal{A}_{\partial D}$ hence, without loss of generality, we can assume that $V_2 = \mathcal{A}_{\partial D}$ and $V_3 = \emptyset$. If V_2 were empty then V would not be perfect join, as it would only have one component, so $V_2 = \{\partial D\}$. It follows that every region in V_1 must be disjoint from ∂D , hence contained in either D or D^c . It follows that $V_1 \subseteq \mathcal{A}_D * \mathcal{A}_{D^c}$, and hence $V \subseteq \lambda(\{D, D^c\})$, which is a contradiction.

On the other hand, if $\mathcal{A}_{\partial D} \subseteq V_1$ then $V_2 * V_3 \subseteq \mathcal{A}_D$. From Lemma 2.9.17 it follows, without loss of generality, that $\mathcal{A}_D \subseteq V_2$ and $V_3 = \emptyset$. If V_2 were empty then V would only one component, hence it would not be a perfect join. It follows that $V_2 = \mathcal{A}_D \neq \emptyset$, hence every region in V_1 must be disjoint from every region of \mathcal{A} represented in D , hence it must be contained in the disc D^c . It follows that $V_1 \subseteq \mathcal{A}_{D^c} \cup \mathcal{A}_{\partial D}$, hence we have that $V \subseteq \lambda(\{D, D^c\})$, contradicting our assumptions on V . This concludes the proof of the maximality of $\lambda(\{D, D^c\})$. \square

From the previous lemma we can deduce that the following map

$$\lambda : \mathcal{PC}_{\nu_A}(\Sigma_m) \longrightarrow \mathcal{MPJ}_{\mathcal{A}}(\Sigma_m),$$

as defined before Lemma 2.9.21, is defined. The key property of this map λ is to be a bijective correspondence, that it the content of the following.

Proposition 2.9.22. *Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions with no vertices with holes, and such that $m \geq 2\nu_A + 1$. Then, the map*

$$\lambda : \mathcal{PC}_{\nu_A}(\Sigma_m) \longrightarrow \mathcal{MPJ}_{\mathcal{A}}(\Sigma_m),$$

is bijective.

The proof of the previous proposition is quite long and convoluted, so we will break it into various statement. The hardest part will be to prove surjectivity of the map λ , which we will now start dealing with.

Let $V = V_1 * V_2 * V_3$ be a maximal perfect join, as in Definition 2.9.19, where the V_i are either empty sets or components of the join. From this join V we want to construct a punctured disc $D \in \mathcal{D}_{\nu_A}(\Sigma_m)$ such that $V_1 * V_2 * V_3 = \lambda(\{D, D^c\})$. Defining such a disc D is the content of the construction to follow.

Without loss of generality we can assume the component V_1 to be infinite and V_2 to be nonempty. Let

$$X = \bigcup_{R \in V_1} R \subseteq \Sigma_m.$$

First, we claim this is a connected subspace. If not, let C be a connected component of X . Then we define

$$\begin{aligned} V_1^C &= \{R \in V_1 \text{ such that } R \subseteq C\} \\ V_1^{C^c} &= \{R \in V_1 \text{ such that } R \cap C = \emptyset\}. \end{aligned}$$

Obviously V_1^C is nonempty, and if X has more than one connected component then $V_1^{C^c}$ is nonempty, as well. Since $V_1 = V_1 * V_1^{C^c}$, it follows that if X is not connected then V_1 is a nontrivial join, which is a contradiction. We have proven that X is connected. In particular each component of X^c is a disc, possibly with marked points.

Let \mathcal{X} be the set of connected components of X^c containing strictly fewer than $\nu_{\mathcal{A}}$ marked points. Let

$$D = X \cup \bigcup_{O \in \mathcal{X}} O :$$

this is a connected subsurface of Σ_m . A region in \mathcal{A} is disjoint from every region in V_1 if and only if it is disjoint from D . We will show that D is a punctured disc.

We observe that $D \neq \Sigma_m$. Indeed, if it were, there would exist no region in \mathcal{A} disjoint from every region in V_1 , hence the join V would only have a single component, thus contradicting the fact that it is perfect.

We remark that every component of D^c is a disc with at least $\nu_{\mathcal{A}}$ marked points, since D is the union of X with the complementary discs containing strictly less than $\nu_{\mathcal{A}}$ marked points.

We observe that, for each of the two component V_2, V_3 , all the regions contained in that component of the join must be contained in the same connected component of D^c . Indeed, without loss of generality, let Z be a connected component of D^c containing a region in V_2 . Similarly to what we have done before we can define the following:

$$\begin{aligned} V_2^Z &= \{R \in V_2 \text{ such that } R \subseteq Z\} \\ V_2^{Z^c} &= \{R \in V_2 \text{ such that } R \cap Z = \emptyset\}. \end{aligned}$$

Every region in V_2 is disjoint from every region in V_1 , and hence contained in one of the complementary components of D . It then holds that $V_2 = V_2^Z * V_2^{Z^c}$. Since V_2 is not a nontrivial join, it follows that $V_2^{Z^c} = \emptyset$, that is every region in V_2 is contained in the same component of D^c .

We can now prove that D is the punctured disc we were looking for.

Lemma 2.9.23. *Under the previous hypotheses the subsurface $D \subseteq \Sigma_m$ is a punctured disc.*

Proof. We will prove that the subspace D^c is connected, which is equivalent to D being a punctured disc. Indeed, let us argue by contradiction and assume there exist two distinct connected components of D^c , say Z, W , such that every region in the component V_2 , which we assumed to be nonempty, is contained in Z . There are two possible cases: either every region in V_3 is also contained in Z or, without loss of generality, every region in V_3 is contained in W .

We observe that there exists at least one marked point which is not contained in either Z or W . If not, components Z and W would be a h -punctured and a l -punctured disc, respectively, with $h + l = m$, hence complementary discs. It would follow that the only region disjoint from both Z and W would be the annular region homotopic to $\partial Z = \partial W$. It would follow that the component V_1 can contain at most one region, contradicting the assumption that it is infinite. It follows that Z^c contains at least $\nu_{\mathcal{A}} + 1$ marked points. It follows that there exists a region $Q \in \mathcal{A}$ such that $Q \subseteq D \cup W$ and it has

nontrivial intersection with both D , W . Moreover, when the component V_3 is nonempty and every region in it is contained in W , we can assume the region Q to nontrivially intersect at least one region in V_3 . In particular $Q \notin V$. Indeed, in Z^c there exists either a $\nu_{\mathcal{A}}$ -punctured discs, or an annulus homotopic to a $\nu_{\mathcal{A}}$ -separating curve satisfying the previous properties, and one of those types of regions is always contained in \mathcal{A} , thanks to Lemma 2.9.12.

We first deal with the case where V_3 is either empty or every region in it is also contained in Z . It follows that $V' = (V_1 \cup \{Q\}) * V_2 * V_3$ is a nontrivial join strictly containing V . We claim that $V_1 \cup \{Q\}$ is not a nontrivial join, hence V' is a perfect join. Let $V_1 \cup \{Q\} = E * F$. Since $V_1 = (V_1 \cap E) * (V_1 \cap F)$ and V_1 is not a nontrivial join, as it is a component of V , it follows that, without loss of generality, we have $F \cap V_1 = \emptyset$, hence $V_1 \subseteq E$. It follows that either $F \cap \{Q\} = \emptyset$, hence $V_1 \cup \{Q\}$ is a trivial join, or $F = \{Q\}$. However, the latter case is impossible, since Q is not disjoint from D , so there exists at least one vertex in V_1 which is not adjacent to it. We have proven that V' is a perfect join strictly containing V , which is a contradiction to the maximality of V . This concludes the proof in this case.

We are now left with the case where V_3 is nonempty, and every region in it is contained in W . We claim that the graph induced by $V_1 \cup V_3 \cup \{Q\}$ is not a nontrivial join. Let $V_1 \cup V_3 \cup \{Q\} = E * F$. Since $V_1 = (E \cap V_1) * (F \cap V_1)$ and V_1 is not a nontrivial join, as it is a component of V , it follows that, without loss of generality, we have $F \cap V_1 = \emptyset$, hence $V_1 \subseteq E$. From an analogous argument we deduce that either $V_3 \subseteq E$ or $V_3 \subseteq F$. In the former case, either $F = \emptyset$, hence the join $E * F$ is trivial, or $F = \{Q\}$, which is also impossible, since Q was not disjoint from D , hence from at least one region in V_1 . On the other hand, suppose $V_3 \subseteq F$: then both E and F each contain a region which is not disjoint from Q , hence Q cannot belong to either E or F , providing a contradiction. We have proven that $V_1 \cup V_3 \cup \{Q\}$ is not a nontrivial join hence the join $V'' = (V_1 \cup V_3 \cup \{Q\}) * V_2$ is a perfect join strictly containing V , which contradicts the maximality of V . This concludes the proof. \square

Since the disc D contains at least one region in \mathcal{A} , then it must contain at least $\nu_{\mathcal{A}}$ marked points. Similarly, we have observed that its complementary disc must contain at least $\nu_{\mathcal{A}}$ marked points. It follows that $\{D, D^c\} \in \mathcal{PC}_{\nu_{\mathcal{A}}}(\Sigma_m)$. We are left to show that $\{D, D^c\}$ is the preimage of the maximal perfect join $V = V_1 * V_2 * V_3$ under λ , which is the content of the following.

Lemma 2.9.24. *Under the previous hypotheses we have $\lambda(\{D, D^c\}) = V$.*

Proof. The proof will be split into two cases: first, when there exists another infinite component of the join V other than V_1 ; second, when both V_2 and V_3 are finite (one possibly empty).

We first consider the case when the join V admits another infinite component which we can assume to be V_2 , without loss of generality. From an argument completely analogous to the one employed before, it follows that there exists a punctured disc D' such that a region is disjoint from every region in V_2 if and only if it is disjoint from D' . In particular the disc D' is disjoint from every region in V_1 , hence it is disjoint from D .

We will now prove that $D' = D^c$. We argue by contradiction and suppose we have

$D' \subsetneq D^c$. There exists either a $\nu_{\mathcal{A}}$ -punctured discs, or an annulus homotopic to a $\nu_{\mathcal{A}}$ -separating curve R with the following properties: we have $R \subseteq D^c$, but R is not contained in D' ; the region R has nontrivial intersection with D' , hence with at least one region in V_2 ; moreover, when V_3 is nonempty, then R has nontrivial intersection with at least one region in V_3 ; lastly, $R \notin V_3$, which is possible since V_3 is finite. Thanks to Lemma 2.9.12 we have $R \in \mathcal{A}$. We claim that $V_2 \cup V_3 \cup \{R\}$ is not a nontrivial join. In order to prove the claim let $V_2 \cup V_3 \cup \{R\} = E * F$. Since $V_2 = (E \cap V_2) * (F \cap V_2)$ and V_2 is not a nontrivial join, as it is a component of V , it follows that, without loss of generality we can assume $F \cap V_2 = \emptyset$, hence $V_2 \subseteq E$. Similarly, either $V_3 \subseteq E$ or $V_3 \subseteq F$. In the former case, either $F = \emptyset$, hence the join $E * F$ is trivial, or it must hold $F = \{R\}$, which is impossible, since R is not disjoint from every region in V_2 . Suppose that $V_3 \subseteq F$, and $V_3 \neq \emptyset$: there exists a region in both V_2 and V_3 which is not disjoint from R , hence the region R cannot belong to either E or F . If $V_3 = \emptyset$ then $F = \{R\}$, which is once again impossible. We have proven our claim, and $V_2 \cup V_3 \cup \{R\}$ is not a nontrivial join. It follows that the join $V_1 * (V_2 \cup V_3 \cup \{R\})$ is perfect and strictly contains V , contradicting maximality of V . We have proven that $D' = D^c$.

Since every region in V_3 must be disjoint from both D and $D' = D^c$, it follows that V_3 is either empty or composed of the single vertex given by the annular region homotopic to ∂D . It follows that $V \subseteq \mathcal{A}_D * \mathcal{A}_{D^c} * \mathcal{A}_{\partial D} = \lambda(\{D, D^c\})$ and since V is maximal, and $\lambda(\{D, D^c\})$ is a perfect join thanks to Lemma 2.9.21, equality holds.

We are now left with the case when both V_2 and V_3 are finite. Without loss of generality we can assume $V_2 \neq \emptyset$. We will prove that this implies that D^c contains exactly $\nu_{\mathcal{A}}$ marked points. Indeed, we have already noticed that D^c contains at least $\nu_{\mathcal{A}}$ marked points. If it contained more then there would exist infinitely many regions in \mathcal{A} represented in D^c . In particular there would exist either a $\nu_{\mathcal{A}}$ -punctured discs, or an annulus homotopic to a $\nu_{\mathcal{A}}$ -separating curve R , not contained in either V_2 or V_3 and with nontrivial intersection with at least one region in V_2 . Moreover, if V_3 is nonempty, we can also assume R to have nontrivial intersection with at least one region in V_3 . Thanks to Lemma 2.9.12 we have $R \in \mathcal{A}$. We now claim that $V_2 \cup V_3 \cup \{R\}$ is not a nontrivial join. Indeed, let $V_2 \cup V_3 \cup \{R\} = E * F$. Since $V_2 = (E \cap V_2) * (F \cap V_2)$ and V_2 is not a nontrivial join, as it is a component of V , it follows that, without loss of generality we can assume $F \cap V_2 = \emptyset$, hence $V_2 \subseteq E$. Similarly, either $V_3 \subseteq E$ or $V_3 \subseteq F$. In the former case, either $F = \emptyset$, hence the join $E * F$ is trivial, or it must hold $F = \{R\}$, which is impossible, since R is not disjoint from every region in V_2 . Suppose that $V_3 \subseteq F$, and $V_3 \neq \emptyset$: there exists a region in both V_2 and V_3 which is not disjoint from R , hence the region R cannot belong to either E or F . If $V_3 = \emptyset$ then $F = \{R\}$, which is once again impossible. We have proven our claim, that is $V_2 \cup V_3 \cup \{R\}$ is not a nontrivial join. It follows that the join $V_1 * (V_2 \cup V_3 \cup \{R\})$ is a perfect join strictly containing V , contradicting maximality of V : this proves that D^c contains exactly $\nu_{\mathcal{A}}$ marked points.

Since D^c contains exactly $\nu_{\mathcal{A}}$ marked points then, thanks to Lemma 2.9.12, there exists a unique region represented in D^c . This region is the only vertex of V_2 , which equals either \mathcal{A}_{D^c} or $\mathcal{A}_{\partial D}$, and we have $V_3 = \emptyset$. It follows that $V \subseteq \mathcal{A}_D * \mathcal{A}_{D^c} * \mathcal{A}_{\partial D} = \lambda(\{D, D^c\})$ and, since V is maximal, and $\lambda(\{D, D^c\})$ is a perfect join thanks to Lemma 2.9.21, equality holds. \square

We are now ready to conclude the proof that the map λ is a bijection.

Proof of Proposition 2.9.22. Surjectivity of λ follows from the previous discussion, in particular Lemma 2.9.24.

We will now prove injectivity. Let $\{B, B^c\}, \{C, C^c\} \in \mathcal{PC}_{\nu_A}(\Sigma_m)$ be such that $\lambda(\{B, B^c\}) = \lambda(\{C, C^c\})$. Since the two joins are equal Lemma 2.9.17 implies that the component must be equal in pairs. Without loss of generality we can assume $\mathcal{A}_B = \mathcal{A}_C$, and that this component is infinite. It follows that every region in $\mathcal{A}_{B^c} * \mathcal{A}_{\partial B}$ is disjoint from every region of \mathcal{A} represented in C , hence must be contained in C^c . From this we can deduce that actually $B^c \subseteq C^c$, and an analogous argument proves that $C^c \subseteq B^c$, hence $\{B, B^c\} = \{C, C^c\}$. Injectivity is proven. \square

Let now $g \in \text{Aut}(\mathcal{G}_A(\Sigma_m))$. We can define the map

$$\Phi(g) : \mathcal{PC}_{\nu_A}(\Sigma_m) \longrightarrow \mathcal{PC}_{\nu_A}(\Sigma_m)$$

given by conjugation by λ , that is

$$\Phi(g) = \lambda^{-1} \circ g \circ \lambda.$$

Our goal is to use this map to induce an automorphism of $\mathcal{GPC}_{\nu_A}(\Sigma_m)$. In order to achieve such a result, and thanks to Lemma 2.7.8, we need to express nesting between pairs of complementary discs (Definition 2.7.3) in terms of a combinatorial property of maximal perfect joins. To do so we start by giving the following definition, which is a slight variation to the one presented in [Mc1, §4].

Definition 2.9.25. We say that two joins $V, W \in \mathcal{MPJ}_A(\Sigma_m)$ are *compatible* if they can be written as $V = V_1 * V_2$ and $W = W_1 * W_2$, where V_1 and W_1 are components of V and W , respectively, and $V_1 \subsetneq W_1$.

We remark that the Corollary 2.9.18 proves that if two joins V, W are compatible then $V \neq W$.

The following result relates the notion of nesting for pairs of complementary discs, that is adjacency in the graph $\mathcal{GPC}_{\nu_A}(\Sigma_m)$, and compatibility between maximal perfect joins.

Lemma 2.9.26. *Let $\{B, B^c\}, \{C, C^c\} \in \mathcal{PC}_{\nu_A}(\Sigma_m)$. Then $\{B, B^c\}$ and $\{C, C^c\}$ are nested if and only if $\lambda(\{B, B^c\})$ and $\lambda(\{C, C^c\})$ are compatible.*

Proof. Let $\{B, B^c\}, \{C, C^c\} \in \mathcal{PC}_{\nu_A}(\Sigma_m)$ be nested, that is, without loss of generality $B \subsetneq C$. In particular nestedness implies that $\{B, B^c\} \neq \{C, C^c\}$, hence we have $\lambda(\{B, B^c\}) \neq \lambda(\{C, C^c\})$, thanks to injectivity of λ . Following the notation introduced in Definition 2.9.20, the join $\lambda(\{B, B^c\})$ can be written as either $\mathcal{A}_B * V$ or $\mathcal{A}_{\partial B} * V$, and $\lambda(\{C, C^c\}) = \mathcal{A}_C * W$. Since $B \subsetneq C$ it follows that $\mathcal{A}_B \subsetneq \mathcal{A}_C$ and $\mathcal{A}_{\partial B} \subsetneq \mathcal{A}_C$, and at least one of the left hand sides is nonempty, hence the two joins are compatible.

For the converse, let us assume that $\lambda(\{B, B^c\}) = V * V'$ and $\lambda(\{C, C^c\}) = W * W'$ are distinct maximal joins with V, W components and $V \subsetneq W$.

If $V = \mathcal{A}_B$ then either $\mathcal{A}_B \subsetneq \mathcal{A}_C$ or $\mathcal{A}_B \subsetneq \mathcal{A}_{C^c}$, where equality cannot happen as it would imply $\lambda(\{B, B^c\}) = \lambda(\{C, C^c\})$. Indeed, if it held $\mathcal{A}_B \subsetneq \mathcal{A}_{C^c}$ then \mathcal{A}_B would

have to be empty, as it would be strictly contained in $\mathcal{A}_{\partial C}$, which contains at most one element. It follows that every region contained in B is contained in C , hence $B \subsetneq C$. Strict inclusion follows from the fact that there exists at least a region belonging to \mathcal{A}_C but not to \mathcal{A}_B , hence contained in C but not in B . It follows that $\{B, B^c\}$ and $\{C, C^c\}$ are nested. If $V = \mathcal{A}_{B^c}$ the argument is completely analogous.

We are now left to the case when $V = \mathcal{A}_{\partial B}$. The inclusion $\mathcal{A}_{\partial B} \subsetneq \mathcal{A}_{\partial C}$ is impossible, as either sets contain at most one vertex, so in order to have a proper inclusion the set $\mathcal{A}_{\partial B}$ would have to be empty. It follows that, up to exchanging C with C^c , it must hold that $\mathcal{A}_{\partial B} \subsetneq \mathcal{A}_C$. In particular $\partial B \subseteq C$ hence either $B \subseteq C$ or $B^c \subseteq C$. Equality is impossible since it would imply $\{B, B^c\} = \{C, C^c\}$ and hence $\lambda(\{B, B^c\}) \neq \lambda(\{C, C^c\})$, which is a contradiction to the compatibility of the two joins. It follows that the two pairs are nested. \square

Lemma 2.7.8 and Lemma 2.9.26 prove that there exists a group homomorphism

$$\tilde{\chi} : \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m)) \longrightarrow \text{Aut}(\mathcal{D}_{\nu_{\mathcal{A}}}(\Sigma_m))$$

which is defined, for $g \in \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$ and for $\{D, D^c\} \in \mathcal{PC}_k(\Sigma_m)$, as

$$\tilde{\chi}(g) \cdot \{D, D^c\} = \lambda^{-1}(g \cdot \lambda(\{D, D^c\})).$$

The key property of this map we will use to prove rigidity of the graphs of regions is injectivity, which is the content of the following.

Proposition 2.9.27. *Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions with no vertices with holes and cork pairs such that $m \geq 2\nu_{\mathcal{A}} + 1$. The map*

$$\tilde{\chi} : \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m)) \longrightarrow \text{Aut}(\mathcal{D}_{\nu_{\mathcal{A}}}(\Sigma_m))$$

is a $\text{Mod}^{\pm}(\Sigma_m)$ -equivariant injective group homomorphism.

In order to prove the previous proposition we will make use of the fact that a graph homomorphism which fixes all the minimal complexity regions must be the identity. This result will be a key component in our proof of injectivity (which differs from the one in [Mc1, Proof of Theorem 3]), since the joins corresponding to pairs $\{D, D^c\}$ with D a $\nu_{\mathcal{A}}$ -punctured discs are particularly nice, and indeed characterised as follows.

Lemma 2.9.28. *Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions with no vertices with holes and cork pairs. Given a $\nu_{\mathcal{A}}$ -punctured disc D , the maximal perfect join $\lambda(\{D, D^c\})$ corresponds to the join of an infinite set (\mathcal{A}_{D^c}) with a subgraph made of a single vertex, which is the nonempty set between $\mathcal{A}_D = \{D\}$ and $\mathcal{A}_{\partial D} = \{\partial D\}$.*

We will call a join as in the previous lemma a *cone*.

The proof of the previous result follows directly from the definition of the map λ (Definition 2.9.20) and the classification of minimal complexity surfaces given by Lemma 2.9.12, which is the reason why the absence of cork pairs is required in both the previous and the following lemma.

Lemma 2.9.29. *Let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be a graph of regions with no vertices with holes and no cork pairs such that $m \geq 2\nu_{\mathcal{A}} + 1$. Let $g \in \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$ such that for every region $R \in \mathcal{A}$ with $\nu(R) = \nu_{\mathcal{A}}$ we have $g(R) = R$. Then $g = \text{Id}$.*

Proof. Let $R \in \mathcal{A}$ be a region such that $\nu(R) > \nu_{\mathcal{A}}$. Let D be a complementary disc of R : we claim that D contains at least $\nu_{\mathcal{A}}$ marked points. If not, no region in \mathcal{A} could be represented in D , and the only possibility for R not to have a hole would be if it were an annular surface. In that case, up to isotopy, we would have $R \subseteq D$, hence $\nu(R) \leq \nu_{\mathcal{A}}$, which is a contradiction. We claim that, up to isotopy, the region $g(R)$ is disjoint from D . Indeed, if $g(R)$ had nontrivial intersection with D , which contains at least $\nu_{\mathcal{A}}$ marked points, there would exist a minimal complexity region $P \in \mathcal{A}$ such that $P \subseteq D$ but $P \cap g(R) \neq \emptyset$. Since P and R are disjoint, hence joined by an edge in $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$, the regions $g(P) = P$ and $g(R)$ must be disjoint, as well: this is a contradiction, and the claim is proven.

We have just shown that, up to isotopy, we have $g(R) \subseteq R$. If $\nu(g(R)) > \nu_{\mathcal{A}}$ we can apply the previous argument once again to the automorphism g^{-1} , obtaining that $R = g^{-1}(g(R)) \subseteq g(R)$, hence $g(R) = R$. If $\nu(g(R)) = \nu_{\mathcal{A}}$ then g^{-1} fixes $g(R)$, hence we have that $R = g^{-1}(g(R)) = g(R)$. This proves that $g(R) = R$ for every region $R \in \mathcal{A}$, hence $g = \text{Id}_{\mathcal{G}_{\mathcal{A}}(\Sigma_m)}$. \square

We can now prove the injectivity of the homomorphism $\tilde{\chi}$.

Proof of Proposition 2.9.27. Let $g \in \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$ such that $\tilde{\chi}(g) = \text{Id}_{\mathcal{GPC}_{\mathcal{A}}(\Sigma_m)}$. We claim that $g = \text{Id}_{\mathcal{G}_{\mathcal{A}}}$. Thanks to Lemma 2.9.29 we only have to prove that g fixes every minimal complexity vertex.

Let $R \in \mathcal{A}$ such that $\nu(R) = \nu_{\mathcal{A}}$. Thanks to Lemma 2.9.12 R is either a $\nu_{\mathcal{A}}$ -punctured disc or an annular surface representing a $\nu_{\mathcal{A}}$ -separating curve. In the first case let $D = R$, in the second one let $D \subseteq \Sigma_m$ be the only $\nu_{\mathcal{A}}$ -punctured disc containing R . As we have noticed the join $\lambda(\{D, D^c\})$ has two components: one is infinite and the other is composed of the single vertex R . Since by definition we have that

$$\{D, D^c\} = \tilde{\chi}(g) \cdot \{D, D^c\} = \lambda^{-1}(g \cdot \lambda(\{D, D^c\}))$$

it follows that $\lambda(\{D, D^c\}) = g \cdot \lambda(\{D, D^c\})$. Thanks to Lemma 2.9.17 the component of $\lambda(\{D, D^c\})$ composed of the single vertex R is mapped onto a component of $\lambda(\{D, D^c\})$: since it cannot be exchanged with the infinite component, it has to be fixed. From this it follows that $g(R) = R$ and the proposition is proven. \square

We are now ready to state Theorem D once again and prove it.

Theorem D. *Let $\nu_{\mathcal{A}} \subseteq \mathcal{R}(\Sigma_m)$ and let $\mathcal{G}_{\mathcal{A}}(\Sigma_m)$ be the associated graph of regions. Then the graph is rigid, i.e. the natural homomorphism*

$$\rho : \text{Mod}^{\pm}(\Sigma_m) \longrightarrow \text{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m))$$

is an isomorphism, if and only if $m \geq 2\nu_{\mathcal{A}} + 1$ and the graph has no vertices with holes and no cork pairs.

Proof. Corollary 2.9.15 proves the “only if” direction of the statement.

For the “if” direction Proposition 2.9.27 provides us with the following commutative diagram

$$\begin{array}{ccc}
 \mathrm{Mod}^{\pm}(\Sigma_m) & \xrightarrow{\mathrm{Id}} & \mathrm{Mod}^{\pm}(\Sigma_m) \\
 \eta \downarrow & & \downarrow \rho \\
 \mathrm{Aut}(\mathcal{G}_{\mathcal{A}}(\Sigma_m)) & \xrightarrow{\tilde{\chi}} & \mathrm{Aut}(\mathcal{D}_{\nu_{\mathcal{A}}}(\Sigma_m))
 \end{array}$$

where the vertical maps are the homomorphism induced by the natural actions. Since $m \geq 2\nu_{\mathcal{A}} + 1$ Theorem *B* proves that ρ is an isomorphism, hence the composition $\tilde{\chi} \circ \eta$ is an isomorphism as well. It follows that $\tilde{\chi}$ must be surjective, hence an isomorphism. This implies that η is an isomorphism, as well. \square

Chapter 3

Exhaustions by Finite Rigid Sets

3.1 Outline of the Chapter

In this chapter we will prove Theorem *E* and its direct corollary Theorem *F*. Moreover, we recall the definition of rigidity we will use throughout this chapter.

Definition 3.1.1. Let Γ be a graph. A subgraph $X \subseteq \Gamma$ is said to be *rigid* if for every graph embedding $i : X \hookrightarrow \Gamma$ there exists a graph automorphism $\varphi \in \text{Aut}(\Gamma)$ such that $i = \varphi|_X$. A graph which is rigid, as a subgraph of itself, is said to be *co-Hopfian*.

Theorem E. Let $S = \Sigma_7$ or $S = \Sigma_8$. Then the strongly separating curve graph $\mathcal{C}_{ss}(S)$ admits an exhaustion by finite rigid sets

$$\bigcup_{i \in \mathbb{N}} X_i = \mathcal{C}_{ss}(S).$$

Moreover, every subgraph X_i has trivial pointwise stabiliser.

Theorem F. Let $S = \Sigma_7$ or $S = \Sigma_8$. Then, for every injective graph self-embedding $i : \mathcal{C}_{ss}(S) \hookrightarrow \mathcal{C}_{ss}(S)$, there exists an extended mapping class $f \in \text{Mod}^\pm(S)$ such that $i = f$, that is the self embedding i coincides with the map induce on the curve graph by the mapping class f .

The objects we will mostly deal with, throughout this chapter, will be the strongly separating curve graphs for the 7-holed sphere $\mathcal{C}_{ss}(\Sigma_7)$ and 8-holed sphere $\mathcal{C}_{ss}(\Sigma_8)$, introduced in Definition 1.3.3. We recall that these graphs correspond to the 3-separating curve graphs $\mathcal{C}_3(\Sigma_7)$ and $\mathcal{C}_3(\Sigma_8)$, as in Definition 1.3.6.

We observe that every curve in the graph $\mathcal{C}_{ss}(\Sigma_7)$ is 3-separating, hence such graph is the same as the strict 3-separating curve graph $\mathcal{C}_{(3)}(\Sigma_7)$, also introduced in Definition 1.3.6. On the other hand, the strongly separating curve graph of the 8-holed sphere includes both 3-separating and 4-separating curves (see Definition 1.3.5 for the definition), hence the strict 3-separating curve graph $\mathcal{C}_{(3)}(\Sigma_8)$, which is induced by the 3-separating curves only, is a proper subgraph of $\mathcal{C}_{ss}(\Sigma_8)$. This graph will play an important role in Section 3.5, and we will prove that it also admits an exhaustion by finite rigid sets.

In order to expand such subgraph we will employ the action of a suitable finite collection G of mapping classes, such that $\langle G \rangle = \text{Mod}^\pm(\Sigma_m)$. The exhaustion we are

looking for is given by the recurring definition $X_{i+1} = G \cdot X_i$. The fact that G generates the entire mapping class group, hence it acts transitively on vertices in the relevant ambient graph, will imply that the sequence we have defined exhaust the entire graph.

The hardest part to prove is the rigidity of the subgraphs X_i . Indeed, such a subgraph can easily be written as a union of rigid subgraphs. Unions of rigid subgraphs are not rigid, in general, but, when the stabilisers of the pairwise intersections are trivial, they are. The arguments described so far can be stated in purely graph-theoretical terms, and we will develop such general machinery in Section 3.2.

Aside from the graph-theoretical arguments, the most part of our proof will actually be dedicated to the choice of a suitable collection of mapping classes, and the proof of the various rigidity property needed to apply the aforementioned general results. These arguments will be both combinatorial, as they involve the study of projections to Kneser graphs, and topological. These are the instances in which the differences between the rigid subgraphs will make our arguments, although analogous in spirit, divergent from a technical viewpoint.

The structure of the chapter will be as follows. In Section 3.2 we will introduce the graph-theoretical machinery we will use to produce exhaustions. Moreover, we will prove that graphs which admit an exhaustion by rigid sets are co-Hopfian, which is another purely graph-theoretical result, hence proving that Theorem *F* immediately follows from Theorem *E*.

In Sections 3.4 and 3.5 we will verify the hypotheses in order to apply the aforementioned graph-theoretical machinery, for the case of the 7-holed sphere and the 8-holed sphere respectively. From the arguments in these section another proof of the combinatorial rigidity of such graphs, that is Theorem 1.3.4 for $g = 0$ and $p = 7, 8$, will follow.

Lastly, in Section 3.6 we will produce an example of a graph which admits many nice metric and combinatorial properties but which is not co-Hopfian, and deduce that it does not admit any exhaustion by rigid sets. The example we will use will be a graph of loops on a infinite-type surface, and our observations allow for many similar examples to be developed.

3.2 Graph-Theoretical Machinery

In this section we will introduce the graph-theoretical machinery we will use to produce exhaustions, and prove that Theorem *E* implies Theorem *F*.

From now on let Γ be a graph. The idea of our construction is to start with a rigid subgraph $X_0 \subseteq \Gamma$ and a finite collection of graph automorphism $H \subset \text{Aut}(\Gamma)$ and let it repeatedly act on said subgraph X to obtain an exhaustion.

We will mostly be interested in subgraphs with trivial pointwise stabiliser for the action of the group of graph automorphisms. A subgraph $X \subseteq \Gamma$ with this property is such that for every graph automorphism $\varphi \in \text{Aut}(\Gamma)$ such that $\varphi|_X = \text{Id}_X$ we have $\varphi = \text{Id}_\Gamma$. This property is what Aramayona and Leininger call *weak rigidity* in [AL2, Definition 3.6]. However, since this property is not implied by rigidity (as in Definition 1.6.1; it would be if a uniqueness property were also required in that definition) we will avoid such terminology. In what follows, when we refer to the stabiliser of a subgraph, unless otherwise stated, it has to be interpreted as the pointwise stabiliser for the action of the group of graph automorphisms.

The following useful lemma is an immediate observation.

Lemma 3.2.1. *Let Γ be a graph. Let $X \subset Y \subset \Gamma$ be subgraphs, such that X has trivial pointwise stabiliser. Then Y has trivial pointwise stabiliser.*

Opposite to what happens with the previous proposition a superset of a rigid subgraph is in general not rigid (for an example of this behaviour see [AL2, Proposition 3.2]). This makes expanding rigid sets to bigger rigid sets tricky in general, but luckily the following statement is what will make it possible in our cases.

Lemma 3.2.2. *Let Γ be a graph. Let $X, Y \subset \Gamma$ be rigid subgraphs such that $X \cap Y$ has trivial pointwise stabiliser. Then $X \cup Y$ is rigid and has trivial pointwise stabiliser.*

Proof. Let $i : X \cup Y \hookrightarrow \Gamma$ be an injective graph homomorphism. Since X, Y are rigid there exists two automorphisms $\varphi, \psi \in \text{Aut}(\Gamma)$ such that $\varphi|_X = i|_X$ and $\psi|_Y = i|_Y$. Since $\varphi|_{X \cap Y} = i|_{X \cap Y} = \psi|_{X \cap Y}$ the triviality of the stabiliser of $X \cap Y$ proves that $\varphi = \psi$. It follows that the embedding i is the restriction of the automorphism $\varphi = \psi$, hence $X \cup Y$ is rigid. Since $X \cap Y \subseteq X \cup Y$ triviality of the stabiliser follows immediately from Lemma 3.2.1. \square

In what follows we will use the following corollary, which is a generalised version of the previous result.

Corollary 3.2.3. *Let Γ be a graph. Let $X_1, \dots, X_n \subset \Gamma$ be rigid subgraphs such that, for every $i \neq j$, the graph $X_i \cap X_j$ has trivial pointwise stabiliser. Then the subgraph $X_1 \cup \dots \cup X_n$ is rigid and has trivial pointwise stabiliser.*

Proof. We will proceed on induction on n . If $n = 2$ then the statement is exactly Lemma 3.2.2. Let us now suppose the thesis to be true for n and prove we have for $n + 1$. In this case $X_1 \cup \dots \cup X_{n+1} = (X_1 \cup \dots \cup X_n) \cup X_{n+1}$ where $Y = (X_1 \cup \dots \cup X_n)$ is rigid by inductive hypothesis. Moreover $X_1 \cap X_{n+1} \subset Y \cap X_{n+1}$ hence the right hand side has trivial pointwise stabiliser due to Lemma 3.2.1. An immediate application of

Lemma 3.2.2 now proves that $Y \cup X_{n+1}$ is rigid. Triviality of the pointwise stabiliser follows immediately from Lemma 3.2.1. \square

We are now ready to sum up the previous observation into the machinery we will use to produce rigid exhaustions, which is a graph-theoretical version of [AL2, Proposition 3.13].

Proposition 3.2.4. *Let Γ be a graph. Let $H \subseteq \text{Aut}(\Gamma)$ be a finite set of automorphisms, such that $\text{Id}_\Gamma \in H$ and $H^{-1} = H$. Let $X \subset \Gamma$ be a rigid subgraph such that for every $g \in H$ the graph $g \cdot X \cap X$ has trivial pointwise stabiliser. Then the sequence defined by*

$$X_0 = X \text{ and } X_{i+1} = H \cdot X_i$$

is an ascending sequence of rigid sets with trivial pointwise stabilisers.

Moreover, if X intersects every orbit of vertices of Γ under the action of $\langle H \rangle$ the aforementioned sequence is an exhaustion.

Proof. First we prove that (X_i) is an ascending sequence of rigid sets, by proceeding by induction on i . The graph X_0 is rigid by hypothesis. We will now assume X_i to be rigid and that for every $j \leq i$ we have $X_j \subset X_i$. For every $g \in H$ it follows immediately that $g \cdot X_i$ is rigid, as it is an isomorphic copy of X_i . Moreover for every $g' \in H$ we have that $(g')^{-1} \in H$, and hence

$$g \cdot X_0 \cap X_0 \subseteq g \cdot X_i \cap g' \cdot (g')^{-1} \cdot X_{i-1} \subseteq g \cdot X_i \cap g' \cdot X_i,$$

as $g' \cdot X_{i-1} \subseteq X_i$, hence $g \cdot X_i \cap g' \cdot X_i$ has trivial pointwise stabiliser thanks to Lemma 3.2.1, so X_i does as well.

We can now apply Corollary 3.2.3 to the collection $\{g \cdot X_i\}_{g \in H}$ proving that X_{i+1} is rigid. Moreover we have

$$X_i = \bigcup_{g \in H} g \cdot X_{i-1} \subseteq \bigcup_{g \in H} g \cdot X_i = X_{i+1}$$

and the ascending property is proven.

Let us now suppose that X intersects every orbit of vertices of Γ under the action of $\text{Aut}(\Gamma)$. We will now prove that we have

$$\bigcup X_i = \Gamma.$$

Indeed, for every vertex $v \in \Gamma$ there exist a vertex $x \in X$ and a graph automorphism $\varphi \in \langle H \rangle$ such that $\varphi(x) = v$. Therefore there exists some $n \in \mathbb{N}$ such that $\varphi \in H^n$, where H^n is the set of products of n elements of H . It follows that

$$v \in H^n \cdot X_0 \subseteq X_n.$$

This concludes the proof. \square

We will apply the previous result to the case of graphs of curves in the following sections.

We will now conclude the section by showing that Theorem *F* follows from Theorem *E*. That will follow from the following more general result, which will also be useful in §3.6.

Proposition 3.2.5. *Let Γ be a graph admitting an exhaustion by rigid sets (X_j) with trivial pointwise stabiliser. Then Γ is co-Hopfian, that is, for every injective graph homomorphism $i : \Gamma \hookrightarrow \Gamma$ there exists a graph automorphism $\varphi \in \text{Aut}(\Gamma)$ such that $i = \varphi$.*

Proof. Since every subgraph X_j is rigid for every j there exists a graph automorphism $\varphi_j \in \text{Aut}(\Gamma)$ such that $i|_{X_j} = \varphi_j|_{X_j}$. Such graph automorphism is unique thanks to the fact that the pointwise stabilisers of the graphs X_j are trivial. For $j < k$ we have $X_j \subseteq X_k$ hence we have that $\varphi_k|_{X_j} = \varphi_j|_{X_j}$, and thanks to the uniqueness property we have that $\varphi_j = \varphi_k$. From now on we will denote this automorphism by φ .

Let v be a vertex of Γ : since the sequence (X_j) is an exhaustion there exists an index k such that $v \in X_k$. Hence it follows that $i(v) = \varphi_k(v) = \varphi(v)$. Since this holds for every vertex we conclude that $i = \varphi$. \square

The following statement is an alternative version of Proposition 3.2.4 and Proposition 3.2.5, expressed in terms of group actions on the graphs. Its proof only requires minor changes in the arguments presented in this section.

Proposition 3.2.6. *Let Γ be a graph. Let $H \subseteq \text{Aut}(\Gamma)$ be a finite set of automorphisms, such that $\text{Id}_\Gamma \in H$ and $H^{-1} = H$. Let $X \subset \Gamma$ be a subgraph such that for every graph embedding $i : X \hookrightarrow \Gamma$ there exists a graph automorphism $\varphi \in \langle H \rangle$ such that $i = \varphi|_X$ (we will denote this property by saying that X is $\langle H \rangle$ -rigid). Moreover, let us assume that if $g \in H$ the graph $g \cdot X \cap X$ has trivial stabiliser under the action of $\langle H \rangle$. Then the sequence defined by*

$$X_0 = X \text{ and } X_{i+1} = H \cdot X_i$$

is an ascending sequence of $\langle H \rangle$ -rigid sets with trivial stabilisers for the action of $\langle H \rangle$.

Moreover, if X intersects every orbit of vertices of Γ under the action of $\langle H \rangle$ the aforementioned sequence is an exhaustion. In that case we have $\langle H \rangle = \text{Aut}(\Gamma)$.

We can observe that in the previous result we have required the stronger hypothesis of X being $\langle H \rangle$ -rigid, and not only rigid, but at the same time obtained a combinatorial rigidity result for the graph in the equality $\langle H \rangle = \text{Aut}(\Gamma)$. In the two cases of interest to us we will be able to apply Proposition 3.2.6, hence getting another proof of combinatorial rigidity for the strongly separating curve graphs of the 7 and the 8-holed sphere, that is Theorem 1.3.4. However, we will only cursorily remark about the application of the previous proposition, since the application of Proposition 3.2.4 is more straightforward and still suffices our needs.

3.3 One-Third Dehn Twists

One major point the cases of the 7 and 8-holed sphere cases will have in common is the types of maps used in the collection of mapping classes used to generate the exhaustion. Indeed, such collection, will be composed of (right) *half Dehn twists* and some other roots of Dehn twists, which we will later introduce and will call *one-third Dehn twists*.

We recall that, given a 2-separating curve γ on Σ_m the (right) half Dehn twist around γ is well defined, and we will denote it with H_γ . The half-Dehn twist around a 2-separating curve is a self-homeomorphism of the surface that is supported in a twice-punctured disc, which can be visualised as lying in a plane, fixes its boundary, and exchanges the two marked points by rotating them of half a twist around a central pivot. For a precise definition we refer to [FM, §9.1.3]. This mapping class constitutes a square root of the Dehn-twist T_γ around γ .

Similarly, a *one-third Dehn twist* around a curve γ is defined when such curve is 3-separating, and constitutes a cubic root of the twist T_γ . In this case, however, the choice of a curve and an orientation is not enough to provide a unique definition of such twists, hence we will need more information in order to define such a mapping class. What follows is an informal description of one-third Dehn twists.

Since the curve γ bounds a subsurface of type $S_{0,3}^1$ the goal is to define a mapping class which fixes γ and rotates the three marked points by an angle of $\frac{2\pi}{3}$, that is it cyclically permutes them. If the bounded surface is thought as an Euclidean disc centred at the origin with the three marked points being the vertices of an equilateral triangle centred at the origin, this map is easily understood: the boundary curve is fixed, the triangle and its inside is rotated by a $\frac{2\pi}{3}$ angle, and every point in between is rotated of an angle which continuously decreases the closer to the boundary curve.

The issue with the previously highlighted construction is that the identification of a surface of type $S_{0,3}^1$ with an Euclidean disc with three marked points is far from being canonical. It follows that, in stark contrast with the case of a half Dehn-twist, which was determined by a simple choice of the left or right direction, some extra data is needed in order to define a one-third Dehn-twist. In particular, the information which is needed are the three “sides” of the triangle, that is three simple arcs which cyclically connect the marked points and only pairwise intersect in one of the endpoints. Given this, a homeomorphism with the Euclidean disc, which maps the arcs to the edges of the triangle, is now well defined up to isotopy (this fact is not completely trivial and will be proven in Lemma 3.3.2), and the one-third Dehn-twist can be defined by using this identification.

We now formally give the definition of one-third Dehn twists.

Definition 3.3.1. Let $\overline{B(0,2)} \subset \mathbb{C}$ be the closed disc around the origin of the complex plane \mathbb{C} of radius 2, with marked points $\bar{x}_k = e^{\frac{2\pi i}{3}k}$ for $k = 0, 1, 2$, and consider the oriented arcs $\bar{a}_k : [0, 1] \rightarrow \mathbb{C}$ to be $a_k(t) = e^{\frac{2\pi i}{3}(k+t)}$. We define the homeomorphism $R : \overline{B(0,2)} \rightarrow \overline{B(0,2)}$ as

$$R(z) = \begin{cases} e^{\frac{2\pi i}{3}} z & \text{for } \|z\| \leq 1 \\ e^{\frac{2\pi i}{3}(2-\|z\|)} z & \text{for } \|z\| > 1. \end{cases}$$

Let now S be a surface. Let $\gamma \subset S$ be a 3-separating curve. Let X be the closure of the component of $S \setminus \gamma$ of type $S_{0,3}^1$. Let x_0, x_1, x_2 be the three marked points of X , and, for $k = 0, 1, 2$, let a_k be three pairwise disjoint oriented simple parametrised arcs such that a_k has x_k as its first endpoint and x_{k+1} as the second one, where subscripts are taken modulo 3. Up to a homeomorphism we can identify X with $\overline{B(0,2)}$, sending marked points to marked points and arcs to arcs. Any two such homeomorphisms are isotopic due to Lemma 3.3.2. This identification induces a homeomorphism $\overline{R} : X \rightarrow X$ such that $\overline{R}|_{\partial X} = \text{Id}_{\partial X}$. We can now define the *one-third Dehn-twist* around γ along arcs a_0, a_1, a_2 as

$$R_\gamma^{a_0, a_1, a_2} : S \rightarrow S$$

by extending the previously defined map \overline{R} to be the identity on X^c . This homeomorphism is well defined up to isotopy, hence as a mapping class $R_\gamma^{a_0, a_1, a_2} \in \text{Mod}(S)$.

The technical fact we needed in the previous definition is the following.

Lemma 3.3.2. *We will follow the notation we have introduced in Definition 3.3.1. Let $f, g : X \rightarrow \overline{B(0,2)}$ be two homeomorphisms such that, for every $k = 0, 1, 2$, we have $f(x_k) = g(x_k) = \bar{x}_k$ and $f(a_k) = g(a_k) = \bar{a}_k$, preserving the orientation. Then the homeomorphisms $h = g \circ f^{-1} : \overline{B(0,2)} \rightarrow \overline{B(0,2)}$ is isotopic to the identity relatively to the three marked points $\bar{x}_0, \bar{x}_1, \bar{x}_2$ (but not to the boundary).*

Proof. First, we notice that, defined $C = \partial B(0,1)$, we have $h(C) = C$. From this it also follows that $h(B(0,1)) = B(0,1)$, and hence, defining $A = \overline{B(0,2)} \setminus B(0,1)$, we have $h(A) = A$. From the homogeneity of $B(0,1)$ it follows that, up to an isotopy supported in the interior of $B(0,1)$, hence relative to C , and hence to the three marked points, we can assume that $h(0) = 0$. We will use this property later in the proof.

Second, we claim that there exists an isotopy $G : C \times [0,1] \rightarrow C$ such that we have $G(0, \cdot) = h|_C$ and $G(1, \cdot) = \text{Id}_C$, relative to the marked points. Indeed, for every arc \bar{a}_k such an isotopy for the restriction map $h|_{\bar{a}_k}$ relative to the boundary, which is $\{\bar{x}_k, \bar{x}_{k+1}\}$, exists. These isotopies can be glued together to obtain the isotopy G .

We now claim that, up to an isotopy of the entire disc $\overline{B(0,2)}$ relative to the marked points, we can assume that $h|_C = \text{Id}_C$. A homotopy doing the trick is, for instance, the map $F : \overline{B(0,2)} \times [0,1] \rightarrow \overline{B(0,2)}$, defined as follows:

$$F(z, t) = \begin{cases} 0 & \text{for } z = 0 \\ \frac{G(\frac{z}{\|z\|}, t)}{h(\frac{z}{\|z\|})} h(z) & \text{for } z \neq 0 \end{cases}$$

which is well defined as $h(w) \neq 0$ for every $w \neq 0$. Indeed, for $z \in C$ we have that $\frac{z}{\|z\|} = z$ hence, for $t = 0$ we have that $F(0, z) = h(z)$ as $G\left(\frac{z}{\|z\|}, 0\right) = h\left(\frac{z}{\|z\|}\right)$. For $t = 1$, on the other hand, we have $G\left(\frac{z}{\|z\|}, 0\right) = z$, hence $F(z, t) = z$. Since the existence of a homotopy between two homeomorphisms implies the existence of an isotopy, due to general results (see §1.1.3), the claim is proven.

Since we can now assume that $h|_C = \text{Id}_C$, it follows that there exists an isotopy of the disc $\overline{B(0,1)}$, relative to the boundary C , between $h|_{\overline{B(0,1)}}$ and the identity $\text{Id}_{\overline{B(0,1)}}$. Moreover, there exists an isotopy of the annulus A , relative to C (but not to the other

boundary component), between h_A and Id_A . By gluing these isotopies we get a global isotopy between h and $\text{Id}_{\overline{B(0,2)}}$, relative to C , hence to the three marked points. The proof of the lemma is complete. \square

Practically speaking, we are mostly interested in the action these one-third Dehn-twists have on other curves. For a picture of this see, for instance, Figure 3.8. In all of the cases we will be interested in we will study the action of a one-third Dehn-twist around a curve γ only on curves α such that $i(\gamma, \alpha) = 2$. It will be enough to understand how a one-third Dehn-twist acts on a nontrivial arc b with both endpoints on γ . Let us fix an arbitrary orientation on the arc. We will keep on using the notation introduced in Definition 3.3.1. Since the arc is non trivial it will first intersect one of the arcs a_h and then another arc a_l . Up to reversing orientations we can assume that $l = h + 1$ modulo 3. The arc $R_\gamma^{a_1, a_2, a_3}(b)$ will have the same endpoints but will now intersect the arc a_{h+1} first and then the arc a_{h+2} , while “rotating” of an angle $\frac{2\pi}{3}$ outside of the triangle given by the three arcs a_k in the direction provided by the orientation of such arcs.

We remark that

$$(R_\gamma^{a_0, a_1, a_2})^{-1} = R_\gamma^{a_2^*, a_1^*, a_0^*}$$

where a_k^* is the arc a_k with its orientation reversed.

3.4 Exhaustion for $\mathcal{C}_{ss}(S_{0,7})$

In this section we will prove Theorem *E* for the 7-holed sphere $\Sigma_7 = S_{0,7}$.

In all the figures which will appear from now on, we will always draw the 7-holed sphere as the doubling of a heptagon, for instance see Figure 3.1. The vertices of the heptagon in the figure will be the seven marked points of the surface. As it usually happens, a dotted arc in the picture has to be intended as an arc on the “back” of the doubled heptagon, whereas a continuous arc is to be intended as on the front. The only exception to this will be for the curves γ_i in Figure 3.1, which will be the doubling of the arcs depicted in the picture, as we will now explain.

We are now ready to define the rigid graph which will be the starting point of our argument. We first recall that, given a set X , its *disjoint union* with itself is the set $X \sqcup X = (X \times \{0\}) \cup (X \times \{1\})$.

Definition 3.4.1. Let \mathcal{H} be a regular hexagon in \mathbb{R}^2 , with vertices $\{v_1, \dots, v_7\}$ cyclically ordered. Let l_i be the side between v_i and v_{i+1} .

Let $S = \mathcal{H} \sqcup \mathcal{H} / \sim$ be the doubling of the heptagon, where $(x, 0) \sim (y, 1)$ if and only if $x = y \in \partial \mathcal{H}$. The marked surface $(S, \{v_1, \dots, v_7\})$ is a 7-holed sphere.

Let d_i be the segment in \mathcal{H} from the midpoint of l_{3i-3} and the midpoint of l_{3i} , where the subscripts are taken modulo 7. Let γ_i be the curve isotopic to the doubling of d_i .

We will denote the graph in $\mathcal{C}_{ss}(\Sigma_7)$ induced by the curves γ_i as \mathcal{H} .

The curves γ_i are represented in Figure 3.1. It can be noted that the graph \mathcal{H} is an embedded 7-cycle with no diagonals: we will call such a graph a *heptagon*.

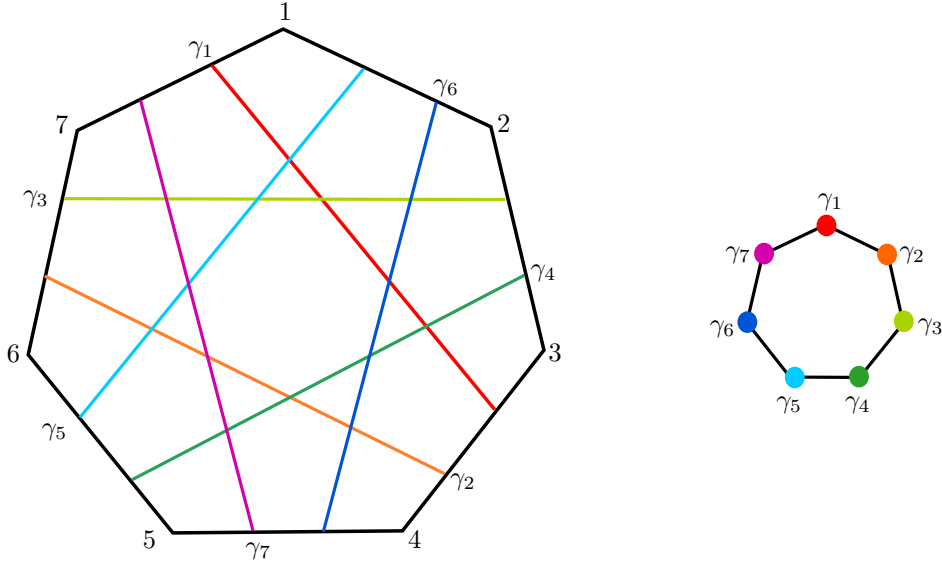


Figure 3.1: The heptagon $\mathcal{H} \subset \mathcal{C}_{ss}(\Sigma_7)$.

The heptagon \mathcal{H} is rigid thanks to [B2, Proposition 3.1].

We will now introduce the mapping classes we will apply Proposition 3.2.4 to. This will be composed of some one-third Dehn twists around the curves γ_i , alongside some half

Dehn twists. In order to define the one-third Dehn twists we need we have to introduce the following notation.

Definition 3.4.2. We will follow the notation we have introduced in Definition 3.4.1. Let $1 \leq i \leq 7$ and $j = 1, 2, 3$. In what follows the indices i are to be considered modulo 7. We define the arcs a_j^i as follows:

- The arc a_1^i is the isotopy class of oriented simple arcs from the marked point v_{3i-2} to v_{3i-1} isotopic to the side l_{3i-2} .
- The arc a_2^i is the isotopy class of oriented simple arcs from the marked point v_{3i-1} to v_{3i} isotopic to the side l_{3i-1} .
- The arc a_3^i is the isotopy class of oriented simple arcs from the marked point v_{3i} to v_{3i-2} isotopic to the segment in the “back” heptagon $\mathcal{H} \times \{1\} \subseteq S$ between the two endpoints.

We define the following one-third Dehn twists, as introduced in Definition 3.3.1,

$$R_i = R_{\gamma_i}^{a_1^i, a_2^i, a_3^i} \in \text{Mod}(\Sigma_7).$$

The oriented arcs a_j^i are represented in Figure 3.2. These arcs are not doubled. The arcs a_3^i , which are dotted, are intended to be on the back side of the sphere. The indices of these arcs a_j^i , which we will only use to define the one-third Dehn twists R_i , should be interpreted as follows: the index $i = 1, \dots, 7$ indicates the curve γ_i around which the twist is made, while the index $j = 1, 2, 3$ refers to the cyclic ordering of the three arcs.

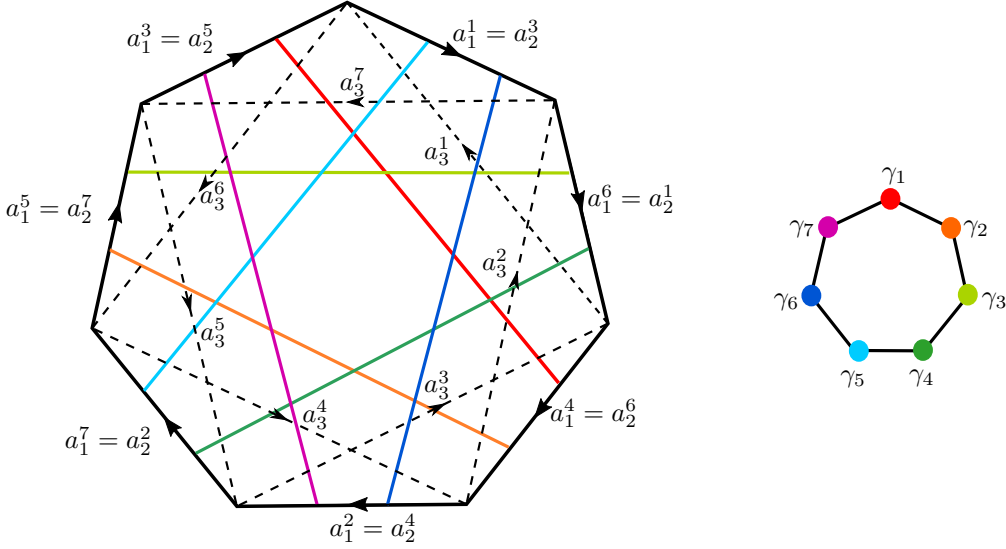


Figure 3.2: The heptagon $\mathcal{H} \subset \mathcal{C}_{ss}(\Sigma_7)$ with the arcs a_j^i .

We will now introduce a collection of 2-separating curves, around which the half-Dehn twists we need will be.

Definition 3.4.3. Let α_i be the unique essential simple closed curve on Σ_7 which is disjoint from every curve in \mathcal{H} apart from γ_{i+3} and γ_{i-3} , where subscripts are taken modulo 8.

We define mapping classes $H_i \in \text{Mod}(\Sigma_7)$ to be the right half-Dehn twists around the curve α_i .

The curves α_i are pictured in Figure 3.3. It is immediate to check that these curves are 2-separating, so the half Dehn-twists around them are defined.

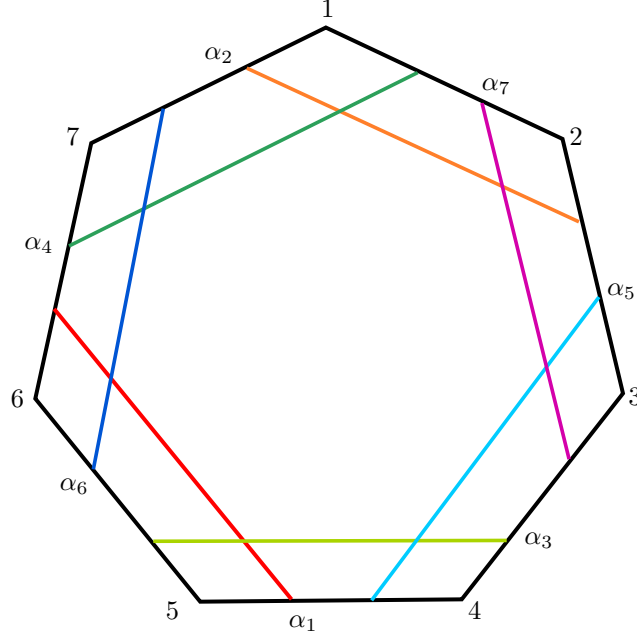


Figure 3.3: The curves α_i .

Throughout both this and the following section we will often encounter statements which are more neatly expressed by condensing various cases which arise from the choice of a sign. In order to avoid ambiguity, in those statement we will use $\varepsilon \in \{1, -1\}$ instead of the more common notation \pm .

We are now ready to define the finite sets of mapping classes which we will need in order to apply Proposition 3.2.4.

Definition 3.4.4. We define

$$G = \{\text{Id}_{\Sigma_7}\} \cup \{R_i^\varepsilon\}_{i=1,\dots,7}^{\varepsilon=\pm 1} \subset \text{Mod}(\Sigma_7).$$

Moreover we define

$$\overline{G} = G \cup \{H_i^\varepsilon\}_{i=1,\dots,7}^{\varepsilon=\pm 1} \subset \text{Mod}(\Sigma_7).$$

The discussion in [FM, § 4.4.4] proves that the collection $\{H_i^\varepsilon\}$, hence \overline{G} , generates the entire mapping class group $\text{Mod}(\Sigma_7)$.

Our final goal is to apply Proposition 3.2.4 with $\Gamma = \mathcal{C}_{ss}(\Sigma_7)$, $H = \overline{G}$ and $X = \overline{G} \cdot \mathcal{H}$. The hardest hypothesis to check will be the rigidity of $\overline{G} \cdot \mathcal{H}$, which is the following statement, whose proof will occupy most of the remainder of the section.

Proposition 3.4.5. *Let \mathcal{H} be the heptagon of $\mathcal{C}_{ss}(\Sigma_7)$ defined in Definition 3.4.1. Let \overline{G} be as in Definition 3.4.4. Let $i : \overline{G} \cdot \mathcal{H} \hookrightarrow \mathcal{C}_{ss}(\Sigma_7)$ be an injective graph homomorphism. Then there exists an extended mapping class $f \in \text{Mod}^\pm(\Sigma_7)$ such that $i = f$.*

In particular, the graph $\overline{G} \cdot \mathcal{H}$ is rigid.

In the previous statement we have willingly highlighted the fact that any embedding of the graph the graph $\overline{G} \cdot \mathcal{H}$ is induced by an extended mapping class that is, in the language of Proposition 3.2.6, $\overline{G} \cdot \mathcal{H}$ is $\text{Mod}^\pm(\Sigma_m)$ -rigid. This stronger property obviously follows from rigidity, once the combinatorial rigidity for $\mathcal{C}_{ss}(\Sigma_7)$, that is Theorem 1.3.4 or Theorem A, is assumed. However, in the proof of the previous statement, we will never have to explicitly use such rigidity result. Hence our proof, alongside Proposition 3.2.6 constitutes another independent proof that $\text{Aut}(\mathcal{C}_{ss}(\Sigma_7)) \cong \text{Mod}^\pm(\Sigma_7)$.

We will start by showing that $\overline{G} \cdot \mathcal{H} = G \cdot \mathcal{H}$, where the collections G, \overline{G} are as defined in Definition 3.4.4. This means that the curves created by the action of the half Dehn-twists H_i are already included in the ones arising from the one-third Dehn-twists T_i^ε , hence it is enough to focus on the action of the latter. The reason why we still have to use the set of mapping classes \overline{G} is the fact that G alone does not generate the entire mapping class group, as it only induces even transpositions on the marked points; it is possible that it still acts transitively on strongly separating curves, but this is currently unknown, and seems a hard problem unlikely to be solved shortly. In order to prove the aforementioned equality we will now prove the following relations between half Dehn-twists and one-third Dehn-twists.

Lemma 3.4.6. *Let the curves γ_i and the mapping classes R_i and H_i be defined as before. Then the following relations hold:*

1. $H_i \gamma_{i+3} = R_{i+3}^{-1} \gamma_{i+1};$
2. $H_i \gamma_{i-3} = R_{i-1} \gamma_{i-3};$
3. $H_i^{-1} \gamma_{i+3} = R_{i+1}^{-1} \gamma_{i+3};$
4. $H_i^{-1} \gamma_{i-3} = R_{i-3} \gamma_{i-1};$

where the subscripts are intended to be modulo 7.

Proof. Equalities 1, 2, 3, 4 are proven in Figures 3.4, 3.5, 3.6, 3.7 respectively. \square

Corollary 3.4.7. *We have $\overline{G} \cdot \mathcal{H} = G \cdot \mathcal{H}$.*

Proof. By definition every curve α_i is disjoint from every curve in \mathcal{H} apart from γ_{i+3} and γ_{i-3} , hence for $j \neq i \pm 3$ it follows that $H_i^\varepsilon \gamma_j = \gamma_j = R_j \gamma_j \in G \cdot \mathcal{H}$. From this observation combined with Lemma 3.4.6 it follows that for every i we have $\alpha_i \cdot \mathcal{H} \subseteq G \cdot \mathcal{H}$, hence $\overline{G} \cdot \mathcal{H} \subseteq G \cdot \mathcal{H}$. The other inclusion is trivial. \square

The next step in the proof of rigidity is to find some relations between different one-third Dehn-twists. Indeed, the graph $G \cdot \mathcal{H}$ is the union of the heptagons $R_i^\varepsilon \cdot \mathcal{H}$, which are glued together according to these relations, which will be the key property to prove rigidity. Without further ado we will now prove the following.

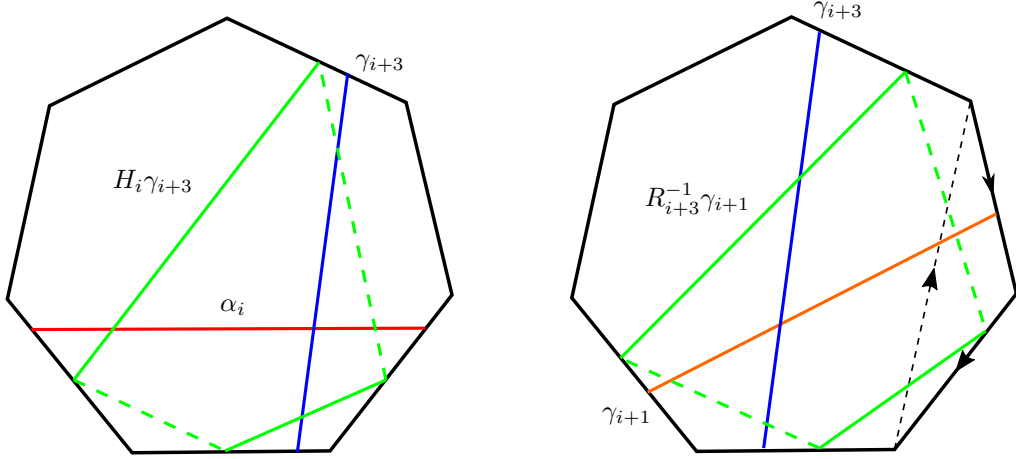


Figure 3.4: Equality 1 of Lemma 3.4.6.

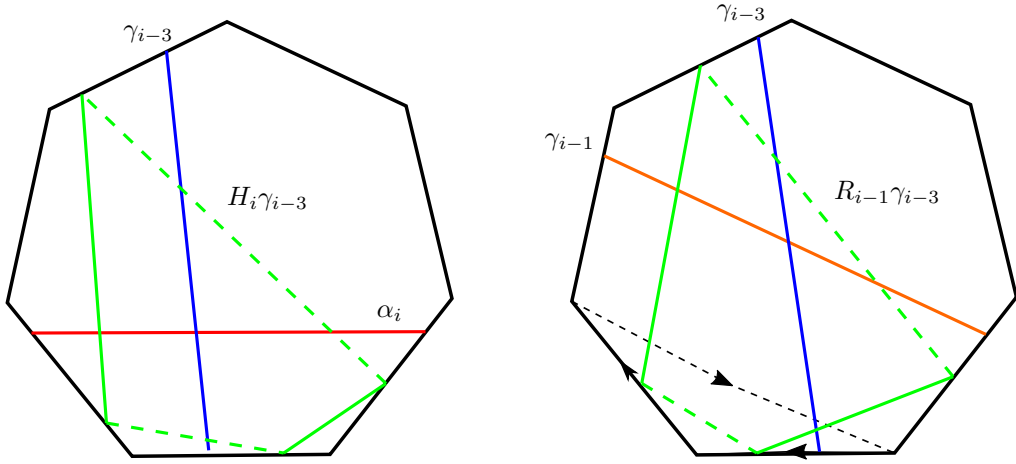


Figure 3.5: Equality 2 of Lemma 3.4.6.

Lemma 3.4.8. *Let the curves γ_i and the mapping classes R_i and H_i be defined as before. The following relations hold:*

1. $R_i \gamma_{i+4} = R_{i+4}^{-1} \gamma_{i+2};$
2. $R_i \gamma_{i+2} = R_{i+4}^{-1} \gamma_i;$
3. $R_i^{-1} \gamma_{i-4} = R_{i-4} \gamma_{i-2};$
4. $R_i^{-1} \gamma_{i-2} = R_{i-4} \gamma_i;$
5. $R_i \gamma_{i+3} = R_{i-3}^{-1} \gamma_{i+1};$

where subscripts are taken modulo 7.

Proof. Equalities 1, 2, 3, 4, 5 are pictured in Figures 3.8, 3.9, 3.10, 3.11, 3.12 respectively. The color scheme employed is chosen to be consistent with curves in Figure 3.1 with $i = 1$. \square

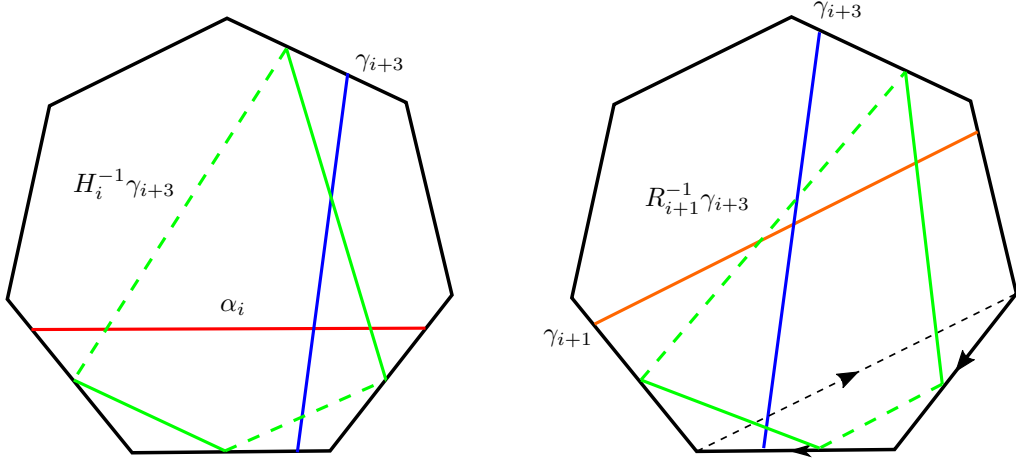


Figure 3.6: Equality 3 of Lemma 3.4.6.

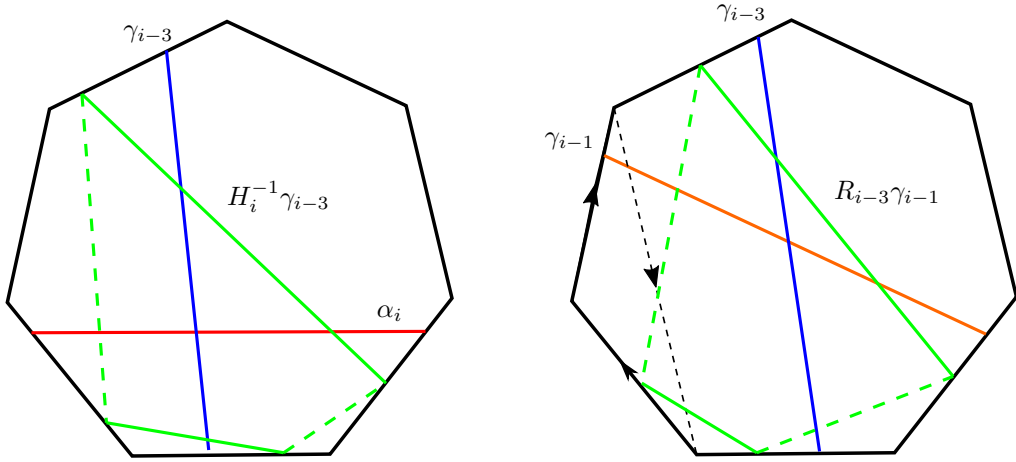


Figure 3.7: Equality 4 of Lemma 3.4.6.

Alongside the relations of the previous lemma we recall that if γ_i and γ_j are disjoint (possibly equal) then $R_i \gamma_j = \gamma_j$. This first implies that every heptagon $R_i^\varepsilon \cdot \mathcal{H}$ intersects \mathcal{H} in three consecutive vertices.

We will approach the rigidity of $G \cdot \mathcal{H}$ by studying the rigidity of some of its subgraphs, at first. In particular the subgraphs we will study will be the union of \mathcal{H} with two other heptagons obtained by applying some one-third Dehn-twist. Indeed, we define the following.

Definition 3.4.9. Let $\varepsilon \in \{+1, -1\}$. For $j = 1, \dots, 7$ let

$$Z_j^\varepsilon = \mathcal{H} \cup (R_j^\varepsilon \cdot \mathcal{H}) \cup (R_{j+4}^{-\varepsilon} \cdot \mathcal{H}),$$

where subscripts are taken modulo 7.

A picture of the graph Z_1^+ is provided in Figure 3.13.

Before we move forward with our argument we need to introduce a particular extended

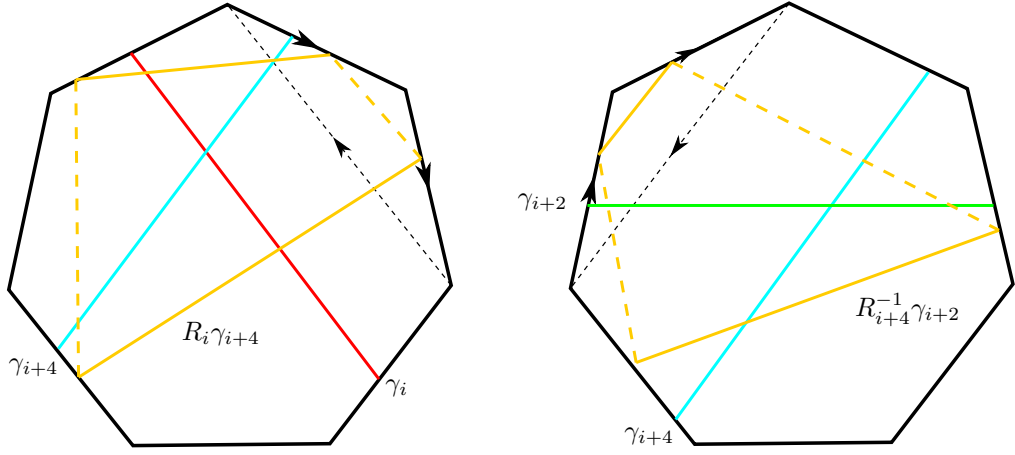


Figure 3.8: Equality 1 of Lemma 3.4.8.

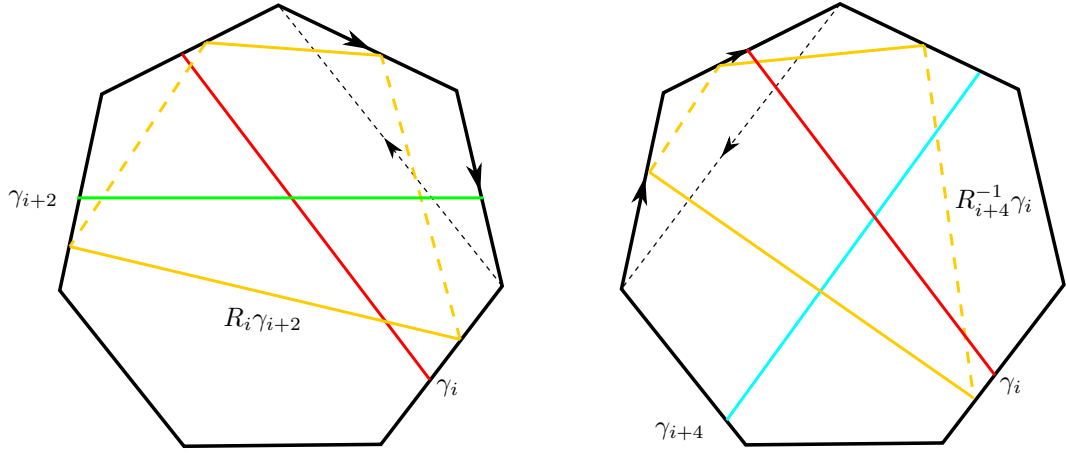


Figure 3.9: Equality 2 of Lemma 3.4.8.

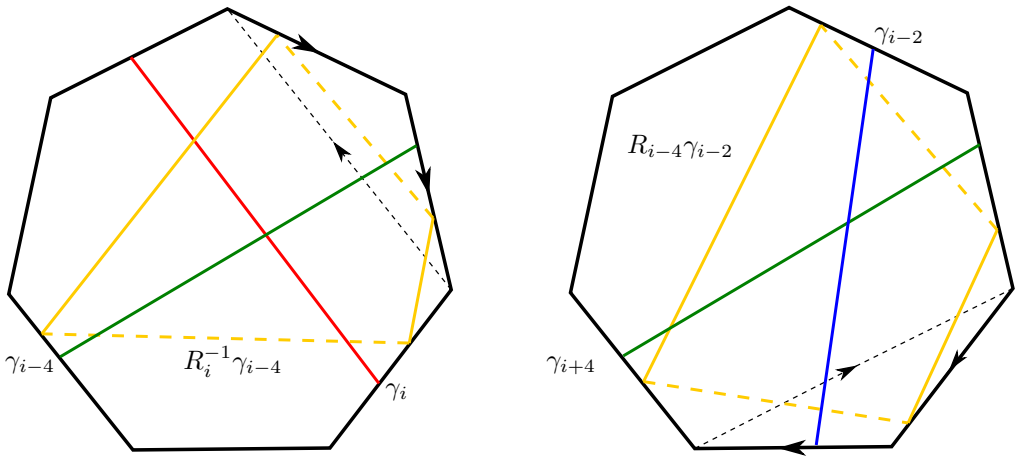


Figure 3.10: Equality 2 of Lemma 3.4.8.

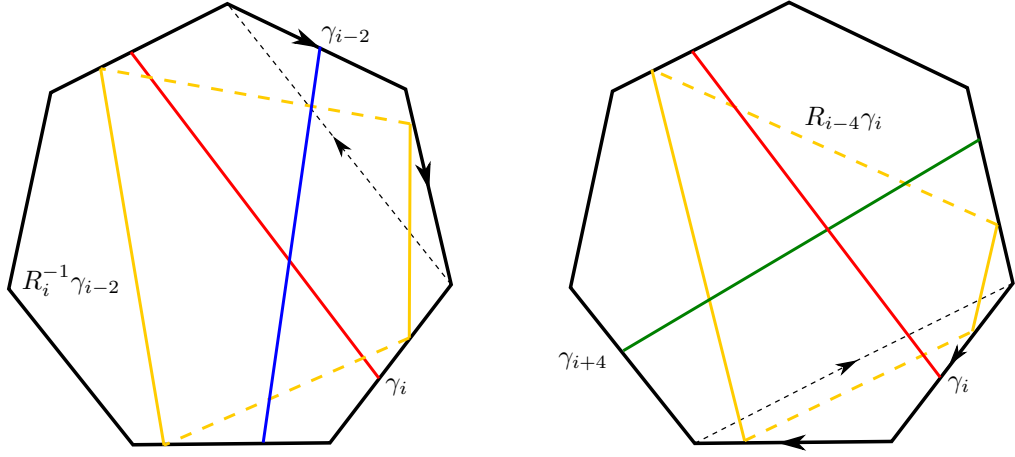


Figure 3.11: Equality 4 of Lemma 3.4.8.

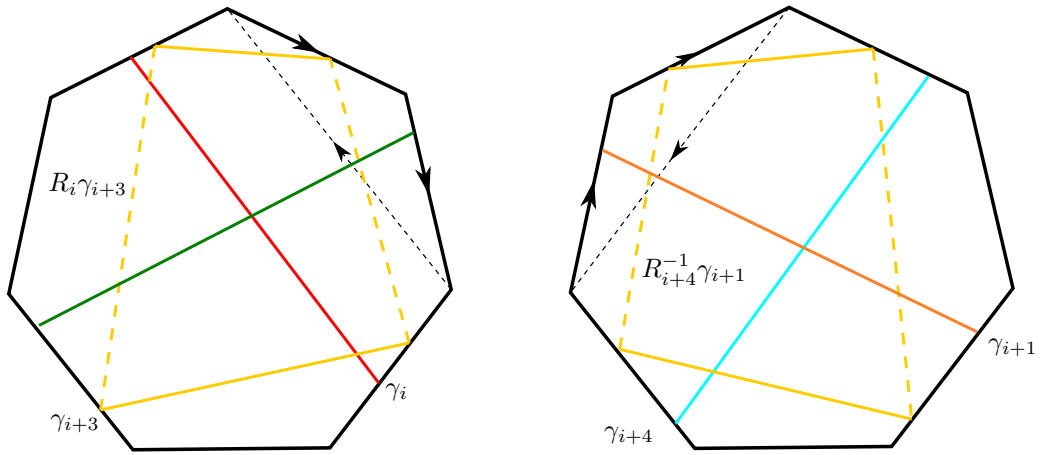


Figure 3.12: Equality 5 of Lemma 3.4.8.

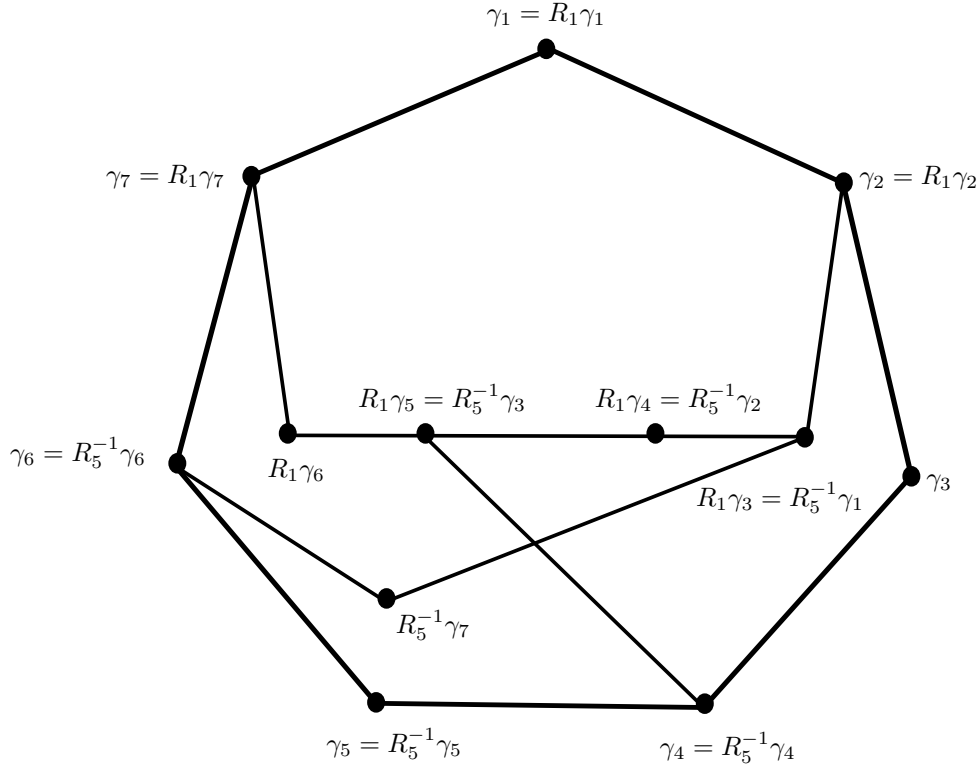


Figure 3.13: The graph Z_1^+ .

mapping class: the *mirror reflection*.

Definition 3.4.10. We follow the notation introduced in Definition 3.4.1. Let $r : S \rightarrow S$ the homeomorphism defined by $r(x, 0) = (x, 1)$ and $r(y, 1) = (y, 0)$. We will denote the isotopy class of r as the *mirror reflection* $\rho \in \text{Mod}^\pm(\Sigma_7)$.

The mirror reflection we have just defined exchanges the “front” and the “back” of the doubled heptagon we are identifying with Σ_7 . This is an orientation-reversing automorphism of order two. Moreover, we notice that the mirror reflection ρ fixes all the curves γ_i .

We observe the following, which is the computation of the stabiliser of the heptagon \mathcal{H} under the action of the extended mapping class group.

Lemma 3.4.11. *Let \mathcal{H} be the heptagon as in Definition 3.4.1. Let $f \in \text{Mod}^\pm(\Sigma_7)$ such that $f|_{\mathcal{H}} = \text{Id}_{\mathcal{H}}$. Then either $f = \text{Id}$, if it preserves the orientation, or $f = \rho$, if it is orientation-reversing.*

Proof. Up to composing f with ρ we can assume f to be orientation-preserving. From now on, we will abuse notation and confuse the mapping class f with one of its representatives. The collection of curves $\gamma_1, \dots, \gamma_7$ satisfies the hypotheses of Proposition 1.2.6.

We claim that f preserves the orientation of every curve γ_i , hence we can apply Proposition 1.2.6 and deduce that f is isotopic to the identity.

The claim follows from the fact that if a surface homeomorphism fixes a curve but changes its orientation, then it must exchange the two connected components of the

complementary of the curve. In particular, for this to be possible the two complementary discs of the curve need to be of the same topological type. However, every curve γ_i bounds a 3-punctured disc on one side and a 4-punctured disc on the other, so f must preserve its orientation. The claim is proven and the proof is complete. \square

We observe the following corollary, which will be useful later to check the hypothesis of Proposition 3.2.4 regarding triviality of the pointwise stabiliser.

Corollary 3.4.12. *Let i, j be such that $j \neq i - 1, i, i + 1$ modulo 7. Then the subgraph $Y = \mathcal{H} \cup \{R_i \gamma_j\}$ has trivial pointwise stabiliser under the action of the extended mapping class group $\text{Mod}^\pm(\Sigma_7)$. In particular Y has trivial pointwise stabiliser for the action of $\text{Aut}(\mathcal{C}_{ss}(\Sigma_7))$.*

Proof. Let $f \in \text{Mod}^\pm(\Sigma_7)$ be an extended mapping class such that $f|_Y = \text{Id}_Y$. Then, since $f|_{\mathcal{H}} = \text{Id}_{\mathcal{H}}$ it follows from Lemma 3.4.11 that either $f = \text{Id}$ or $f = \rho$. Since $\rho(R_i \gamma_j) \neq R_i \gamma_j$ it follows that $f = \text{Id}$. Triviality of the stabiliser follows from the fact that $\text{Aut}(\mathcal{C}_{ss}(\Sigma_7)) \cong \text{Mod}^\pm(\Sigma_7)$, that is Theorem 1.3.4. \square

We remark that the first part of the previous proposition is exactly the property needed in order to apply Proposition 3.2.6, and deduce rigidity from there.

We will from now on turn our attention towards the proof the following, which, combined with [B2, Proposition 3.1], proves rigidity for the graphs $Z_{(i)}^\varepsilon$.

Proposition 3.4.13. *Let $i : Z_j^\varepsilon \hookrightarrow \mathcal{C}_{ss}(\Sigma_7)$ be an injective graph homomorphism such that $i|_{\mathcal{H}} = \text{Id}_{\mathcal{H}}$. Then either $i = \text{Id}_{Z_j^\varepsilon}$ or $i = \rho|_{Z_j^\varepsilon}$.*

From now on, in order to make the notation less cumbersome, during the proofs we will only deal with the graph Z_1^+ , the other cases being completely analogous. The first step to prove the previous rigidity result is to study the combinatorics of the marked points enclosed by curves in the graph. Let

$$\pi : \mathcal{C}_{ss}(\Sigma_7) \longrightarrow K(7, 3)$$

be the projection of the strongly separating curve graph onto the Kneser graph $K(7, 3)$, as defined in §1.5.1. We recall that, given a strongly separating curve $\gamma \subset \Gamma_7$ its projection $\pi(\gamma)$ is the set of three punctures contained in the 3-puncture complementary disc of γ .

The action of a one-third Dehn-twist the Kneser graph is easily understood: indeed, such a mapping class cyclically permutes the three punctures x_1, x_2, x_3 as in Definition 3.3.1, hence it acts on the Kneser graph as the 3-cycle $(x_1 x_2 x_3)$. For example the mapping class T_1 acts on the Kneser graph $K(7, 3)$ as the permutation (123) .

A heptagon in the Kneser graph is rigid thanks to [B2, Proposition 3.3]. We will now prove a rigidity result, similar in spirit to Proposition 3.4.13, for the projection of the graphs Z_i^ε .

Lemma 3.4.14. *Let $i : Z_j^\varepsilon \hookrightarrow \mathcal{C}_{ss}(\Sigma_7)$ be an injective graph homomorphism such that we have $i|_{\mathcal{H}} = \text{Id}_{\mathcal{H}}$. Then $\pi \circ i|_{Z_j^\varepsilon} = \pi|_{Z_j^\varepsilon}$.*

Proof. Thanks to [B2, Proposition 3.3] we can assume, up to a relabelling of punctures, that the heptagon \mathcal{H} has the following projection:

- $\pi(\gamma_1) = \{1, 2, 3\};$
- $\pi(\gamma_2) = \{4, 5, 6\};$
- $\pi(\gamma_3) = \{7, 1, 2\};$
- $\pi(\gamma_4) = \{3, 4, 5\};$
- $\pi(\gamma_5) = \{6, 7, 1\};$
- $\pi(\gamma_6) = \{2, 3, 4\};$
- $\pi(\gamma_7) = \{5, 6, 7\}.$

It is worth noting that, while every labelling is the same up to a permutation, the one we have chosen differs from the one employed in [B2]. Moreover, we observe that [B2, Lemma 3.2] proves that two 2-distant vertices in $K(7, 3)$ have intersection exactly 2. We notice that the fact that $i|_{\mathcal{H}} = \text{Id}_{\mathcal{H}}$ implies that $\pi(i(\gamma_j)) = \pi(\gamma_j)$ for every j .

From this, keeping in mind the structure of Z_1^+ given in Figure 3.13, which is the same graph structure as $i(Z_1^+)$, it follows that $\pi(i(R_1\gamma_6))$ has intersection 2 with the vertices $\pi(i(\gamma_1)) = \{1, 2, 3\}$, $\pi(i(\gamma_4)) = \{3, 4, 5\}$ and $\pi(i(\gamma_6)) = \{2, 3, 4\}$. Since $\pi(i(R_1\gamma_6))$ is also disjoint from $\pi(i(\gamma_7)) = \{5, 6, 7\}$ it follows that

$$\pi(i(R_1\gamma_6)) = \{1, 3, 4\} = (123) \cdot \{2, 3, 4\} = \pi(R_1\gamma_6).$$

Since $\pi(i(R_1\gamma_5))$ is disjoint from both $\pi(i(R_1\gamma_6)) = \{1, 3, 4\}$ and $\pi(i(\gamma_4)) = \{3, 4, 5\}$ it follows that

$$\pi(i(R_1\gamma_5)) = \{2, 6, 7\} = (123) \cdot \{6, 7, 1\} = \pi(R_1\gamma_5).$$

Similarly $\pi(i(R_1\gamma_4))$ has intersection 2 with $\pi(i(\gamma_4)) = \{3, 4, 5\}$, $\pi(i(R_1\gamma_6)) = \{1, 3, 4\}$ and $\pi(i(\gamma_2)) = \{4, 5, 6\}$, while being disjoint from $\pi(i(R_1\gamma_5)) = \{2, 6, 7\}$. It follows that

$$\pi(i(R_1\gamma_4)) = \{1, 4, 5\} = (123) \cdot \{3, 4, 5\} = \pi(R_1\gamma_4).$$

Since $\pi(i(R_1\gamma_3))$ is disjoint from both $\pi(i(R_1\gamma_4)) = \{1, 4, 5\}$ and $\pi(i(\gamma_2)) = \{4, 5, 6\}$ it follows that

$$\pi(i(R_1\gamma_3)) = \{2, 3, 7\} = (123) \cdot \{7, 1, 2\} = \pi(R_1\gamma_3).$$

In the same vein $\pi(i(R_5^{-1}\gamma_7))$ is disjoint from both $\pi(i(R_1\gamma_3)) = \{2, 3, 7\}$ and $\pi(i(\gamma_6)) = \{2, 3, 4\}$, hence it follows that

$$\pi(i(R_1\gamma_7)) = \{1, 5, 6\} = (123) \cdot \{5, 6, 7\} = \pi(R_1\gamma_7).$$

This concludes the proof. \square

Now that we have dealt with the combinatorics it is time to employ topological arguments to prove the rigidity property of the graphs $Z_{(i)}^\varepsilon$. As a preliminary step to this we observe that rigidity of \mathcal{H} implies that, given any two non-adjacent curves β, β'

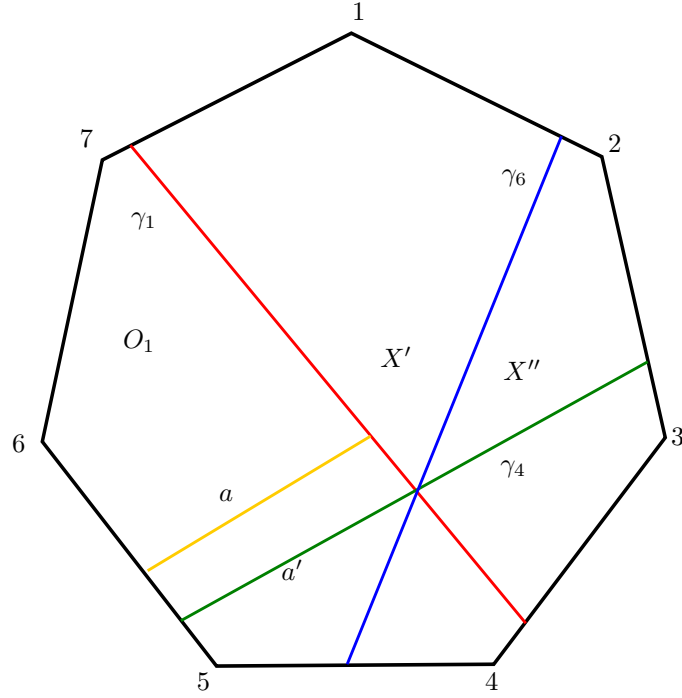


Figure 3.14: The curves are doubling of the pictured arcs.

in a heptagon, we have $i(\beta, \beta') = 2$. Moreover, the following uniqueness of neighbour property, analogous to Lemma 2.5.3, also holds for $\mathcal{C}_{ss}(\Sigma_7)$.

Lemma 3.4.15. *Let α, β be two vertices of $\mathcal{C}_{ss}(\Sigma_7)$. Then there exists at most one vertex $\gamma \in \mathcal{C}_{ss}(\Sigma_7)$ which is adjacent to both α and β .*

Proof. Let γ, γ' be two 3-separating curves disjoint from α and β . Let X_α (resp. $X_\beta, X_\gamma, X_{\gamma'}$) be the unique complementary disc of α (resp. β, γ, γ') containing exactly 3 marked points. Each pair of curves is disjoint if and only if their 3-punctured complementary discs are disjoint. Lemma 2.5.3, applied with $A = X_\alpha$ and $B = X_\beta$ implies that $X_\gamma = X_{\gamma'}$ hence $\gamma = \partial X_\gamma = \partial X_{\gamma'} = \gamma'$, and we are done. \square

Thanks to these observations we can now move to the proof of Proposition 3.4.13.

Proof of Proposition 3.4.13. Let $\beta = i(R_1\gamma_5)$. We first claim that either $\beta = R_1\gamma_5$ or $\beta = \rho(R_1\gamma_5)$. For every j let B_j be the 3-punctured complementary disc of γ_j and let O_j be the 4-punctured one. For a quick way to visualise the relations used throughout the proof we refer to Figure 3.13.

Since we have observed that $i(\beta, \gamma_1) = i(\gamma_4, \gamma_1) = 2$ both $a = \beta \cap O_1$ and $a' = \gamma_4 \cap O_1$ are single arcs with both endpoints on γ_1 . Moreover, Equality 1 in Lemma 3.4.8, applied with $i = 1$, proves that $R_1\gamma_5 = R_5^{-1}\gamma_3$, hence it is adjacent to $R_5^{-1}\gamma_4 = \gamma_4$. It follows that $R_1\gamma_5$ and γ_4 are disjoint, so in particular the arcs a and a' are. Thanks to Lemma 3.4.14 we know that one component of $O_1 \setminus a$ contains the marked points $\{4, 5\}$, while the other contains $\{6, 7\}$, and the same hold for a' . It follows that the two arcs a, a' are isotopic. This construction is shown in Figure 3.14.

From the previous argument it follows that the isotopy type of the arc $b = \beta \cap B_1$ determines the entire curve β . Since b is disjoint from γ_4 it then lies in the closure of one of the two components of $B_1 \setminus \gamma_4$, one of which is a once-punctured disc, the other being a 2-punctured one. If b were contained in the once-punctured disc, then b would either form a bigon with γ_1 , which would violate minimal position, or be isotopic to $\gamma_4 \cap B_1$. In the latter case we would have $\beta = \gamma_4 = i(\gamma_4)$ which is impossible, as $R_1\gamma_5$ and γ_4 are adjacent, hence distinct, vertices in the heptagon $R_1 \cdot \mathcal{H}$, as shown in Figure 3.13, and the map i is injective.

It follows that the arc b must be contained in the region $X = B_1 \cap O_4$, which is a 2-punctured disc containing the marked points $\{1, 2\}$. The curve γ_6 , which intersects both γ_1 and γ_6 exactly twice and separates the marked point 1 from 2, cuts X into two once-punctured discs X', X'' . Without loss of generality we can assume the curves to be arranged in a way such that the curves are still in minimal position and transverse, but there are two triple points, that is the triple intersection $\gamma_1 \cap \gamma_4 \cap \gamma_6$ contains exactly two points, as in Figure 3.14. Let X' be the once-punctured disc such that $\partial X' \subset \gamma_1 \cup \gamma_6$. Let X'' be the component such that $\partial X'' \subset \gamma_4 \cup \gamma_6$. We have $\partial X \cap \partial X'' = \gamma_6 \cap B_1$.

Since $\beta = R_1\gamma_5 = R_5^{-1}\gamma_3$, due to Equality 1 of Lemma 3.4.8 applied with $i = 1$, the curve β is nonadjacent to γ_6 in $R_5^{-1} \cdot \mathcal{H}$, hence $i(\beta, \gamma_6) = 2$ thanks to a previous observation. It follows that $b \cap X''$ is a single nontrivial arc with both endpoints on γ_6 . Since there exists only one nontrivial arc in X'' with both endpoints on γ_6 , the arc $b \cap X''$ is uniquely determined. The component of $X'' \setminus b$ which does not intersect γ_4 contains the marked point 2.

It now follows that $b \cap X'$ is composed of two arcs, each one with an endpoint on γ_1 and the other on γ_6 : we claim these arcs are isotopic. If not, the region of $X' \setminus b$ whose boundary contains the two arcs would contain the marked point 1, hence the two marked points 1, 2 would both belong to the same component of $\Sigma_7 \setminus \beta$, which is impossible, since Lemma 3.4.14 implies that

$$\pi(\beta) = \pi(R_1\gamma_5) = (123) \cdot \{6, 7, 1\} = \{6, 7, 2\}.$$

It follows that there are only two possibilities for b , and hence for β : those possibilities are pictured in Figure 3.15, and it can be noticed that those are exactly the curves $R_1\gamma_5$ or $\rho(R_1\gamma_5)$, as can be seen by comparison with Figure 3.16. The claim is proven.

Up to postcomposing with the mirror image ρ we can now assume $i(R_1\gamma_5) = R_1\gamma_5$. A completely analogous argument can now be applied to $R_5^{-1}\gamma_2 = R_1\gamma_4$, which is represented in Figure 3.17, proving that either $i(R_1\gamma_4) = R_1\gamma_4$ or $i(R_1\gamma_4) = \rho(R_1\gamma_4)$. However, the latter case is impossible since $R_1\gamma_5$ and $\rho(R_1\gamma_4)$ are not disjoint, as shown in Figure 3.18.

We have shown that if $i(R_1\gamma_5) = R_1\gamma_5$ then i also fixes $R_1\gamma_4$. Since $i(\gamma_7) = \gamma_7$ Lemma 3.4.15, applied with $\alpha = R_1\gamma_5$ and $\beta = \gamma_7$ proves that $i(R_1\gamma_6) = R_1\gamma_6$. Since $i(\gamma_2) = \gamma_2$ Lemma 3.4.15, applied with $\alpha = R_1\gamma_4$ and $\beta = \gamma_2$ proves that we have $i(R_1\gamma_3) = R_1\gamma_3$. Equality 2 in Lemma 3.4.8 proves that $R_1\gamma_3 = R_5^{-1}\gamma_1$. Another application of Lemma 3.4.15, applied with $\alpha = R_5^{-1}\gamma_1$ and $\beta = \gamma_6$ proves that we have $i(R_5^{-1}\gamma_7) = R_5^{-1}\gamma_7$. This proves that $i = \text{Id}_{Z_1^+}$.

If $i(R_1\gamma_5) = \rho(R_1\gamma_5)$, instead, composition with ρ proves that $i = \rho|_{Z_1^+}$, and the

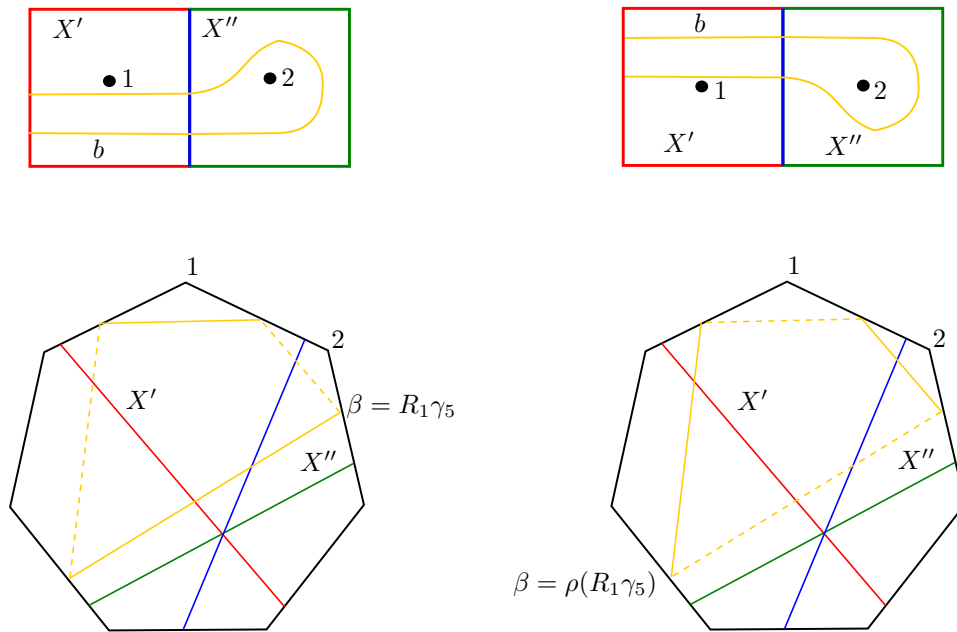


Figure 3.15

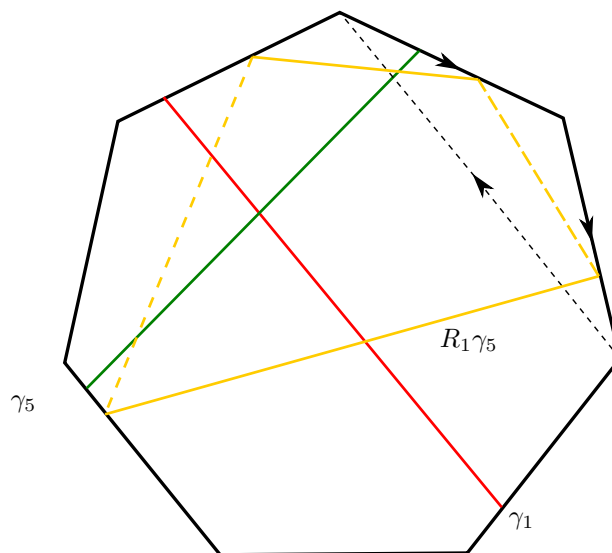


Figure 3.16: The curve in yellow is $R_1\gamma_5$.

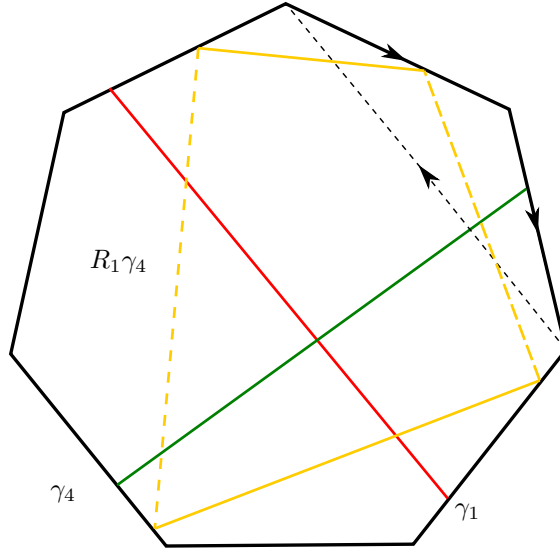


Figure 3.17: The curve in yellow is $R_1\gamma_4$.

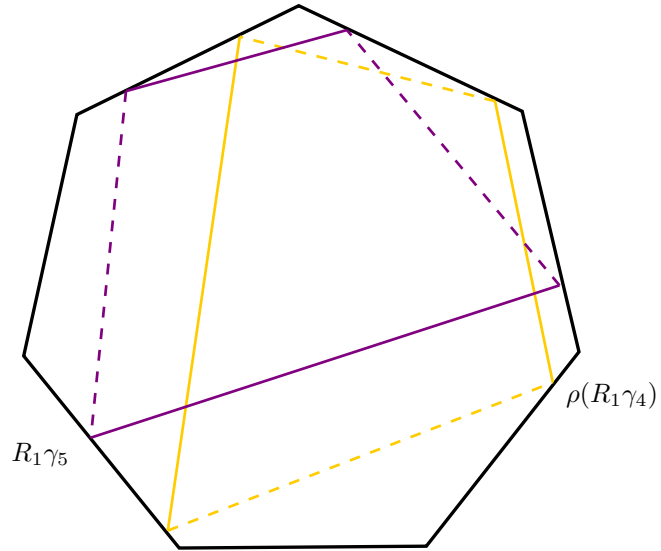


Figure 3.18: The purple curve is $R_1\gamma_5$ while the yellow curve is $\rho(R_1\gamma_4)$.

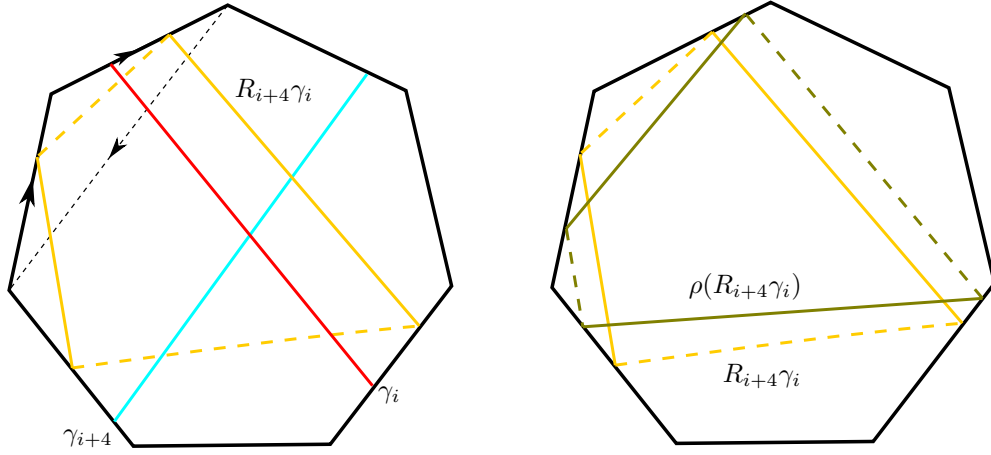


Figure 3.19: The yellow curve is $R_{j+4}\gamma_i$ while the olive one is $\rho(R_{j+4}\gamma_i)$.

proposition is proven. \square

We are now almost ready to prove rigidity of $G \cdot \mathcal{H}$: first we prove the following, which is a “global version” of Proposition 3.4.13.

Lemma 3.4.16. *Let $i : G \cdot \mathcal{H} \hookrightarrow \mathcal{C}_{ss}(\Sigma_7)$ be an injective graph homomorphism such that $i|_{\mathcal{H}} = \text{Id}_{G \cdot \mathcal{H}}$. Then either $i = \text{Id}_Z$ or $i = \rho|_{G \cdot \mathcal{H}}$.*

Proof. Proposition 3.4.13 proves that for every j the injection i restricts to either the identity or the mirror image on Z_j^ε . We claim that these restrictions are consistently either always the identity or the mirror image. In order to prove this we argue by contradiction: if our claim were not true there would exist an index j such that $i|_{Z_j^\varepsilon} = \text{Id}_{Z_j^\varepsilon}$ and $i|_{Z_{j+4}^{-\varepsilon}} = \rho|_{Z_{j+4}^{-\varepsilon}}$. This follows from the fact that the sequence $Z_{j_n}^{\varepsilon_n}$ where $i_{n+1} = i_n + 4$ modulo 7 and $\varepsilon_0 = 1$ and $\varepsilon_{n+1} = -\varepsilon_n$ exhausts all the possible subgraphs Z_j^ε .

Let j be as above. We notice that $R_{j+4}^{-\varepsilon}\gamma_j$ belongs to both Z_j^ε and $Z_{j+4}^{-\varepsilon}$. It follows that

$$R_{j+4}^{-\varepsilon}\gamma_j = i(R_{j+4}^{-\varepsilon}\gamma_j) = \rho(R_{j+4}^{-\varepsilon}\gamma_j).$$

Since $\rho(R_{j+4}^{-\varepsilon}\gamma_i) \neq R_{j+4}^{-\varepsilon}\gamma_j$, as proven in Figure 3.19 (for the $+$ case), we get a contradiction and the lemma is proven. \square

We are now ready to prove the rigidity of $G \cdot \mathcal{H}$, that is Proposition 3.4.5.

Proof of Proposition 3.4.5. Let $i : G \cdot \mathcal{H} \hookrightarrow \mathcal{C}_{ss}(\Sigma_7)$ be an injective graph homomorphism. Then, thanks to [B2, Proposition 3.1], up to postcomposing with an orientation preserving mapping class $\varphi \in \text{Mod}(\Sigma_7)$ we can assume $\varphi \circ i|_{\mathcal{H}} = \text{Id}_{\mathcal{H}}$. It now follows from Lemma 3.4.16 that either $i = \varphi|_{G \cdot \mathcal{H}}^{-1}$ or $i = \rho \circ \varphi|_{G \cdot \mathcal{H}}^{-1}$ and both maps are induced by extended mapping classes. \square

We can now conclude the proof of Theorem *E* for the 7-holed sphere.

Proof of Theorem *E* for Σ_7 . We can now apply Proposition 3.2.4 with $\Gamma = \mathcal{C}_{ss}(\Sigma_7)$, $H = \overline{G}$, and $X = \overline{G} \cdot \mathcal{H}$.

Rigidity of X is Proposition 3.4.5. The triviality of stabilisers follows from the fact that every set subgraph of the form $(g \cdot X) \cap X$ contains \mathcal{H} and another curve of the form $T_i \gamma_j$ for $j \neq i - 1, i, i + 1$, hence it is a superset of a subgraph which has trivial pointwise stabiliser due to Corollary 3.4.12, hence its stabiliser is also trivial thanks to Lemma 3.2.1.

Lastly, $\text{Mod}(\Sigma_7) = \langle H \rangle$ acts transitively on strongly separating curves in Σ_7 , since they are all of the same topological type, hence there is only one orbit of vertices, so the requirement to have an exhaustion is fulfilled. \square

3.5 Exhaustion for $\mathcal{C}_{ss}(S_{0,8})$

In this section we will prove Theorem *E* for the 8-holed sphere Σ_8 .

The proof will closely follow the path of Section 3.4, and many of the arguments will be similar, although the graphs we study will differ and there will be more technical complications. Once again the goal is to verify the hypotheses of Proposition 3.2.4 in order to get an exhaustion.

We will now introduce the objects we will use for the argument: similarly to Section 3.4 the 8-holed sphere will be represented as the doubling of an octagon. We will now introduce new curves γ_i , similarly to what we have done in Definition 3.4.1. We recall that, given a set X , its *disjoint union* with itself is the set $X \sqcup X = (X \times \{0\}) \cup (X \times \{1\})$.

Definition 3.5.1. Let \mathcal{O} be a regular hexagon in \mathbb{R}^2 , with vertices $\{v_1, \dots, v_8\}$ cyclically ordered. Let l_i be the side between v_i and v_{i+1} .

Let $S = \mathcal{O} \sqcup \mathcal{O} / \sim$ be the doubling of the heptagon, where $(x, 0) \sim (y, 1)$ if and only if $x = y \in \partial \mathcal{O}$. The marked surface $(S, \{v_1, \dots, v_8\})$ is a 8-holed sphere.

Let d_i be the segment in \mathcal{O} from the midpoint of l_{3i-3} and the midpoint of l_{3i} , where the subscripts are taken modulo 8. Let γ_i be the curve isotopic to the doubling of d_i .

We will denote the graph in $\mathcal{C}_{ss}(\Sigma_8)$ induced by the curves γ_i as \mathcal{O} .

The curves γ_i are pictured in Figure 3.20. This graph \mathcal{O} is an embedded 8-cycle which also have edges between opposite vertices: we will refer to this graph as an “octagon with (long) diagonals”.

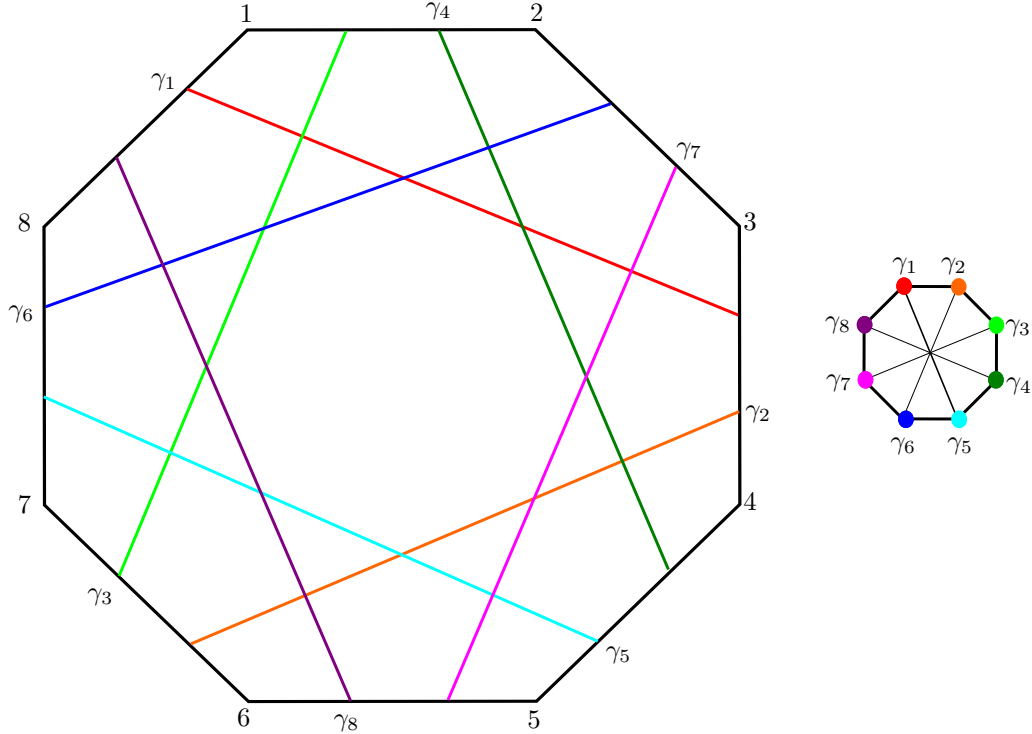


Figure 3.20: The octagon with diagonal $\mathcal{O} \subset \mathcal{C}_{ss}(\Sigma_8)$.

Let $\mathcal{O}' \subseteq \mathcal{O}$ be an embedded 8-cycle. When a property (usually distance between two

vertices) is satisfied in \mathcal{O}' , we say that such property holds “in an octagon inside \mathcal{O} ”. We will also call *copy of \mathcal{O}* a subgraph (in any graph) which is isomorphic to \mathcal{O} .

The curves γ_i are all 3-separating, hence \mathcal{O} is a subgraph of $\mathcal{C}_{(3)}(\Sigma_8)$ and [B2, Lemma 7.2] proves that it is a rigid subgraph there. However, this does not imply *a priori* that the graph \mathcal{O} is rigid as a subgraph of $\mathcal{C}_{ss}(\Sigma_8)$, as a graph embedding could potentially map 3-separating curves to 4-separating ones. We will see that such a rigidity result actually holds, but this will require some extra technical work we will take care of later. Indeed, we will initially prove rigidity of some subgraphs in $\mathcal{C}_{(3)}(\Sigma_8)$, obtaining an exhaustion by rigid sets for that graph along the way, and then use those results to construct the exhaustion for the entire strongly separating curve graph of the 8-holed sphere.

First, we need to define the mapping classes which will form the set H as in the statement of Proposition 3.2.4. Similarly to Definition 3.4.2 we first define the one-third Dehn-twists which will be part of H .

Definition 3.5.2. We will follow the notation we have introduced in Definition 3.5.1. Let $1 \leq i \leq 8$ and $j = 1, 2, 3$. In what follows the indices i are to be considered modulo 8. We define the arcs a_j^i as follows:

- The arc a_1^i is the isotopy class of oriented simple arcs from the marked point v_{3i-2} to v_{3i-1} isotopic to the side l_{3i-2} .
- The arc a_2^i is the isotopy class of oriented simple arcs from the marked point v_{3i-1} to v_{3i} isotopic to the side l_{3i-1} .
- The arc a_3^i is the isotopy class of oriented simple arcs from the marked point v_{3i} to v_{3i-2} isotopic to the segment in the “back” heptagon $\mathcal{O} \times \{1\} \subseteq S$ between the two endpoints.

We define the following one-third Dehn twists, as introduced in Definition 3.3.1,

$$R_i = R_{\gamma_i}^{a_1^i, a_2^i, a_3^i} \in \text{Mod}(\Sigma_8).$$

The oriented arcs a_j^i are represented in Figure 3.21. These arcs are not doubled. The arcs a_3^i , which are dotted, are intended to be on the back side of the sphere.

Analogously to Definition 3.4.3, we will now introduce a collection of 2-separating curves, around which the half-Dehn twists we need will be.

Definition 3.5.3. Let α_i be the unique essential simple closed curve on Σ_7 which is disjoint from every curve in \mathcal{H} apart from γ_i and γ_{i+1} , where subscripts are considered modulo 8.

We define mapping classes $H_i \in \text{Mod}(\Sigma_8)$ to be the right half-Dehn twists around the curve α_i .

The curves α_i are pictured in Figure 3.22. It is immediate to check these curves are 2-separating, so the half Dehn-twists around them are defined.

We are now ready to give the following definition, which is analogous to Definition 3.4.4.

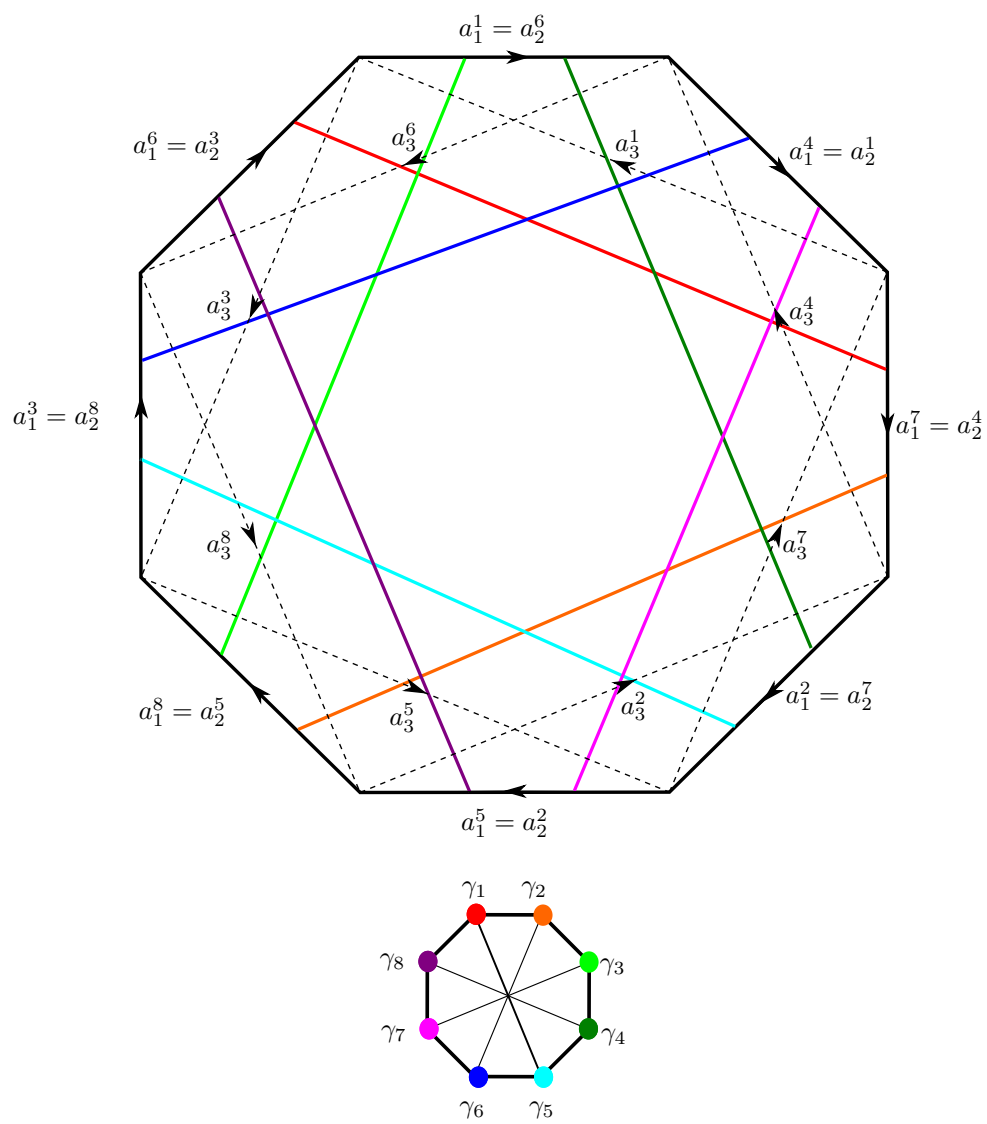


Figure 3.21: The arcs a_i^j .

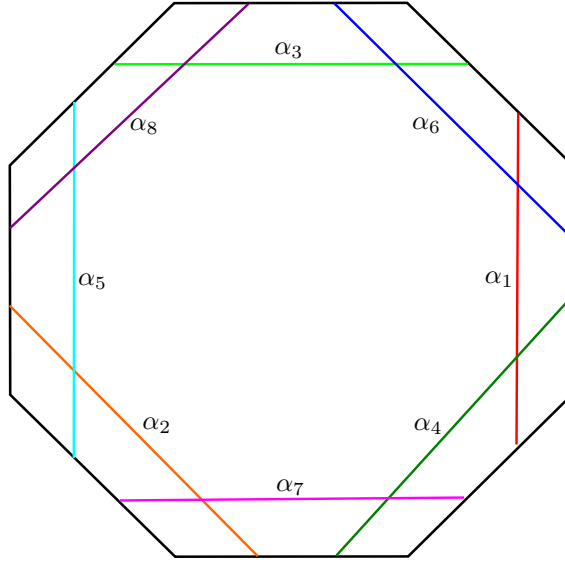


Figure 3.22: The curves α_i .

Definition 3.5.4. We define

$$G = \{\text{Id}_{\Sigma_8}\} \cup \{R_i^\varepsilon\}_{i=1,\dots,8}^{\varepsilon=\pm 1} \subset \text{Mod}(\Sigma_8).$$

Moreover we define

$$\overline{G} = G \cup \{H_i^\varepsilon\}_{i=1,\dots,8}^{\varepsilon=\pm 1} \subset \text{Mod}(\Sigma_8).$$

Once again the discussion in [FM, § 4.4.4] proves that the collection $\{H_i^\varepsilon\}$, hence \overline{G} , generates the entire mapping class group $\text{Mod}(\Sigma_8)$.

Our final goal will be to apply Proposition 3.2.4 with $H = \overline{G}$ and a suitable subgraph $X \subseteq \mathcal{C}_{ss}(\Sigma_8)$, in order to obtain an exhaustion. It should be noticed that the octagon with diagonals \mathcal{O} , or any graph obtained by the action of mapping classes onto it, is unsuitable as a candidate. This is for such a graph only contains 3-separating curves, hence does not intersect all the orbits of vertices under the action of the mapping class group. To take care of this problem we will expand the graph \mathcal{O} to a bigger one. In order to do so we start with the following lemma.

Lemma 3.5.5. *For every $i = 1, \dots, 8$ there exists a unique curve $\zeta \in \mathcal{C}_{ss}(\Sigma_8)$ which is disjoint from $\gamma_i, \gamma_{i+1}, \gamma_{i+4}$, and γ_{i+5} . Moreover, this curve is 4-separating.*

Proof. It is clear from Figure 3.23 that the surface

$$S \setminus (\gamma_i \cup \gamma_{i+1} \cup \gamma_{i+4} \cup \gamma_{i+5})$$

is composed of four surfaces of type $S_{0,1}^1$, two surfaces of type $S_{0,2}^1$, and a surface of type $S_{0,0}^2$, say A . A curve with the required disjointness property must lay in one of those subsurfaces. Any curve lying in a once-punctured disc is isotopic to the puncture, hence it is nonessential. Every essential curve lying in a twice-punctured disc is isotopic

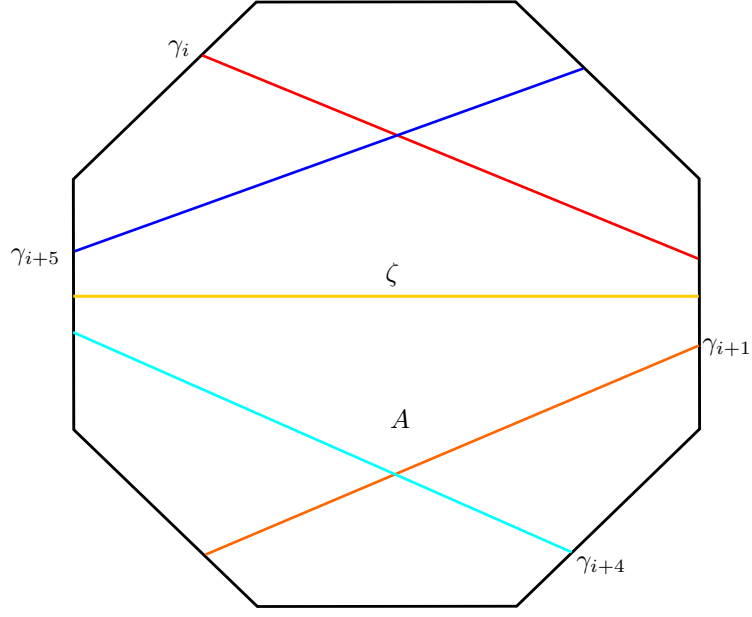


Figure 3.23: The yellow curve is the curve ζ .

to the boundary, hence it bounds a twice-punctured disc itself and so it is not strongly separating. It follows that the only strongly separating curve which is disjoint from the four aforementioned curves is isotopic to the boundary of the annular region, that is the curve in Figure 3.23, which is a 4-separating curve. \square

Given the previous lemma we give the following definition.

Definition 3.5.6. Let γ_i be as defined in Definition 3.5.1. Let ζ be the unique curve disjoint from $\gamma_1, \gamma_2, \gamma_5, \gamma_6$, provided by Lemma 3.5.5.

We now define the subgraph \mathcal{O}^* as the full subgraph of $\mathcal{C}_{ss}(\Sigma_8)$ induced by the vertices representing curves γ_i and ζ .

The graph \mathcal{O}^* is obtained from \mathcal{O} by adding a new vertex and new edges connecting it to $\gamma_1, \gamma_2, \gamma_5, \gamma_6$. A picture of such graph is provided in Figure 3.24.

We can notice that this graph is not symmetric, and the missing three 4-separating curves generated by Lemma 3.5.5 might be added in order to make it so, but our definition will be enough for our purposes and minimises the amount of cumbersome notation and technical checks needed in the proofs.

The graph \mathcal{O}^* now contains both 3 and 4-separating curves, that is it contains curves of every topological type represented in $\mathcal{C}_{ss}(\Sigma_8)$. It follows that the graph $\overline{G} \cdot \mathcal{O}^*$ now intersects every orbit of vertices under the action of $\text{Mod}(\Sigma_8)$, hence it is now a valid candidate for the subgraph X in the hypotheses of Proposition 3.2.4. Our main goal will now be to prove that this graph is rigid, that is the following statement, which is analogous to Proposition 3.4.5.

Proposition 3.5.7. *Let \mathcal{O}^* be as defined in Definition 3.5.6. Let \overline{G} be as in Definition 3.5.4. Let $i : \overline{G} \cdot \mathcal{O}^* \hookrightarrow \mathcal{C}_{ss}(\Sigma_8)$ be an injective graph homomorphism. Then there exists an extended mapping class $f \in \text{Mod}^\pm(\Sigma_8)$ such that $i = f$.*

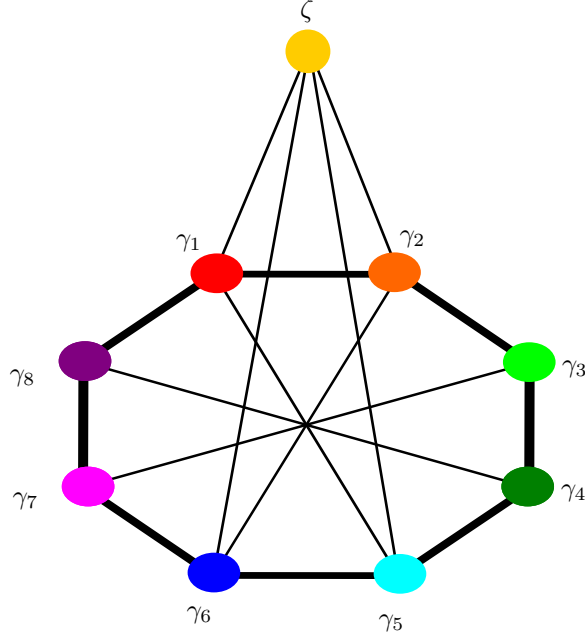


Figure 3.24: The graph \mathcal{O}^* .

In particular the graph $\overline{G} \cdot \mathcal{O}^* \subseteq \mathcal{C}_{ss}(\Sigma_8)$ is rigid.

In order to prove the previous proposition we will first prove rigidity of $\overline{G} \cdot \mathcal{O}$ as a subgraph of $\mathcal{C}_{(3)}(S_{0,8})$. Indeed, the most part of the reminder of this section will ultimately be devoted to the proof of the following.

Proposition 3.5.8. *Let \mathcal{O} be as defined in Definition 3.5.1. Let \overline{G} be as in Definition 3.5.4. Let $i : \overline{G} \cdot \mathcal{O} \hookrightarrow \mathcal{C}_{(3)}(\Sigma_8)$ be an edge-preserving graph homomorphism. Then there exists an extended mapping class $f \in \text{Mod}^\pm(\Sigma_8)$ such that $i_{|\overline{G} \cdot \mathcal{O}} = f_{\overline{G} \cdot \mathcal{O}}$.*

In particular, the graph $\overline{G} \cdot \mathcal{O} \subseteq \mathcal{C}_{(3)}(\Sigma_8)$ is rigid.

As for Proposition 3.4.5, in the last two statements we have willingly remarked that any injection is induced by a mapping class, rather than by a generic graph automorphism. In fact, this would allow us to apply Proposition 3.2.6 and deduce the combinatorial rigidity of the graphs, that is that $\text{Aut}(\mathcal{C}_{ss}(\Sigma_8)) \cong \text{Mod}^\pm(\Sigma_8)$, which already followed from Theorem 1.3.4 or Theorem A, and $\text{Aut}(\mathcal{C}_{(3)}(\Sigma_8)) \cong \text{Mod}^\pm(\Sigma_8)$, which follows from Theorem D.

The structure of the proof of the previous proposition will be almost exactly the same as for the proof of Proposition 3.4.5, the main difference being having to use different graphs, hence different relations between vertices. Similarly to the case of the 7-holed sphere we start by proving that we can reduce to the study of $G \cdot \mathcal{O}$, and in order to do so we state the following lemma, analogous of Lemma 3.4.6.

Lemma 3.5.9. *Let the curves γ_i and the mapping classes R_i and H_i be defined as before. Then the following relations hold:*

1. $H_i \gamma_i = R_{i-2} \gamma_i$;

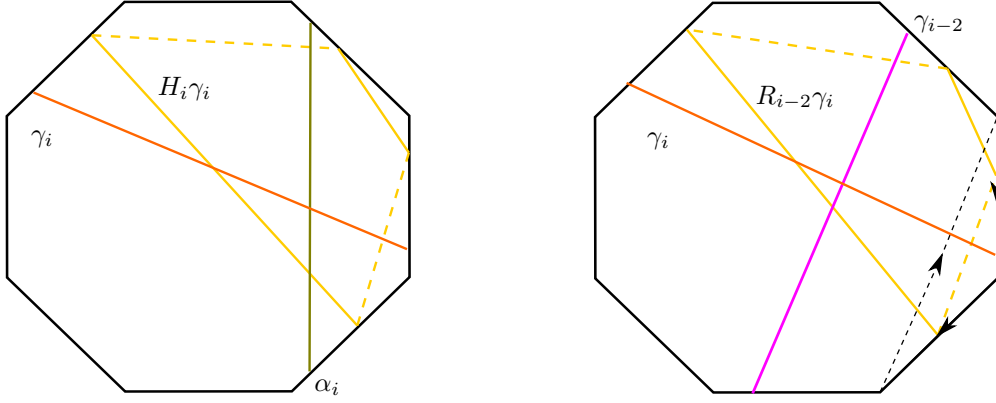


Figure 3.25: Equality 1 of Lemma 3.5.9.

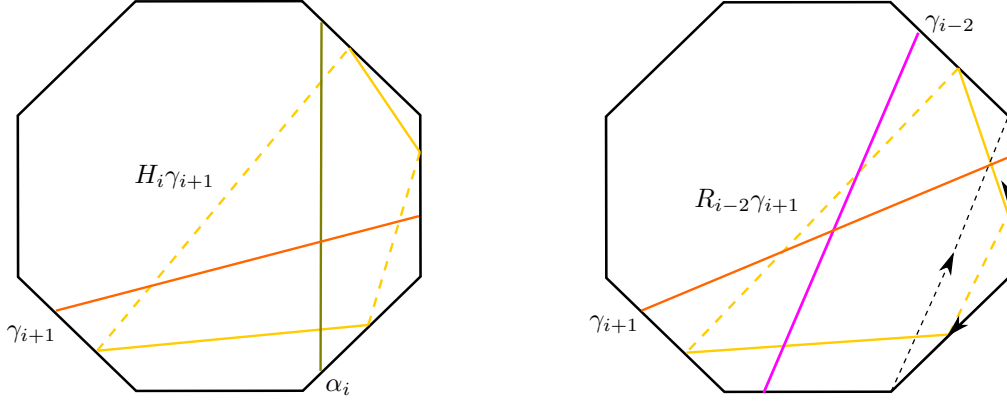


Figure 3.26: Equality 2 of Lemma 3.5.9.

2. $H_i \gamma_{i+1} = R_{i-2} \gamma_{i+1};$
3. $H_i^{-1} \gamma_i = R_{i+3}^{-1} \gamma_i;$
4. $H_i^{-1} \gamma_{i+1} = R_{i+3}^{-1} \gamma_{i+1};$

where the subscripts are intended to be modulo 8.

Proof. For the proof of Equalities 1, 2, 3 and 4, respectively, we refer to Figures 3.25, 3.26, 3.27, 3.28. \square

The following corollary is analogous to Corollary 3.5.10.

Corollary 3.5.10. *We have $\overline{G} \cdot \mathcal{O} = G \cdot \mathcal{O}$.*

Proof. By definition every curve α_i is disjoint from every curve in \mathcal{O} apart from γ_i and γ_{i+1} , hence for $j \neq i, i+1$ it follows that $H_i^\varepsilon \gamma_j = \gamma_j = R_j \gamma_j \in G \cdot \mathcal{O}$. From this observation combined with Lemma 3.5.9 it follows that for every i we have $\alpha_i \cdot \mathcal{O} \subseteq G \cdot \mathcal{O}$, hence $\overline{G} \cdot \mathcal{O} \subseteq G \cdot \mathcal{O}$. The other inclusion is trivial. \square

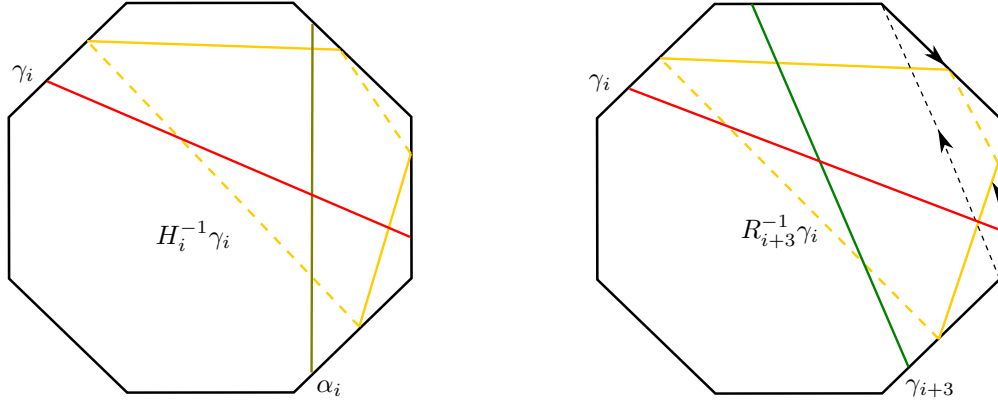


Figure 3.27: Equality 2 of Lemma 3.5.9.

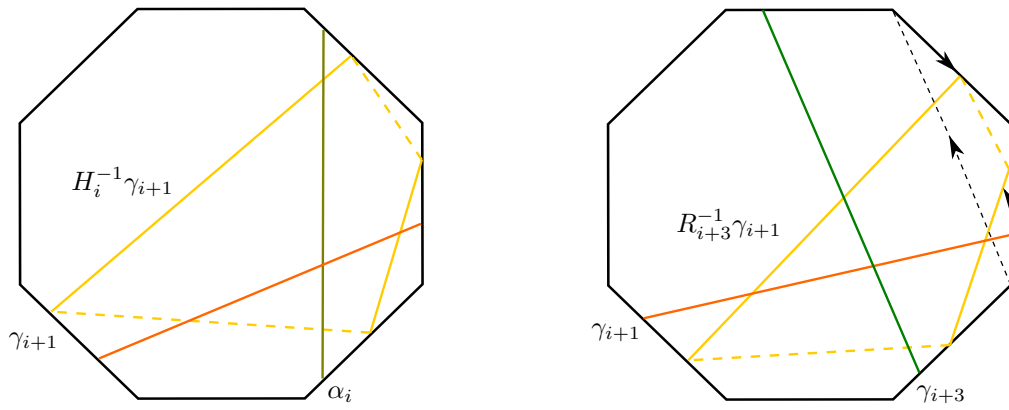


Figure 3.28: Equality 4 of Lemma 3.5.9.

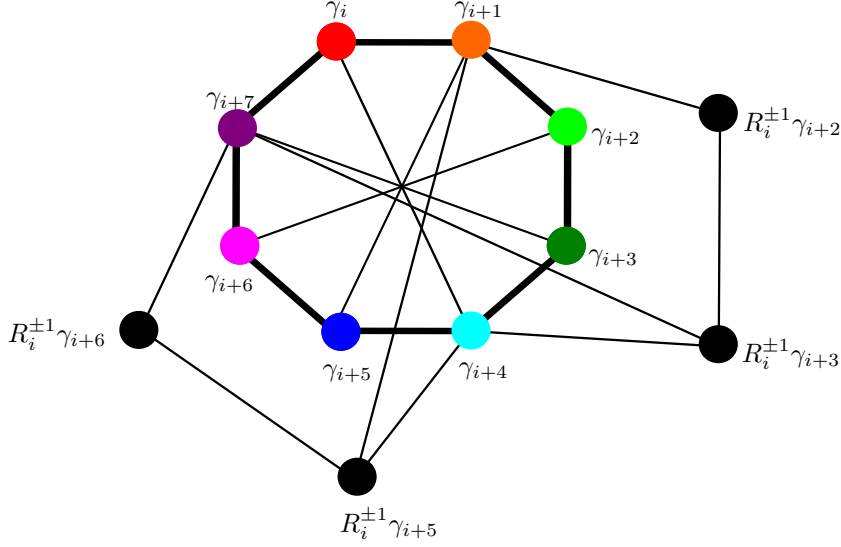


Figure 3.29: The graph W_j^ε .

Similarly to the previous section we observe that if γ_i and γ_j are disjoint (possibly equal) then $R_i \gamma_j = \gamma_j$. This implies that $\mathcal{O} \subset G \cdot \mathcal{O}$, and that every octagon with diagonal $R_i^\varepsilon \cdot \mathcal{O}$ will intersect with \mathcal{O} in four vertices, which are γ_{i-1} , γ_i , γ_{i+1} , and γ_{i+4} . Note that three of these vertices are consecutive, while the forth is the one “opposite” to the middle of the consecutive ones.

Similarly to the case we dealt with in the previous section, rigidity of $G \cdot \mathcal{O}$ will follow from the rigidity of some of its subgraphs. The subgraphs we will study, which will play the same role as those in Definition 3.4.9 while being slightly more simply defined, are the following.

Definition 3.5.11. Let \mathcal{O} be as in Definition 3.5.1. Let R_i be the one-third Dehn-twists defines in Definition 3.5.2. Let $\varepsilon \in \{+1, -1\}$. For $i = 1, \dots, 8$ let

$$W_j^\varepsilon = \mathcal{O} \cup (R_j^\varepsilon \cdot \mathcal{O}).$$

A picture of W_j^ε is provided in Figure 3.29.

Analogously to Definition 3.4.10, we will now define the *mirror reflection* mapping class.

Definition 3.5.12. We follow the notation introduce in Definition 3.5.1. Let $r : S \longrightarrow S$ the homeomorphism defined by $r(x, 0) = (x, 1)$ and $r(y, 1) = (y, 0)$. We will denote the isotopy class of r as the *mirror reflection* $\rho \in \text{Mod}^\pm(\Sigma_8)$

The mirror reflection we have just defined exchanges the “front” and the “back” of the doubled heptagon we are identifying with Σ_8 . This is an orientation-reversing automorphism of order two. Moreover, we notice that the mirror reflection ρ fixes all the curves γ_i .

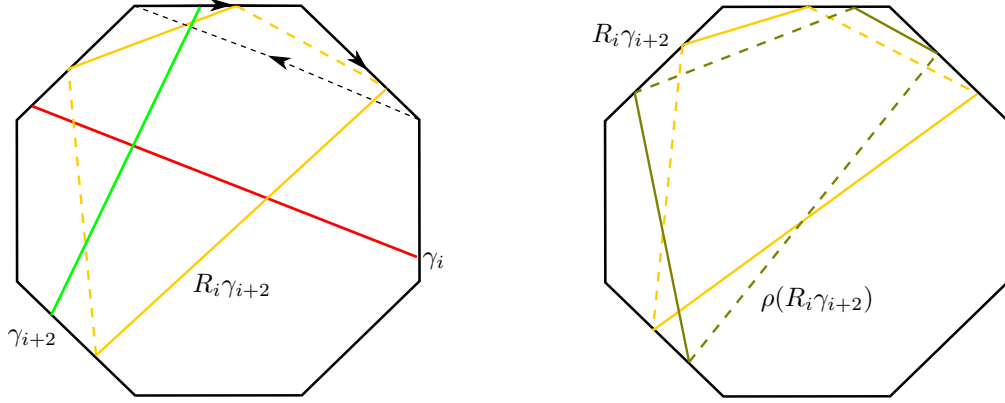


Figure 3.30: The yellow curve is $R_i\gamma_{i+2}$ while the olive one is its mirror image $\rho(R_i\gamma_{i+2})$.

We observe the following, which is the computation of the stabiliser of the octagon with diagonal \mathcal{O} .

Lemma 3.5.13. *Let \mathcal{O} be the octagon with diagonals defined in Definition 3.5.1. Let $f \in \text{Mod}^\pm(\Sigma_8)$ be an extended mapping class such that $f|_{\mathcal{O}} = \text{Id}_{\mathcal{O}}$. Then either $f = \text{Id}$, if it preserves the orientation, or $f = \rho$, if it is orientation-reversing.*

Proof. The proof is very similar to the one for Lemma 3.4.11.

Up to composing f with ρ we can assume f to be orientation-preserving. From now on, we will abuse notation and confuse the mapping class f with one of its representatives. The collection of curves $\gamma_1, \gamma_4, \gamma_5, \gamma_7$ satisfies the hypotheses of Proposition 1.2.6.

We claim that f preserves the orientation of every one of the aforementioned curves, hence we can apply Proposition 1.2.6 and deduce that f is isotopic to the identity.

The claim follows from the fact that if a surface homeomorphism fixes a curve but changes its orientation, then it must exchange the two connected components of the complementary of the curve. In particular, for this to be possible the two complementary discs of the curve need to be of the same topological type. However, every curve γ_i is 3-separating, hence it bounds a 3-punctured disc on one side and a 5-punctured disc on the other, so f must preserve its orientation. The claim is proven. \square

Analogously to Corollary 3.4.12, we have the following corollary, which will be useful later to check hypothesis of Proposition 3.2.4 regarding the triviality of the stabilisers.

Corollary 3.5.14. *Let i, j be such that $j \neq i - 1, i, i + 1, i + 4$ modulo 8. Then the subgraph $Y = \mathcal{O} \cup \{R_i\gamma_j\}$ has trivial pointwise stabiliser under the action of $\text{Mod}^\pm(\Sigma_8)$. In particular, such subgraph has trivial pointwise stabiliser.*

Proof. Let $f \in \text{Mod}^\pm(\Sigma_8)$ be an extended mapping class such that $f|_Y = \text{Id}_Y$. Then, since $f|_{\mathcal{O}} = \text{Id}_{\mathcal{O}}$ it follows from Lemma 3.4.11 that either $f = \text{Id}$ or $f = \rho$. Since $\rho(R_i\gamma_j) \neq R_i\gamma_j$, as can be seen in Figures 3.30, 3.31, 3.32, 3.33 for $j = i+2, i+3, i+5, i+6$, respectively, it follows that $f = \text{Id}$. \square

We will focus our attention towards the proof the following proposition, which will be the key intermediate step toward the proof of rigidity for $G \cdot \mathcal{O}$, and is a sort of rigidity

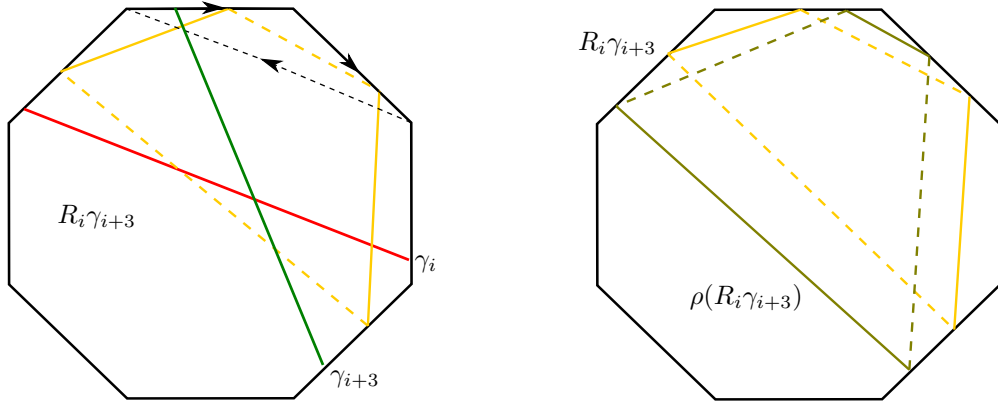


Figure 3.31: The yellow curve is $R_i\gamma_{i+3}$ while the olive one is its mirror image $\rho(R_i\gamma_{i+3})$.

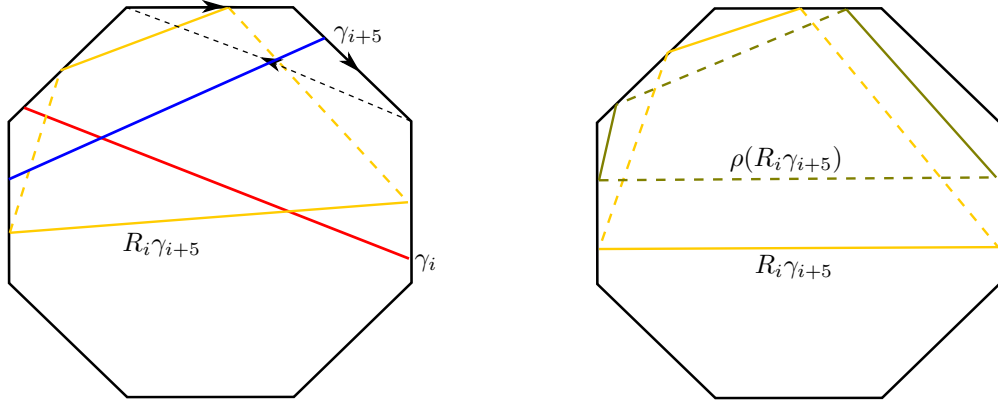


Figure 3.32: The yellow curve is $R_i\gamma_{i+5}$ while the olive one is its mirror image $\rho(R_i\gamma_{i+5})$.

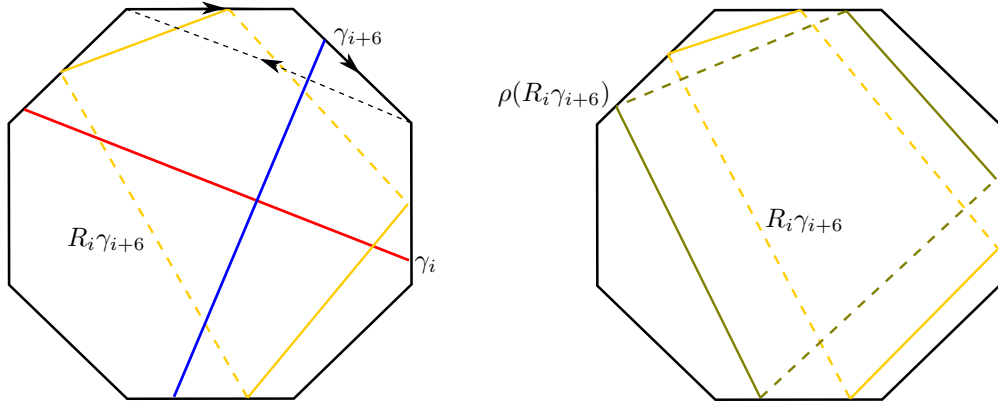


Figure 3.33: The yellow curve is $R_i\gamma_{i+6}$ while the olive one is its mirror image $\rho(R_i\gamma_{i+6})$.

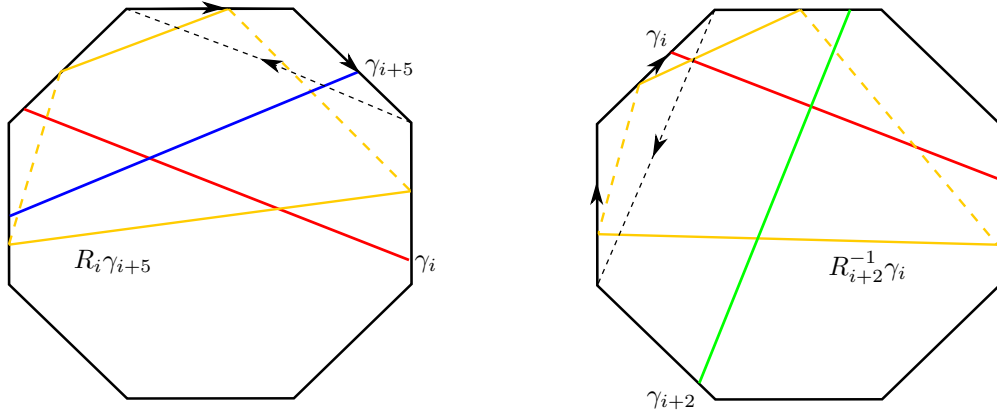


Figure 3.34: Equality 1 in Lemma 3.5.16.

result for the subgraphs W_i^ε . However, we will not prove that such subgraphs are rigid in the sense of Definition 1.6.1, as we will only prove that the embeddings which are restrictions of embeddings of the entire $G \cdot \mathcal{O}$ are topologically induced. However, this will suffice for our needs.

Proposition 3.5.15. *Let $i : G \cdot \mathcal{O} \hookrightarrow \mathcal{C}_{(3)}(S_{0,8})$ be an injective graph homomorphism such that $i|_{\mathcal{O}} = \text{Id}_{\mathcal{O}}$. Then for every W_j^ε we have that either $i|_{W_j^\varepsilon} = \text{Id}_{W_j^\varepsilon}$ or $i|_{W_j^\varepsilon} = \rho|_{W_j^\varepsilon}$.*

Before we proceed we will highlight some relations between the actions of some one-third Dehn-twists on some curves, which will be needed later. The following statement is similar to Lemma 3.4.8, although it contains fewer relations.

Lemma 3.5.16. *Let the curves γ_i and the mapping classes R_i and H_i be defined as before. The following relations hold:*

1. $R_i \gamma_{i+5} = R_{i+2}^{-1} \gamma_i$;
2. $R_i \gamma_{i+2} = R_{i+2}^{-1} \gamma_{i+5}$;

where subscripts are taken modulo 8.

Proof. For the proof of Equalities 1 and 2, respectively, we refer to Figures 3.8 and 3.9. The colour scheme is chosen to be consistent with the one employed in Figure 3.20, for $i = 1$. \square

The consequences of the previous lemma which will really be useful for us are contained in the following corollary.

Corollary 3.5.17. *Let the curves γ_i and the mapping classes R_i be defined as before. The following facts hold:*

1. The curves $R_i \gamma_{i+2}$ and γ_{i+2} are non-adjacent vertices in a copy of \mathcal{O} inside $G \cdot \mathcal{O}$;
2. The curves $R_i \gamma_{i+2}$ and γ_{i+3} are 2-distant vertices in an octagon of a copy of \mathcal{O} inside $G \cdot \mathcal{O}$;

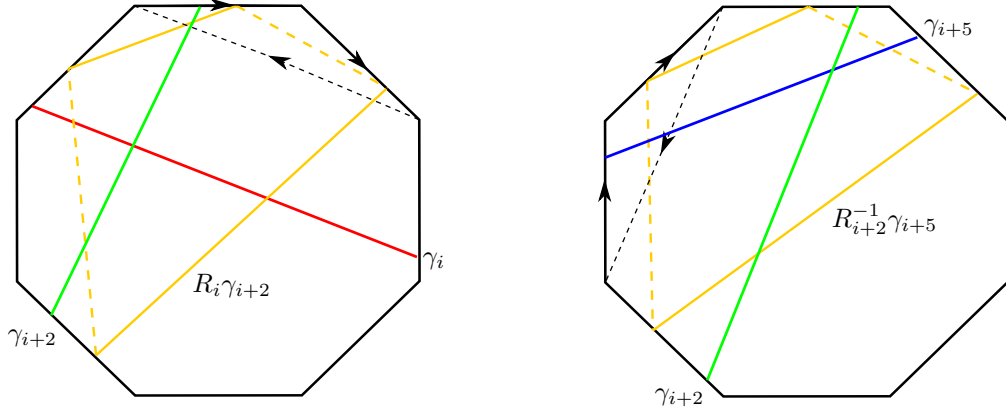


Figure 3.35: Equality 2 in Lemma 3.5.16.

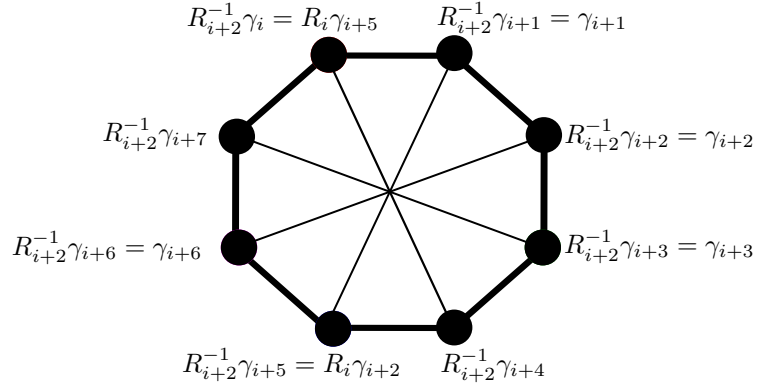


Figure 3.36: The graph $R_{i+2}^{-1} \cdot \mathcal{O}$.

3. The curves $R_i \gamma_{i+5}$ and γ_{i+6} are 2-distant vertices in an octagon of a copy of \mathcal{O} inside $G \cdot \mathcal{O}$;
4. The curves $R_i \gamma_{i+2}$ and γ_{i+6} are adjacent in $G \cdot \mathcal{O}$.

Proof. The proof follows immediately from the application of the relations we have stated in Lemma 3.5.16 to the octagon with diagonals $R_{i+2}^{-1} \cdot \mathcal{O} \subseteq G \cdot \mathcal{O}$, as pictured in Figure 3.36. \square

From now on, in order to make the notation less cumbersome, during the proofs we will only deal with the graph W_1^+ , the other cases being completely analogous. The first step to prove the previous rigidity result is to study the projection to the Kneser graph

$$\pi : \mathcal{C}_{(3)}(S_{0,8}) \longrightarrow K(8, 3)$$

as defined in §1.5.1.

As in the case of the 7-holed sphere one-third Dehn-twists acts on the Kneser graph as 3-cycles: for example the mapping class R_1 acts on the Kneser graph $K(8, 3)$ as the permutation (123).

A copy of \mathcal{O} in the Kneser graph is rigid thanks to [B2, Lemma 7.1]. Similarly to what we did in Lemma 3.4.14 we will now prove a sort of rigidity result for the projection of the graphs W_i^ε , once again only considering restrictions of embeddings of the entire $G \cdot \mathcal{O}$. Indeed, as in Proposition 3.5.15, we will only consider restriction of embeddings of the entire $G \cdot \mathcal{O}$.

Lemma 3.5.18. *Let $i : G \cdot \mathcal{O} \hookrightarrow \mathcal{C}_{(3)}(S_{0,8})$ be an edge-preserving injection such that $i|_{\mathcal{O}} = \text{Id}_{\mathcal{O}}$. Then $\pi \circ i|_{W_j^\varepsilon} = \pi|_{W_j^\varepsilon}$.*

Proof. In order to simplify our notation we will prove the lemma for W_1^+ , the other cases being completely analogous. Thanks to [B2, Lemma 7.1] we can assume, using a different labelling of punctures, that the octagon with diagonals \mathcal{O} has the following projection:

- $\pi(\gamma_1) = \{1, 2, 3\};$
- $\pi(\gamma_2) = \{4, 5, 6\};$
- $\pi(\gamma_3) = \{7, 8, 1\};$
- $\pi(\gamma_4) = \{2, 3, 4\};$
- $\pi(\gamma_5) = \{5, 6, 7\};$
- $\pi(\gamma_6) = \{8, 1, 2\};$
- $\pi(\gamma_7) = \{3, 4, 5\};$
- $\pi(\gamma_8) = \{6, 7, 8\}.$

From this we observe that two vertices which are 2-distant in an octagon inside a copy of \mathcal{O} have intersection exactly 1.

Moreover we notice that the fact that $i|_{\mathcal{H}} = \text{Id}_{\mathcal{H}}$ implies that $\pi(i(\gamma_j)) = \pi(\gamma_j)$ for every j .

Keeping in mind the structure of W_1^+ pictured in Figure 3.29 we can observe that the vertex $\pi(i(R_1\gamma_3))$ is adjacent to both vertices $\pi(i(\gamma_2)) = \{4, 5, 6\}$ and $\pi(i(\gamma_7)) = \{3, 4, 5\}$, hence

$$\pi(i(R_1\gamma_3)) \subset \{1, 2, 7, 8\}.$$

Similarly the vertex $\pi(i(T_1\gamma_4))$ is adjacent to both the vertex $\pi(i(\gamma_5)) = \{5, 6, 7\}$ and $\pi(i(\gamma_8)) = \{6, 7, 8\}$, so

$$\pi(i(R_1\gamma_4)) \subset \{1, 2, 3, 4\}.$$

The vertices $\pi(i(R_1\gamma_3))$ and $\pi(i(R_1\gamma_4))$ are 2-distant in an octagon in a copy of \mathcal{O} due to Fact 2 in Corollary 3.5.17. Thanks to a previous observation, it follows that they have nontrivial intersection, so in particular $2 \in \pi(i(R_1\gamma_3))$. Since $\pi(i(R_1\gamma_3))$ and $\pi(i(R_1\gamma_4))$ are disjoint it follows that $2 \notin \pi(i(R_1\gamma_4))$, hence

$$\pi(i(R_1\gamma_4)) = \{1, 3, 4\} = (123) \cdot \{2, 3, 4\} = \pi(R_1\gamma_4).$$

For the same reason, since $1 \in \pi(R_1\gamma_4)$ it follows that $1 \notin \pi(i(R_1\gamma_3))$, hence we have that

$$\pi(i(R_1\gamma_3)) = \{2, 7, 8\} = (123) \cdot \{7, 8, 1\} = \pi(R_1\gamma_3).$$

In a completely analogous fashion the vertex $\pi(i(R_1\gamma_6))$ is adjacent to both the vertex $\pi(i(\gamma_5)) = \{5, 6, 7\}$ and $\pi(i(\gamma_2)) = \{4, 5, 6\}$, so

$$\pi(i(R_1\gamma_6)) \subset \{1, 2, 3, 8\}.$$

Moreover, the vertex $\pi(i(R_1\gamma_7))$ is adjacent to both the vertex $\pi(i(\gamma_8)) = \{6, 7, 8\}$ and $\pi(i(R_1\gamma_3)) = \{2, 7, 8\}$, as we have just proven, so

$$\pi(i(R_1\gamma_7)) \subset \{1, 3, 4, 5\}.$$

The vertices $i(R_1\gamma_6)$ and $i(\gamma_7)$ are 2-distant in an octagon in a copy of \mathcal{O} due to Corollary 3.5.17 3, hence their projection have nontrivial intersection, from which it follows that $3 \in \pi(i(R_1\gamma_6))$. From an argument completely similar the one employed before, it follows that $1 \in \pi(i(R_1\gamma_7))$. We now have the following:

$$\begin{aligned}\pi(i(R_1\gamma_6)) &= \{2, 3, 8\} = (123) \cdot \{8, 1, 2\} = \pi(R_1\gamma_6); \\ \pi(i(R_1\gamma_7)) &= \{1, 4, 5\} = (123) \cdot \{3, 4, 5\} = \pi(R_1\gamma_7),\end{aligned}$$

and the lemma is proven. \square

Similarly to what we did in the previous section now that we have dealt with the combinatorics it is time to move to topological arguments to prove rigidity of the graphs W_i^ε .

Beforehand, we observe that, thanks to the rigidity of \mathcal{O} inside $\mathcal{C}_{(3)}(S_{0,8})$, that is [B2, Lemma 7.2], every two curves which represent non-adjacent vertices in a copy of \mathcal{O} intersect exactly twice. In particular this holds for the curves in Corollary 3.5.17 1)-3).

Moreover, we state the following result, concerning uniqueness of common neighbours in $\mathcal{C}_{(3)}(S_{0,8})$ for vertices of \mathcal{O} .

Lemma 3.5.19. *Let α, β be two 2-distant curves in an octagon of a copy of \mathcal{O} inside $\mathcal{C}_{(3)}(S_{0,8})$. Then there exists a unique curve in $\mathcal{C}_{(3)}(S_{0,8})$ which is disjoint from both α and β .*

Proof. Let X_α (resp. X_β) be the complementary disc of α (resp. β) containing exactly 3 marked points. Since \mathcal{O} is rigid ([B2, Lemma 7.2]) from Figure 3.20 (for example by looking at γ_1 and γ_3) we can observe that the region $D = X_\alpha \cup X_\beta$ is a 5-punctured disc, hence the complementary disc D^c contains exactly 3 marked points. Every curve disjoint from both α and β is either contained in D , or in either of the 2-punctured disc $X_\alpha \setminus X_\beta$ or $X_\beta \setminus X_\alpha$. It follows that the only 3-separating curve which is disjoint from both α and β is ∂D . \square

We are now ready to proceed to the proof of Proposition 3.5.15

Proof of Proposition 3.5.15. In order to make the notation less cumbersome we will only prove the result for W_1^+ , the other cases being completely analogous.

We first fix some notation, as in Proposition 3.4.13. For every j let B_j be the 3-punctured complementary disc of γ_j and let O_j be the 5-punctured one. For a quick way to visualise the relations used throughout the proof we refer to Figure 3.29.

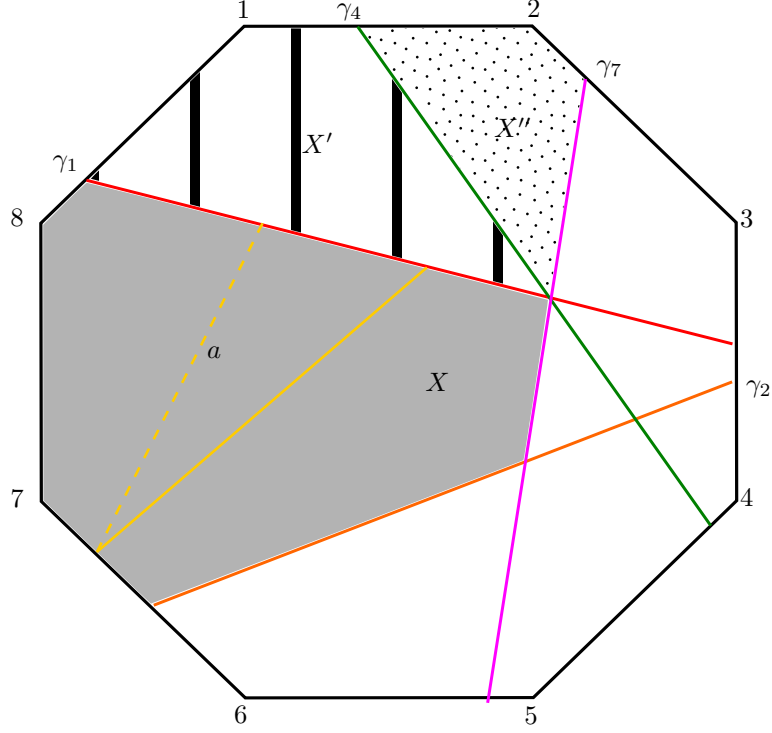


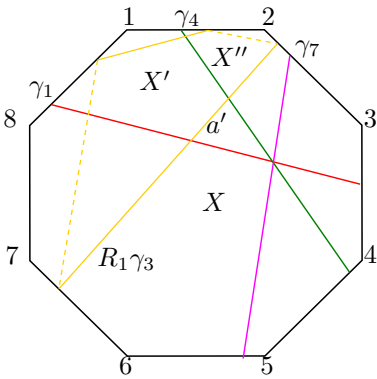
Figure 3.37: The gray area is the region X , the striped area is the region X' , the dotted one is the region X'' .

Let $\beta = i(R_1\gamma_3)$. We first claim that either $\beta = R_1\gamma_3$ or $\beta = \rho(R_1\gamma_3)$. We have observed that $i(\gamma_1, \beta) = 2$, since the two curves are non-adjacent vertices in $R_1 \cdot \mathcal{O}$. It follows that $a = \beta \cap O_1$ is a single arc with both endpoints on γ_1 and disjoint from γ_2 . Moreover, Fact 4 in Corollary 3.5.17, alongside the fact that the map i is defined on the entire $G \cdot \mathcal{O}$, proves that β , and hence a , is disjoint from γ_7 .

Since $i(\gamma_7, \gamma_1) = i(\gamma_7, \gamma_2) = 2$ it follows that the region $X = O_1 \cap O_2 \cap O_7$ is a twice punctured disc containing the marked points $\{7, 8\} \subseteq \pi(\beta)$, thanks to Lemma 3.5.18. Moreover ∂X is composed of a single arc of γ_1 , a single arc of γ_2 and two disjoint arcs of γ_7 .

Let now B_β be the complementary disc of β containing exactly 3 marked points. The disc B_β intersects γ_1 while it is disjoint from both γ_2 and γ_7 . It follows that the closure of the subspace $X \setminus B_\beta$ does not contain any marked point, and its boundary contains an arc of γ_2 and two arcs of γ_7 . It follows that a is isotopic to the closure of $\partial X \setminus \gamma_1$, under an isotopy supported inside X , and which keeps the endpoints on γ_1 at every time. This arc is depicted in Figure 3.37; in particular such arc is uniquely determined.

In order to determine β we are now left to study the topological type of the arc $a' = \beta \cap B_1$. Since β and γ_7 are disjoint it follows that $a' \subset B_1 \cap O_7$, which is a twice-punctured disc containing the set of marked points $\{1, 2\}$. The curve γ_4 , which intersects both γ_1 and γ_7 exactly twice and separates the marked point 1 from 2, cuts $B_1 \cap O_7$ into two once-punctured discs X', X'' . Without loss of generality we can assume the curves to be arranged in a way such that they are still in minimal position and transverse, but there are two triple points, that is the triple intersection $\gamma_1 \cap \gamma_4 \cap \gamma_7$ contains exactly



We have that $i(\beta, \gamma_4) = 2$ as they are two 2-distant vertices in an octagon of a copy of \mathcal{O} thanks to Fact 2 of Corollary 3.5.17. From this it follows that $a' \cap X''$ is a single nontrivial arc with both endpoints on γ_4 , hence it is completely determined. Moreover, the component of $X'' \setminus a$ which does not intersect γ_7 contains the marked point 2. It now follows that $a' \cap X'$ is composed of two arcs each with an endpoint on γ_1 and the other on γ_4 : we claim these arcs are isotopic. If not, the region of $X' \setminus a'$ whose boundary contains the two arcs would contain the marked point 1, hence the two marked points 1, 2 would both belong to the same component of $\Sigma_8 \setminus \beta$, which is impossible as Lemma 3.5.18 implies that

$$\pi(\beta) = \pi(R_1\gamma_3) = (123) \cdot \{7, 8, 1\} = \{2, 7, 8\}.$$

Up to postcomposing with the mirror image ρ we can now assume $i(R_1\gamma_3) = T_1\gamma_3$. A completely analogous argument can now be applied to $R_1\gamma_7$, proving that either $i(R_1\gamma_7) = R_1\gamma_7$ or $i(R_1\gamma_7) = \rho(R_1\gamma_7)$. However, the latter case is impossible since $R_1\gamma_3$ and $\rho(R_1\gamma_7)$ are not disjoint, as shown in Figure 3.40. It follows that $i(R_1\gamma_7) = R_1\gamma_7$.

152

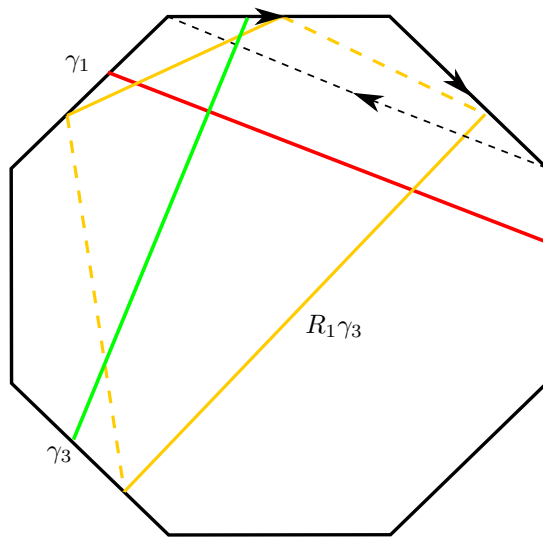


Figure 3.39: The yellow curve is $R_1\gamma_3$.

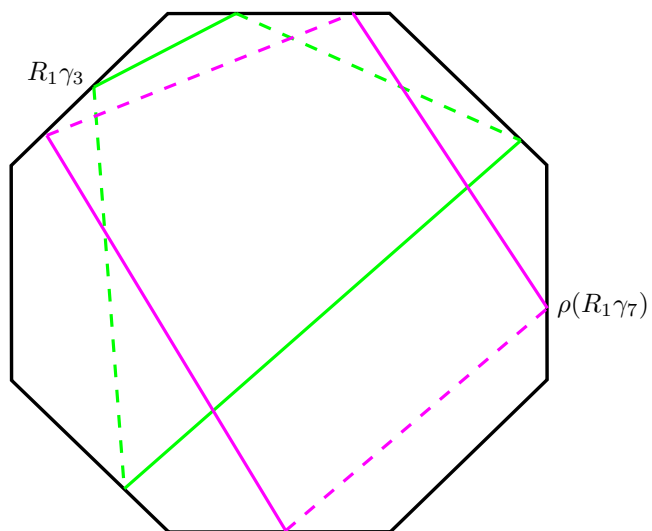


Figure 3.40: The green curve is $T_1\gamma_3$, the pink one is $\rho(T_1\gamma_7)$.

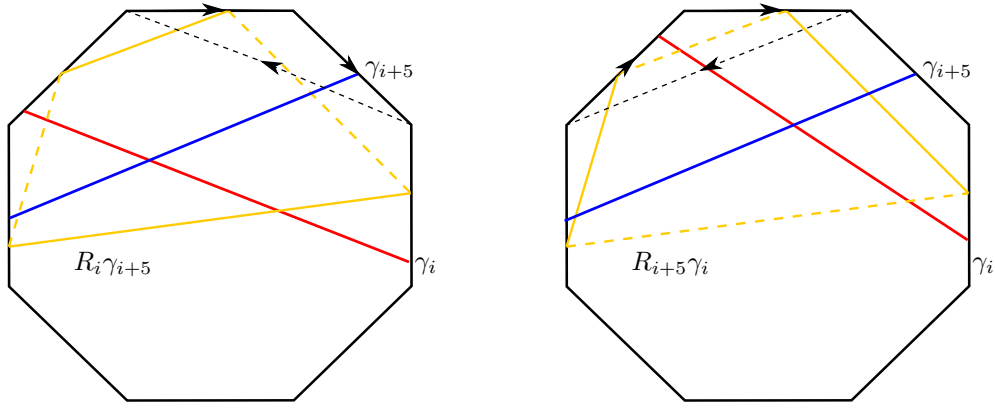


Figure 3.41: The curves $R_i \gamma_{i+5}$ and $R_{i+5} \gamma_i$, which are mirror images of one another.

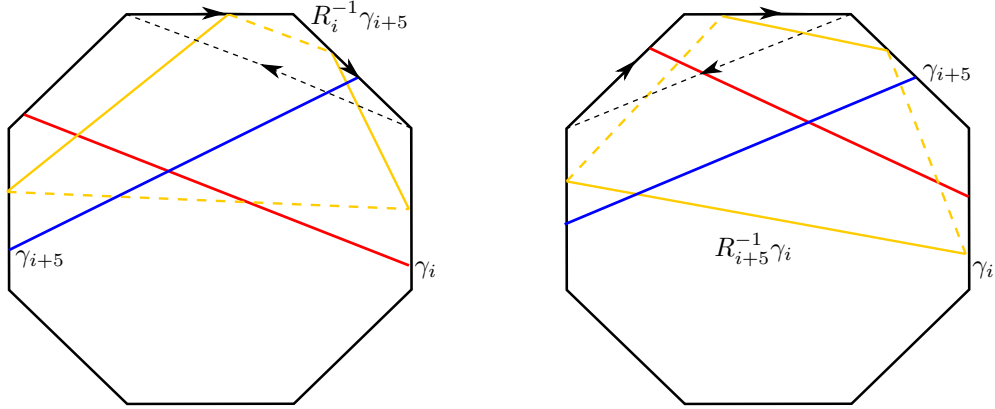


Figure 3.42: The curves $R_i^{-1} \gamma_{i+5}$ and $R_{i+5}^{-1} \gamma_i$, which are mirror images of one another.

We can now conclude that if $i(R_1 \gamma_3) = R_1 \gamma_3$ then $i|_{W_1^+} = \text{Id}_{W_1^+}$. Otherwise, if $i(R_1 \gamma_3) = \rho(R_1 \gamma_3)$, composition with ρ proves that $i|_{W_1^+} = \rho|_{W_1^+}$, and the proposition is proven. \square

We can now move to the proof of rigidity for the entire subgraph $G \cdot \mathcal{O} \subset \mathcal{C}_{(3)}(\Sigma_8)$. Before we do so, we state and prove the following useful lemma.

Lemma 3.5.20. *Let γ_i and R_i defined as before. Let ρ be the mirror reflection map. The following relation holds:*

$$R_i^\varepsilon \gamma_{i+5} = \rho(R_{i+5}^\varepsilon \gamma_i)$$

where subscripts are intended modulo 8.

Proof. We refer to Figure 3.41 and Figure 3.42 for the $\varepsilon = +1$ and $\varepsilon = -1$ cases, respectively. \square

We are now ready to prove that the restriction of embeddings $G \cdot \mathcal{O} \hookrightarrow \mathcal{C}_{(3)}(S_{0,8})$ to subgraphs W_i^ε are consistently either always the identity or always the mirror reflection.

Lemma 3.5.21. *Let $i : G \cdot \mathcal{O} \hookrightarrow \mathcal{C}_{(3)}(S_{0,8})$ be an injective graph homomorphism such that $i|_{\mathcal{O}} = \text{Id}_{\mathcal{O}}$. Then either $i = \text{Id}_{G \cdot \mathcal{O}}$ or $i = \rho|_{G \cdot \mathcal{O}}$.*

Proof. Thanks to Proposition 3.5.15, the map i induces either the identity or the mirror reflection of each of the subgraphs W_j^ε : as in the proof Lemma 3.4.16 the argument will revolve around proving that the induced map is the same for all graphs.

First, we claim that $i|_{W_j^\varepsilon} = \text{Id}_{W_j^\varepsilon}$ if and only if $i|_{W_{j+2}^{-\varepsilon}} = \text{Id}_{W_{j+2}^{-\varepsilon}}$. If this were not the case, up to postcomposing with ρ , from Lemma 3.5.16 1 we would have

$$R_j^\varepsilon \gamma_{j+5} = i(R_j^\varepsilon \gamma_{j+5}) = i(R_{j+2}^{-\varepsilon} \gamma_i) = \rho(R_{j+2}^{-\varepsilon} \gamma_i) = \rho(R_j^\varepsilon \gamma_{j+5})$$

which is a contradiction, as can be seen in Figure 3.32. The claim is proven.

Second, we now claim that $i|_{W_j^\varepsilon} = \text{Id}_{W_j^\varepsilon}$ if and only if $i|_{W_{j+5}^\varepsilon} = \text{Id}_{W_{j+5}^\varepsilon}$. If not, up to postcomposing with ρ , thanks to Lemma 3.5.20 we would have

$$i(R_j^\varepsilon(\gamma_{j+5})) = R_j^\varepsilon(\gamma_{j+5}) = \rho(R_{j+5}^\varepsilon \gamma_j) = i(R_{j+5}^\varepsilon \gamma_j)$$

contradicting injectivity of i , as $R_j^\varepsilon(\gamma_{j+5}) \neq R_{j+5}^\varepsilon \gamma_j$ as can be deduced from Figures 3.41 and 3.42.

The set of subgraphs $\{W_j^\varepsilon\}$ can be partitioned into the following four subsets:

$$\begin{aligned} P_1 &= \left\{ W_{2k+1}^{(-1)^{k+1}} \right\}; \\ P_2 &= \left\{ W_{2k}^{(-1)^{k+1}} \right\}; \\ P_3 &= \left\{ W_{2k+1}^{(-1)^k} \right\}; \\ P_4 &= \left\{ W_{2k}^{(-1)^k} \right\}. \end{aligned}$$

The first claim proves that all for all the subgraphs in the same P_h the restriction of i is consistently the identity or the mirror reflection. Moreover we observe that $W_1^+ \in P_1$, $W_6^+ = W_{1+5}^+ \in P_2$, $W_3^+ = W_{6+5}^+ \in P_3$ and $W_8^+ = W_{3+5}^+ \in P_4$. The second claim hence implies that the choice for the induced map is consistently the same for all the subgraphs W_j^ε . The lemma is proven. \square

We are now ready to prove rigidity of $\overline{G} \cdot \mathcal{O}$ inside $\mathcal{C}_{(3)}(S_{0,8})$, that is Proposition 3.5.8.

Proof of Proposition 3.5.8. Let $i : \overline{G} \cdot \mathcal{O} \hookrightarrow \mathcal{C}_{(3)}(S_{0,8})$ be an injective graph homomorphism. Then, thanks to [B2, Lemma 7.2], up to postcomposing with an orientation preserving mapping class $\varphi \in \text{Mod}(S_{0,8})$, we can assume that $\varphi \circ i|_{\mathcal{O}} = \text{Id}_{\mathcal{O}}$. It now follows from Lemma 3.5.21 that either $i = \varphi|_{\overline{G} \cdot \mathcal{O}}^{-1}$ or $i = \rho \circ \varphi|_{\overline{G} \cdot \mathcal{O}}^{-1}$ and both maps are induced by extended mapping classes. \square

From the results we have proven so far we can deduce the following result, which is the existence of an exhaustion for the graph $\mathcal{C}_{(3)}(S_{0,8})$.

Theorem 3.5.22. *The graph $\mathcal{C}_{(3)}(\Sigma_8)$ admits an exhaustion by finite rigid sets with trivial pointwise stabilisers.*

Proof. We can apply Proposition 3.2.4 with $\Gamma = \mathcal{C}_{(3)}(\Sigma_8)$, $H = \overline{G}$, and $X = \overline{G} \cdot \mathcal{O}$.

Rigidity of X is Proposition 3.5.8. The triviality of stabilisers follows from the fact that every set subgraph of the form $(g \cdot X) \cap X$ contains \mathcal{O} and another curve of the form $R_i \gamma_j$ for $j \neq i-1, i, i+1, i+4$, hence it is a superset of a subgraph which has trivial pointwise stabiliser due to Corollary 3.5.14, hence its stabiliser is itself trivial. This also implies that X has trivial pointwise stabiliser, hence every subgraph of the exhaustion has trivial pointwise stabiliser due to Lemma 3.2.1.

Lastly, $\text{Mod}(\Sigma_8) = \langle H \rangle$ acts transitively on 3-separating curves in Σ_8 , since they are all of the same topological type, hence there is only one orbit of vertices, so the requirement to have an exhaustion is fulfilled. \square

We can now move towards the proof of the rigidity of $\overline{G} \cdot \mathcal{O}^* \subset \mathcal{C}_{ss}(\Sigma_8)$, where \mathcal{O}^* is the augmentation of the octagon with diagonal as in Definition 3.5.6. In order to complete this proof we will need to make sure that every embedding of \mathcal{O} in the strongly separating curve graph is actually contained in the subgraph $\mathcal{C}_{(3)}(\Sigma_8)$. This is the content of the following.

Proposition 3.5.23. *Let $i : \mathcal{O} \hookrightarrow \mathcal{C}_{ss}(\Sigma_8)$ be an injective graph homomorphism. Then we have $i(\mathcal{O}) \subset \mathcal{C}_{(3)}(\Sigma_8)$.*

Before we move to the proof of the previous result we state and prove a technical lemma which will be useful for the proof of the previous proposition.

Lemma 3.5.24. *Let $\gamma \subset \Sigma_8$ be a 4-separating curve, and let A, B be the two complementary discs of γ . Let $\alpha, \alpha' \subset \Sigma_8$ be two different 3-separating curves both disjoint from γ . Let $\delta \subset \Sigma_8$ be a 3 or 4-separating curve disjoint from both α and α' .*

The curves α and α' are disjoint if and only if, up to relabelling the components, we have $\alpha \subset A$ and $\alpha' \subset B$.

If α and α' are not disjoint from each other then either $\alpha, \alpha' \subset A$ or $\alpha, \alpha' \subset B$. Moreover, the curve δ is contained in the other component, it is disjoint from γ and, if it is 4-separating, it is isotopic to γ .

Proof. For the first claim it is enough to notice that the regions A, B are 4-punctured discs so if one of them contained both α and α' it would have to contain two disjoint 3-punctured discs, which is clearly impossible.

For the second claim it is clear that both α, α' belong to the same complementary disc of γ , since they intersect with each other. Without loss of generality let us assume that $\alpha, \alpha' \subset A$. Let us argue by contradiction and suppose that $i(\delta, \gamma) > 0$. Assuming that δ and γ are in minimal position it follows that the closure of every component of $A \setminus \delta$ is a possibly punctured disc and, since the intersection was not trivial, at least two discs contain at least one puncture, so no one contains more than 3. Since both α and α' must be contained in the same component it follows that they are both contained in a k -punctured disc D , with $k \leq 3$. If $k < 3$ such a containment is straight up impossible; on the other hand, if $k = 3$ then it must hold $\alpha = \partial D = \alpha'$, which is a contradiction. It follows that δ is disjoint from γ and hence contained in B . In order to conclude our argument we observe that the only 4-separating curve contained in B is isotopic to $\partial B = \gamma$. The lemma is proven. \square

We will now move to the proof of Proposition 3.5.23.

Proof of Proposition 3.5.23. Before we commence the proof we observe that three 3-separating curves on Σ_8 cannot be pairwise disjoint (this is equivalent to saying there are no triangles in $\mathcal{C}_{(3)}(\Sigma_8)$). In particular it follows that if a curve γ is disjoint from two disjoint 3-separating curves then γ must be 4-separating. Moreover we remark that two disjoint 4-separating curves in Σ_8 must be isotopic. Speaking in terms of the curve graph this means that in any triangle inside $\mathcal{C}_{ss}(\Sigma_8)$ exactly one vertex is a 4-separating curve. Moreover, Lemma 3.5.24 implies that the two 3-separating curves in such a triangle lie in different complementary components of the 4-separating one.

Let $\beta_1 \dots, \beta_8$ be an embedding of \mathcal{O} into $\mathcal{C}_{ss}(\Sigma_8)$, that is eight different vertices of $\mathcal{C}_{ss}(\Sigma_8)$ such that β_i is adjacent to β_j for $j = i - 1, i + 1, i + 4$, where subscripts are taken modulo 8. We will argue by contradiction and suppose, up to a change of labels, that the curve β_1 is 4-separating. Thanks to the previous observation this implies that the curves β_2, β_5 and β_8 are 3-separating. For the entire proof let A, B be the two complementary discs of γ_1 , each of which contains exactly 4 marked points.

We will first deal with the case when the curve β_4 is 4-separating and prove this leads to a contradiction. The case where β_6 is 4-separating is completely analogous up to a relabelling of the vertices. This case implies that the curve β_3 is 3-separating. We claim that the curves β_5 and β_8 are disjoint: if not, an application of Lemma 3.5.24 with $\gamma = \beta_1$, $\alpha = \beta_5$, $\alpha' = \beta_8$ and $\delta = \beta_4$ would imply that $\beta_1 = \beta_4$, which is impossible. Up to exchanging the components A, B , if needed, it follows from Lemma 3.5.24 that $\beta_5 \subset A$ and $\beta_8 \subset B$.

The curve β_2 is disjoint from β_1 , hence is contained in either A or B . We start by dealing with the case where $\beta_2 \subset A$. A picture with the edges relevant for the proof (but which does not include all the edges between the pictured vertices) is provided in Figure 3.43. An application of Lemma 3.5.24 with $\gamma = \beta_1$, $\alpha = \beta_2$ and $\alpha' = \beta_5$ and $\delta = \beta_6$ proves that β_6 is 3-separating and $\beta_6 \subset B$. Similarly, β_7 is disjoint from both β_6 and β_8 so $\beta_7 \subset A$. Since β_3 is disjoint from both β_2 and β_7 it also follows that $\beta_3 \subset B$. Since β_3 and β_8 are both contained in B , another application of Lemma 3.5.24 with $\gamma = \beta_1$, $\alpha = \beta_3$ and $\alpha' = \beta_8$ and $\delta = \beta_4$ now proves that $\beta_4 = \beta_1$, which is a contradiction.

We will now deal with the case when $\gamma_2 \subset B$. In particular this implies that β_2 and β_5 are disjoint. Since $\beta_2, \beta_5, \beta_6$ is a triangle it follows that β_6 is 4-separating, hence β_7 is 3-separating. From Lemma 3.5.24, applied with $\gamma = \beta_6$, $\alpha = \beta_2$ and $\alpha' = \beta_5$ and $\delta = \beta_7$, it follows that β_7 is disjoint from either β_2 or β_5 . In the first case the three curves $\beta_2, \beta_3, \beta_7$ form a triangle of 3-separating curves, which is impossible. A picture containing the relevant edges in this case (the other being similar) is provided in Figure 3.44. In the other case, as $\beta_5 \subset A$ and $\beta_8 \subset B$, and in particular they are disjoint, the curves $\beta_5, \beta_7, \beta_8$ form a triangle of 3-separating curves, which is a contradiction.

We are now left to the case where both β_4 and β_6 are 3-separating curves. We will first deal with the case where, after switching components and/or β_2 with β_8 if need be, we have $\beta_2 \subset A$ and $\beta_5 \subset B$. In particular this implies that β_2 and β_5 are disjoint, hence the curves $\beta_2, \beta_5, \beta_6$ form a triangle of 3-separating curves, which is a contradiction. This case is pictured in Figure 3.45.

We are now left with the case when, up to switching components, $\beta_2, \beta_5, \beta_8 \subset A$. A picture with the edges that are relevant to this case is displayed in Figure 3.46. From

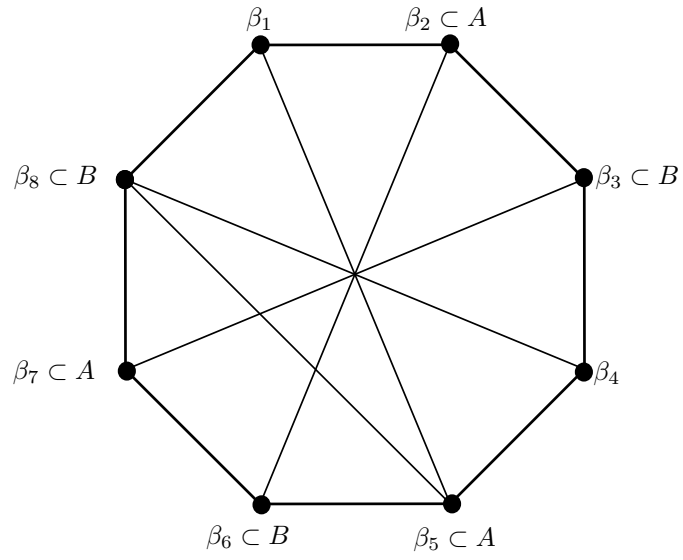


Figure 3.43: The curves β_1, β_4 are 4-separating.

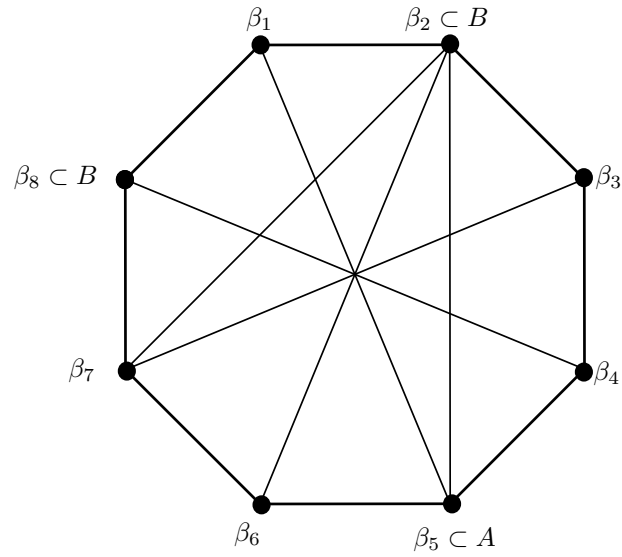


Figure 3.44: The curves $\beta_1, \beta_4, \beta_6$ are 4-separating.

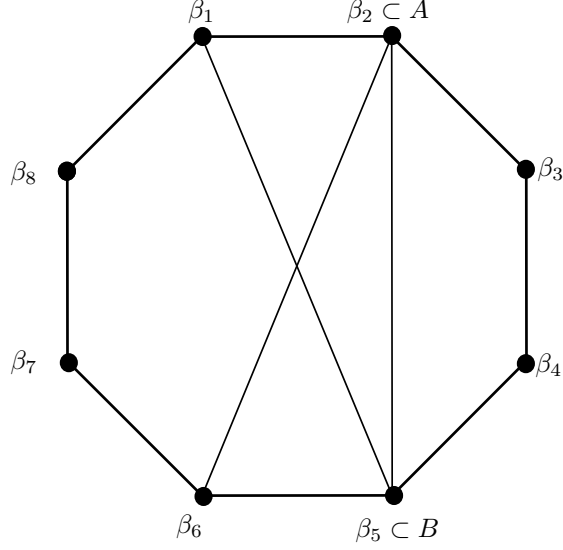


Figure 3.45: The curve β_1 is 4-separating.

Lemma 3.5.24, applied with $\gamma = \beta_1$, $\alpha = \beta_2$ and $\alpha' = \beta_5$ and $\delta = \beta_6$, it follows that $\beta_6 \subset B$, in particular it is disjoint from β_8 , since $\beta_8 \subset A$. The curves $\beta_6, \beta_7, \beta_8$ form a triangle, hence from the preliminary observation it follows that β_7 is 4-separating. A completely analogous argument proves that $\beta_4 \subset B$ and β_3 is 4-separating. It follows that the two distinct curves β_3 and β_7 are both 4-separating and disjoint, hence $\beta_3 = \beta_7$ which is a contradiction. The proof is now complete. \square

Now that we have taken care of the technical details we are ready to prove the rigidity of $\overline{G} \cdot \mathcal{O}^* \subset \mathcal{C}_{ss}(\Sigma_8)$.

Proof of Proposition 3.5.7. Let $i : \overline{G} \cdot \mathcal{O}^* \hookrightarrow \mathcal{C}_{ss}(\Sigma_8)$ be an edge-preserving injection. Since $\overline{G} \cdot \mathcal{O}$ is the union of copies of \mathcal{O} it follows from Proposition 3.5.23 that the restriction $i_{\overline{G} \cdot \mathcal{O}}$ is an edge-preserving embedding $\overline{G} \cdot \mathcal{O} \hookrightarrow \mathcal{C}_{(3)}(\Sigma_8)$. Then, thanks to Proposition 3.5.8, up to postcomposing with an extended mapping class $\varphi \in \text{Mod}^\pm(\Sigma_8)$ we can assume $\varphi \circ i_{\overline{G} \cdot \mathcal{O}} = \text{Id}_{\overline{G} \cdot \mathcal{O}}$.

By Definition 3.5.6 the curve ζ is the unique curve in $\mathcal{C}_{ss}(\Sigma_8)$ disjoint from $\gamma_1, \gamma_2, \gamma_5$ and γ_6 . It follows that for every j the curve $R_j \zeta$ is disjoint from $R_j \gamma_1, R_j \gamma_2, R_j \gamma_5$ and $R_j \gamma_6$. The curve $\varphi \circ i(R_j \zeta)$ (which is ζ itself for $j = 1, 2, 5, 6$) must be disjoint from $R_j \gamma_1, R_j \gamma_2, R_j \gamma_5$ and $R_j \gamma_6$: Lemma 3.5.5 proves that these curves admit a unique common neighbour, which is $R_j \zeta$ hence $\varphi \circ i(R_j \zeta) = R_j \zeta$. The same argument shows that for every j we have $\varphi \circ i(H_j \zeta) = H_j \zeta$. It follows that $\varphi \circ i = \text{Id}_{\overline{G} \cdot \mathcal{O}^*}$, hence $i = \varphi_{|\overline{G} \cdot \mathcal{O}^*}^{-1}$ which was our claim. \square

We can now conclude the proof of Theorem E for Σ_8 .

Proof of Theorem E for Σ_8 . We can now apply Proposition 3.2.4 with $\Gamma = \mathcal{C}_{ss}(\Sigma_8)$, $H = \overline{G}$, and $X = \overline{G} \cdot \mathcal{O}^*$.

Rigidity of X is Proposition 3.5.7. The triviality of stabilisers follows from the fact that every set of the form $g \cdot \mathcal{O}^* \cap \mathcal{O}^*$ contains \mathcal{O} and another curve $R_i \gamma_j$ such that

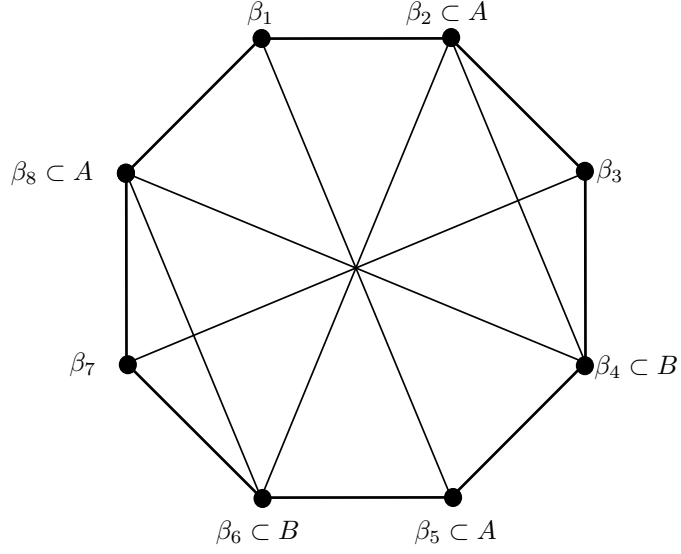


Figure 3.46: The curves $\beta_1, \beta_3, \beta_7$ are 4-separating.

$j \neq i - 1, i, i + 1, i + 4$, hence it is a superset of a set with trivial pointwise stabilisers due to Corollary 3.5.14, hence its stabiliser is itself trivial. This also implies that \mathcal{O}^* has trivial stabiliser, hence every subgraph of the exhaustion has trivial pointwise stabiliser due to Lemma 3.2.1.

Lastly, the graph \mathcal{O}^* contains both 3-separating and 4-separating curves, hence it intersects every orbit of curves under the action of $\text{Mod}(\Sigma_8)$. It follows that the requirement to have an exhaustion is fulfilled, too. \square

3.6 Examples of non-Exhaustable Graphs

In this section we will briefly present a couple of examples of graphs which are combinatorially rigid, and with nice metric properties, for which the co-Hopfian property as in Proposition 3.2.5 fails. We can hence deduce that those graphs do not admit any exhaustion by rigid sets with trivial stabilisers.

The graphs we will use will be graphs of curves and arcs on infinite-type surfaces (that is surfaces S such that the fundamental group $\pi_1(S)$ is not finitely generated). Indeed, under quite weak hypotheses, these graphs are combinatorially rigid, that is every automorphism is induced by an extended mapping class. Moreover, we can chose graphs such that the action of the extended mapping class group is transitive on both vertices and edges. Lastly, it is possible to choose examples which are metrically nice, in particular which have infinite diameter and are Gromov-hyperbolic (for a definition see [B1, §6.1], for instance).

The main property of infinite-type surfaces we will use is the fact that they admit proper subsurfaces which are homeomorphic to the entire surface, but not isotopic to it. This can be considered a failure of the co-Hopfian property for the category of infinite-type surfaces, so it should not be surprising that graphs arising from such objects fail to be co-Hopfian, as well.

The examples we are about to present can be adapted to a large number of infinite-type surfaces. We will, however, just restrict to the possibly simplest case, that is the one of a surface with no genus and a Cantor set of marked points. Such a surface will be called a *Cantor tree-surface* and is defined as follow.

Definition 3.6.1. A *Cantor set* is a totally disconnected compact metrisable space such that every point is an accumulation point (see [Mo, p.83]).

Let $C \subset S_{0,0}$ be a Cantor set on the 2-dimensional sphere. The surface (S, C) , that is a sphere with the Cantor set C as set of marked points, is called a *Cantor tree-surface*.

Before we move forward, let us notice that the embedding of a Cantor set into \mathbb{R}^2 , hence into the sphere $S_{0,0}$, is unique up to homeomorphisms. Indeed, [Mo, Theorem 8, § 12] proves that every two Cantor set are (abstractly) homeomorphic, while [Mo, Theorem 7, §13] stated that every homeomorphism between two Cantor sets embedded into \mathbb{R}^2 extends to a homeomorphism of the entire plane. This implies that the Cantor tree-surface is unique up to homeomorphism.

We will now define the graph for which we are going to show that the co-Hopfian property does not hold, as defined in [S-C, §5]

Definition 3.6.2. Let (S, C) be the Cantor tree-surface. Let $p \in S$ be point on $S \setminus X$. Let $\mathcal{A}(S, p)$ be the *loop graph* on S based at the point p , that is the graph such that:

Vertices There is a vertex for every isotopy class of simple parametrised arcs with both endpoints at p , and whose interior is disjoint from p , where the isotopy is relative to the endpoints;

Edges There is an edge between two arcs if they admit representatives only intersecting at p .

The following result holds.

Theorem 3.6.3. *Let (S, C) be the Cantor tree-surface. Let $p \in S \setminus C$. Let $\mathcal{A}(S, p)$ be the loop graph based at p . The following facts hold:*

1. *The graph $\mathcal{A}(S, p)$ is vertex-transitive, that is, the action of $\text{Aut}(\mathcal{A}(S, p))$ on vertices is transitive;*
2. *The graph $\mathcal{A}(S, p)$ is edge-transitive, that is, the action of $\text{Aut}(\mathcal{A}(S, p))$ on edges is transitive;*
3. *The graph $\mathcal{A}(S, p)$ is rigid, that is $\text{Aut}(\mathcal{A}(S, p)) \cong \text{Mod}^\pm(S)$;*
4. *The graph $\mathcal{A}(S, p)$ has infinite diameter and is Gromov-hyperbolic;*
5. *The graph $\mathcal{A}(S, p)$ does not admit an exhaustion by rigid sets with trivial pointwise stabilisers.*

Proof. Properties 1 and 2 follow from the fact that every complementary component of any arc on S is a disc with a Cantor set of marked points, and every two arcs only intersecting in p separate S into three discs with a Cantor set of marked points.

Property 3 is [S-C, Theorem 1].

Property 4 is [AFP, Theorem 1.1]. It is worth nothing (as also done by the authors) that the proof of hyperbolicity is essentially the same as the one for hyperbolicity of the curve graph on finite type surfaces given in [PS], as that proof did not use at any point the fact that surfaces are of finite type.

Property 5 follow from the fact that $\mathcal{A}(S, p)$ admits an embedding onto a proper subgraph: indeed, if $\mathcal{A}(S, p)$ admitted an exhaustion by rigid sets with trivial pointwise stabilisers then Proposition 3.2.5 would imply that every embedding of the graph is an automorphism. To see that $\mathcal{A}(S, p)$ admits an embedding onto a proper subgraph let γ be a curve on S disjoint from p . Let $\Gamma \subset \mathcal{A}(S, p)$ be the subgraph of loops disjoint from a . Let S' be the component of $S \setminus \gamma$ containing p . Then $\Gamma = \mathcal{A}(S', p)$. We have that S' is homeomorphic to S (we are abusing notation and treat the puncture of S' as a marked point). In particular it follows that the entire loop graph $\mathcal{A}(S, p)$ is isomorphic to $\mathcal{A}(S', p) = \Gamma$, which is a proper subgraph. \square

Another closely related example of a graph having an embedding onto a proper subgraph, hence not admitting an exhaustion by finite rigid set with trivial pointwise stabilisers, is the curve graph of the Cantor tree-surface $\mathcal{C}(S)$.

Since the surface S is of infinite-type, the curve graph $\mathcal{C}(S)$ has diameter 2, so in particular it is more trivial than the loop graph, at least from a metrical viewpoint. Apart from this, though, the other properties which we have highlighted for the loop graph in the previous theorem still hold for $\mathcal{C}(S)$.

Theorem 3.6.4. *Let (S, X) be the Cantor tree-surface. Let $\mathcal{C}(S)$ be the curve graph of S . The following facts hold:*

1. *The graph $\mathcal{C}(S)$ is vertex-transitive;*
2. *The graph $\mathcal{C}(S)$ is edge-transitive;*

3. *The graph $\mathcal{C}(S)$ is rigid, that is $\text{Aut}(\mathcal{C}(S)) \cong \text{Mod}^\pm(S)$;*
4. *The graph $\mathcal{C}(S)$ does not admit an exhaustion by rigid sets with trivial pointwise stabilisers.*

Proof. Properties 1,2, and 4 follow from analogous arguments as Properties 1,2, and 5 in the proof of Theorem 3.6.3.

Property 3 is [HMV, Theorem 2]. □

Bibliography

- [AFP] J. Aramayona, A. Fossas, H. Parlier. *Arc and curve graphs for infinite-type surfaces*. Proc. Amer. Math. Soc. 145 (2017), no. 11, 4995–5006.
- [AL1] J. Aramayona, C. J. Leininger. *Finite rigid sets in curve complexes*. J. Topol. Anal. 5 (2013), no. 2, 183–203.
- [AL2] J. Aramayona, C. J. Leininger. *Exhausting curve complexes by finite rigid sets*. Pacific J. Math. 282 (2016), no. 2, 257–283.
- [BHS] J. Behrstock, M. F. Hagen, A. Sisto. *Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups*. Geom. Topol. 21 (2017), no. 3, 1731–1804.
- [B1] B. H. Bowditch. *A course on geometric group theory*. Mathematical Society of Japan Memoirs, volume 16, 2006.
- [B2] B. H. Bowditch. *Rigidity of the strongly separating curve graph*. Michigan Math. J. 65 (2016), no. 4, 813–832.
- [B3] B. H. Bowditch. *Large-scale rank and rigidity of the Weil-Petersson metric*. Groups Geom. Dyn. 14 (2020), no. 2, 607–652.
- [BM1] T. E. Brendle, D. Margalit. *Commensurations of the Johnson kernel*. Geom. Topol. 8 (2004), 1361–1384.
- [BM2] T. E. Brendle, D. Margalit. *Normal subgroups of mapping class groups and the metaconjecture of Ivanov*. J. Amer. Math. Soc. 32 (2019), no. 4, 1009–1070.
- [BCM] J. F. Brock, R. D. Canary, Y. N. Minsky. *The classification of Kleinian surface groups, II: The ending lamination conjecture*. (English summary) Ann. of Math. (2) 176 (2012), no. 1, 1–149.
- [De] Max Dehn. *Papers on Group Theory and Topology*. Springer-Verlag, New York, 1987.
- [Di] Valentina Disarlo. *Combinatorial rigidity of arc complexes*. Preprint, arXiv:1505.08080, 2015.
- [Do] M. P. Do Carmo. *Differential geometry of curves and surface*. Prentice-Hall Inc., Englewood Cliffs, New Jersey, 1976.
- [FM] B. Farb, D. Margalit. *A primer on mapping class group*. Princeton Mathematical Series, volume 49, Princeton University Press, 2012.

- [Hal] N. Halpern. *A proof of the collar lemma*. Bull. London Math. Soc. 13 (1981), no. 2, 141–144.
- [Hara] F. Harary. *Graph theory*. Addison-Wesley Series in Mathematics, Addison-Wesley, Reading, 1969.
- [Hare1] J. L. Harer. *Stability of the homology of the mapping class groups of orientable surfaces*. Ann. of Math. (2) 121 (1985), no. 2, 215–249.
- [Hare2] J. L. Harer. *The virtual cohomological dimension of the mapping class group of an orientable surface*. Invent. Math. 84 (1986), no. 1, 157–176.
- [Harv1] W. J. Harvey. *Boundary structure of the modular group*. Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978), pp. 245–251, Ann. of Math. Stud., 97, Princeton Univ. Press, Princeton, N.J., 1981.
- [Hat] A. Hatcher. *Algebraic Topology*. Cambridge University Press, New York, 2002.
- [He] J. Hernández Hernández. *Exhaustion of the curve graph via rigid expansions*. Glasg. Math. J. 61 (2019), no. 1, 195–230.
- [HMV] J. Hernández Hernández, I. Morales, F. Valdez. *Isomorphisms between curve graphs of infinite-type surfaces are geometric*. Rocky Mountain J. Math. 48 (2018), no. 6, 1887–1904.
- [HLM] J. Hernández Hernández, C. J. Leininger, R. Maungchang. *Finite rigid subgraphs of pants graphs*. Geom. Dedicata 212 (2021), 205–223.
- [Hi] M. W. Hirsch. *Differential Topology*. Graduate Texts in Mathematics, volume 33, Springer, New York, 1976.
- [Ir] Elmas Irmak. *Complexes of nonseparating curves and mapping class groups*. Michigan Math. J., 54(1):81–110, 2006.
- [IrM] Elmas Irmak and John D. McCarthy. *Injective simplicial maps of the arc complex*. Turkish J. Math., 34(3):339–354, 2010.
- [Iv1] N. V. Ivanov. *Automorphisms of complexes of curves and of Teichmüller spaces*. Progress in knot theory and related topics, 113–120, Travaux en Cours, 56, Hermann, Paris, 1997.
- [Iv2] N. V. Ivanov. *Fifteen problems about the mapping class groups*. Problems on mapping class groups and related topics, 71–80, Proc. Sympos. Pure Math., 74, Amer. Math. Soc., Providence, RI, 2006.
- [J] M. C. Jordan. *Cours d’analyse de l’École Polytechnique*. Gauthier-Villars, Paris (1887), pp. 587–594.
- [Ke] L. Keen. *Collars on Riemann surfaces*. Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland, College Park, Md., 1973), pp. 263–268. Ann. of Math. Studies, No. 79, Princeton Univ. Press, Princeton, N.J., 1974.

- [Ki] Y. Kida. *Automorphisms of the Torelli complex and the complex of separating curves*. J. Math. Soc. Japan 63 (2011), no. 2, 363–417.
- [Kn] M. Kneser. *Aufgabe 360*. Jahresbericht der DMV (1955), 58(2): 27.
- [Kor] M. Korkmaz. *Automorphisms of complexes of curves on punctured spheres and on punctured tori*. Topology Appl. 95 (1999), no. 2, 85–111.
- [KP] M. Korkmaz, A. Papadopoulos. *On the arc and curve complex of a surface*. Math. Proc. Cambridge Philos. Soc., 148(3):473–483, 2010.
- [Kos] A. A. Kosinski. *Differential manifolds*. Pure and Applied Mathematics, volume 138, Academic Press Inc, San Diego, 1993.
- [Li] W. B. R. Lickorish. *A finite set of generators for the homeotopy group of a 2-manifold*. Proc. Cambridge Philos. Soc. 60 (1964), 769–778.
- [Lu] F. Luo. *Automorphisms of the complex of curves*. Topology 39 (2000), no. 2, 283–298.
- [Ma] G. E. Martin. *The foundations of geometry and the non-Euclidean plane*. Undergraduate Texts in Mathematics, Springer, New York, 1975.
- [MM1] H. A. Masur, Y. N. Minsky. *Geometry of the complex of curves. I. Hyperbolicity*, Invent. Math. 138 (1999), no. 1, 103–149.
- [MM2] H. A. Masur, Y. N. Minsky. *Geometry of the complex of curves. II. Hierarchical structure*. Geom. Funct. Anal. 10 (2000), no. 4, 902–974.
- [MP] J. D. McCarthy, A. Papadopoulos. *Simplicial actions of mapping class groups*. In Handbook of Teichmüller theory. Volume III, volume 17 of IRMA Lect. Math. Theor. Phys., pages 297–423. Eur. Math. Soc., Zürich, 2012.
- [Mc1] A. McLeay. *Normal subgroups of the braid group and the metaconjecture of Ivanov*. Preprint, arXiv:1801.05209, 2018.
- [Mc2] A. McLeay. *Geometric normal subgroups in mapping class groups of punctured surfaces*. New York J. Math. 25 (2019), 839–888.
- [Mo] E. E. Moise. *Geometric topology in dimensions 2 and 3*. Graduate text in Mathematics, volume 47, Springer, New York, 1977.
- [PS] P. Przytycki, A. Sisto. *A note on acylindrical hyperbolicity of mapping class groups*. (English summary) Hyperbolic geometry and geometric group theory, 255–264, Adv. Stud. Pure Math., 73, Math. Soc. Japan, Tokyo, 2017.
- [Pu] A. Putman. *A note on the connectivity of certain complexes associated to surfaces*. Enseign. Math. (2) 54 (2008), no. 3-4, 287–301.
- [R] T. Radó. *Über den Begriff der Riemannsche Fläche*. Acta Univ. Szeged 2 (1924-26) 101-121.

- [Sha] K. J. Shackleton. *Combinatorial rigidity in curve complexes and mapping class groups*. Pacific J. Math. 230 (2007), no. 1, 217–232.
- [Shi] E. Shinkle. *Finite rigid sets in arc complexes*. Algebr. Geom. Topol. 20 (2020), no. 6, 3127–3145.
- [Sp] E. H. Spanier. *Algebraic topology*. Springer-Verlag, New York, 1966.
- [S-C] A. Schaffer-Cohen. *Automorphisms of the loop and arc graph of an infinite-type surface*. preprint, arXiv:1912.06774, 2019.
- [SU] E. R. Scheinerman, D.H. Ullman. *Fractional graph theory. A rational approach to the theory of graphs*. Wiley, New York, 1997.
- [Tho] R. Thom. *Des variétés triangulées aux variétés différentiables*. 1960 Proc. Internat. Congress Math. 1958 pp. 248–255 Cambridge Univ. Press, New York.
- [Thu] W. P. Thurston. *Three dimensional geometry and topology*. Silvio Levy (ed), Princeton University Press, Princeton, 1997.
- [Wh] J. H. C. Whitehead. *On $C1$ -Complexes*. The Annals of Mathematics. Second Series. 41 (4) (October 1940): 809–824.
- [Wo] S. A. Wolpert. *Geodesic length functions and the Nielsen problem*. J. Differential Geom. 25 (1987), no. 2, 275–296.